

# Looking beyond Pixels: Theory, Algorithms and Applications of Continuous Sparse Recovery

THÈSE N° 8340 (2018)

PRÉSENTÉE LE 1<sup>ER</sup> JUIN 2018

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS  
LABORATOIRE DE COMMUNICATIONS AUDIOVISUELLES  
PROGRAMME DOCTORAL EN INFORMATIQUE ET COMMUNICATIONS

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Hanjie PAN

acceptée sur proposition du jury:

Prof. S. Süsstrunk, présidente du jury  
Prof. M. Vetterli, Prof. T. Blu, directeurs de thèse  
Dr J.-L. Starck, rapporteur  
Prof. A.-J. van der Veen, rapporteur  
Prof. D. Van De Ville, rapporteur



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE

Suisse  
2018



# Acknowledgments

I would like to express my greatest gratitudes to my two thesis advisors: Martin and Thierry, without whom I would never have embarked the exciting journey through the finite rate of innovation jungle. Martin and Thierry are probably the best advisors any PhD student could dream to have, despite their very different styles.

I am grateful to Thierry for introducing me to the finite rate of innovation world during my master study. Thanks for your mathematical genius, which saved me time and time again. You taught me to persevere and believe in my own research.

Martin, thanks for giving me the freedom to explore different research directions. I could not thank you enough for your patience and support to guide me through the adventure. Your guidance ensured that I never lose sight of the big picture and be aware of the potential stagnation over the research plateau. Thanks to your encouragement, I took an internship at IBM Research Zürich. The invaluable experience made me aware of how useful the theoretical work I developed in this thesis is for a practical problem.

I am thankful to my thesis committee, Jean-Luc Starck, Alle-Jan van der Veen, Dimitri Van De Ville and the jury president Sabine Süsstrunk, for reading my thesis thoroughly and giving constructive feedbacks.

I would like to thank the real bosses of the lab: Jacqueline, Heidi and Nicoletta. Your proficiency ensured that I can focus on the research side without getting stressed out by administrative issues.

During my LCAV years, I am fortunate to interact with friends in and outside LCAV. Robin, thank you for the joint work on the (yet to be finished) acoustic direction of arrival estimation, for trips and for your enlightening Python advices. Thank you Eric, for your passion and hard work on refactoring my messy code for the direction of arrival estimation and the collaboration on ultrasound imaging. Thank you Gilles, my officemate, for sharing your passion on skiing and lending a hand on the final touches of illustrations in papers and slides. Thanks Mihailo for numerous ski trips and your black humor. Thank you Adam, Michalina and Dalia for sharing lunches across the campus. Thank you Zhou and Niranjana for sharing with me your inspiring and wild dreams on all sorts of applications related to computational imaging and your passion for coffee. Thanks to everyone who has made this voyage awesome: Alireza, Benjamín, Frederike, Golnoosh, Jay, Lionel, Loïc, Luc, Marta, Miranda, Mitra, Paolo, Reza and Runwei.

I am extremely lucky to be able to work with Matthieu, Sepand, Paul, Mervé and Lucien during my internship at IBM. Thank you all to bridge the gap between the theoretical work and a complicated real world application, like radio astronomy source estimation.

I am grateful to friends I met in Lausanne: Xifan, Jingjing, Tian, Ruofan, Chun-Lam, Su, Fan, Kaicheng, Xinchao, Xiaolu, Bin, Hao, Cong, Kévin, Rhicheck and Saeid. Thank you all for extending my life in Lausanne to an additional dimension than research. The world would be plainer without you.

Last but not the least, I would like to thank my parents, Xinying and Dachang for their

---

unconditional love and support, for nurturing my scientific curiosity early on in my childhood. You provided everything I needed to reach the current stage and beyond. This thesis is dedicated to them.

# Abstract

Sparse recovery is a powerful tool that plays a central role in many applications, including source estimation in radio astronomy, direction of arrival estimation in acoustics or radar, super-resolution microscopy, and X-ray crystallography. Conventional approaches usually resort to discretization, where the sparse signals are estimated on a pre-defined grid. However, the sparse signals do not line up conveniently on any grid in reality. While the discrete setup usually leads to a simple optimization problem that can be solved with standard tools, e.g.,  $\ell_1$ -based sparse recovery, there are two noticeable drawbacks as a consequence of discretization: (i) Because of the model mismatch, the effective noise level in the measurements is increased; (ii) The minimum resolution a discrete algorithm can reach is limited by the artificially imposed grid step-size, which may not be sufficient in an imaging system, e.g., microscopy. Because of the limitations of the conventional grid-based methods, it is essential to develop a technique that estimates sparse signals in the *continuous-domain* — in essence looking beyond pixels.

The aims of this thesis are (i) to further develop a continuous-domain sparse recovery framework based on finite rate of innovation (FRI) sampling on both theoretical and algorithmic aspects; (ii) adapt the proposed generic sparse recovery technique to several applications, namely radio astronomy point source estimation, direction of arrival estimation in acoustics, and single image up-sampling; (iii) show that the continuous-domain sparse recovery approach can surpass the instrument resolution limit and achieve super-resolution.

We propose a continuous-domain sparse recovery technique by generalizing the FRI sampling framework to cases with *non-uniform* measurements. We achieve this by identifying a set of unknown uniform sinusoidal samples (which are related to the sparse signal parameters to be estimated) and the linear transformation that links the uniform samples of sinusoids to the measurements. It is shown that the continuous-domain sparsity constraint can be equivalently enforced with a discrete convolution equation of these sinusoidal samples. Then, the sparse signal is reconstructed by minimizing the fitting error between the given and the re-synthesized measurements (based on the estimated sparse signal parameters) subject to the sparsity constraint. Further, we develop a *multi-dimensional* sampling framework for Diracs in two or higher dimensions with *linear* sample complexity. This is a significant improvement over previous methods, which have a complexity that increases exponentially with space dimension. An efficient algorithm has been proposed to find a valid solution to the continuous-domain sparse recovery problem such that the reconstruction (i) satisfies the sparsity constraint; and (ii) fits the given measurements (up to the noise level).

We validate the flexibility and robustness of the FRI-based continuous-domain sparse recovery in both simulations and experiments with real data. Specifically, we show that the proposed method surpasses the diffraction limit of radio telescopes, which is inversely proportional to the maximum baseline, with both realistic simulation and real data from the LOFAR radio telescope. In addition, FRI-based sparse reconstruction requires fewer measurements and smaller baselines to reach a similar reconstruction quality compared with conventional methods. Next, we apply

the proposed approach to direction of arrival estimation in acoustics. We show that accurate off-grid source locations can be reliably estimated from microphone measurements with arbitrary array geometries in both 2D and 3D setups. Finally, we demonstrate the effectiveness of the continuous-domain sparsity constraint in regularizing an otherwise ill-posed inverse problem, namely single-image super-resolution. By incorporating image edge models, which are estimated from the given low-resolution image, the up-sampled image retains sharp edges and is free from ringing artifacts.

**Keywords:** Finite rate of innovation (FRI), approximation, sparse reconstruction, irregular sampling, continuous-domain sparsity, multi-dimension, radio interferometry, point source reconstruction, direction of arrival, edge modeling, curve annihilation, image upsampling, alternating direction method of multipliers (ADMM).

# Résumé

La reconstruction parcimonieuse (anglais: sparse recovery) est un outil puissant jouant un rôle central dans plusieurs applications comme l'estimation de sources ponctuelles en radio-astronomie, l'estimation de la direction d'arrivée en acoustique ou en radar, la microscopie à super-résolution, ainsi que la cristallographie par rayons X. Les techniques conventionnelles consistent généralement par une discrétisation du problème pour reconstruire la solution parcimonieuse sur une grille prédéfinie. Ce dernier prend généralement la forme d'un problème d'optimisation simple qui peut être résolu par des méthodes standards telle la minimisation en norme  $\ell_1$ . Néanmoins, la formulation discrète n'est pas sans conséquences: D'une part, la discrétisation du problème induit artificiellement une élévation du niveau du bruit. D'autre part, la reconstruction parcimonieuse est limitée à la résolution de la grille, alors qu'en pratique de tels signaux ne tombent jamais exactement sur ces-derniers. Pour des applications en microscopie par exemple, cette résolution artificielle est insuffisante. En vue de ces limitations, il est essentiel d'étudier des procédures qui retrouvent la solution parcimonieuse à l'échelle continue — pour voir au-delà des pixels.

Les objectifs de cette thèse sont (i) de poursuivre le développement de la reconstruction parcimonieuse à l'échelle continue, basée sur l'échantillonnage par taux d'innovation fini (anglais: FRI) du point de vue théorique et algorithmique; (ii) d'adapter ce modèle générique de reconstruction parcimonieuse à divers applications, notamment l'estimation de sources ponctuelles en radio-astronomie, l'estimation de la direction d'arrivée en acoustique, et du suréchantillonnage par image unique; (iii) démontrer que la reconstruction parcimonieuse à l'échelle continue permet de surpasser les limites de l'instrument par super-résolution.

Nous proposons une technique de reconstruction parcimonieuse à l'échelle continue en généralisant le cadre d'échantillonnage par taux d'innovation fini au cas où les mesures sont non-uniformes. Nous réalisons cela en identifiant un ensemble de mesures de forme sinusoïdale qui sont liées aux mesures initiales, ainsi que la transformation linéaire qui les lie à ces derniers. Nous montrons que la contrainte de parcimonie à l'échelle continue peut être représentée par une convolution discrète des mesures de forme sinusoïdale. Ensuite, le signal parcimonieux est estimé en minimisant l'erreur entre les mesures réelles et celles induites par les paramètres du signal parcimonieux, en respectant la contrainte de parcimonie. De plus, nous étendons cette procédure à l'échantillonnage multi-dimensionnelle de Diracs ayant une complexité qui croît linéairement avec le nombre d'échantillons. Cela donne une amélioration significative par rapport aux méthodes ultérieures dont la complexité augmente de manière exponentielle avec le nombre de dimensions. Un algorithme efficace est proposé pour trouver une solution au problème de reconstruction parcimonieuse de signaux à l'échelle continue de sorte que (i) la solution respecte la condition de parcimonie et (ii) coïncide avec les mesures (à hauteur du niveau du bruit).

Nous vérifions la flexibilité et robustesse de notre méthode de reconstruction parcimonieuse à l'échelle continue en simulation et sur de vraies données. En particulier nous montrons, en simulation et sur de vraies données issues du télescope LOFAR, que notre méthode permet de dépasser la limite de diffraction des radio-télescopes, ce dernier étant inversement proportionnelle

---

à la distance maximale séparant les antennes. De plus, notre méthode de reconstruction parcimonieuse peut se contenter de moins de mesures et d'un instrument de plus petite taille pour obtenir des estimées équivalentes aux méthodes conventionnelles. Ensuite, nous appliquons notre méthode au problème d'estimation de la direction d'arrivée en acoustique. Nos résultats montrent que la direction de sources non-alignées avec une grille discrète peuvent être déterminées précisément, sans poser de contraintes sur la géométrie des microphones en 2D ou 3D. Finalement, la contrainte de parcimonie à l'échelle continue peut être utilisée comme méthode de régularisation pour le problème de super-résolution par image unique. En formant un modèle des bordures d'une image à partir d'une version à faible résolution de cette dernière, cette information peut être utilisée dans l'image à haute résolution pour éliminer les effets de bords.

**Mots-Clés:** Taux d'innovation fini, approximation, reconstruction parcimonieuse, échantillonnage non-uniforme, parcimonie à l'échelle continue, signaux multi-dimensionnels, interférométrie par ondes radio, estimation de sources ponctuelles, modélisation de bords, suréchantillonnage d'images.



# Contents

<b>Acknowledgments</b>	<b>iii</b>
<b>Abstract</b>	<b>v</b>
<b>Résumé</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.1.1 Source estimation in radio astronomy . . . . .	2
1.1.2 Direction of arrival estimation with microphone arrays . . . . .	3
1.1.3 Image up-sampling . . . . .	4
1.2 Methods . . . . .	4
1.2.1 Sensing sparsity through annihilation . . . . .	4
1.2.2 A generic FRI reconstruction algorithm . . . . .	6
1.3 Thesis outline and main contributions . . . . .	7
<b>2 Continuous Sparse Recovery with Finite Rate of Innovation Sampling</b>	<b>13</b>
2.1 Introduction . . . . .	13
2.2 The classic FRI reconstruction framework . . . . .	16
2.3 A generic FRI reconstruction algorithm . . . . .	17
2.3.1 FRI reconstruction as a constrained optimization . . . . .	17
2.3.2 Essential ingredients . . . . .	20
2.3.3 An iterative algorithm to solve (P2) . . . . .	22
2.4 Four algorithmic examples for various types of FRI signals . . . . .	24
2.4.1 Stream of Diracs with irregular time domain samples . . . . .	24
2.4.2 Weighted sum of Diracs with irregular Fourier domain samples . . . . .	26
2.4.3 FRI curves . . . . .	28
2.4.4 Diracs on the sphere . . . . .	30
2.5 Application to radio interferometry . . . . .	32
2.5.1 Data acquisition and signal model . . . . .	32
2.5.2 Reconstruction of point sources . . . . .	33
2.5.3 Simulation results . . . . .	35
2.6 Conclusion . . . . .	35
2.A Derivation of the equivalent formulation of (P1) . . . . .	36
2.B Equivalent form for the solution of (2.3) . . . . .	37
2.C Exact interpolation with Dirichlet kernel . . . . .	37

2.D	Spherical harmonics of Diracs on the sphere . . . . .	38
2.E	Coefficients of the canonical basis in (2.13) . . . . .	39
<b>3</b>	<b>Multi-dimensional FRI Sampling with Linear Sample Complexity</b>	<b>41</b>
3.1	Introduction . . . . .	41
3.2	Methods . . . . .	43
3.2.1	FRI reconstruction in 2D . . . . .	43
3.2.2	Algorithm and implementation . . . . .	45
3.2.3	Generalization to higher dimensions . . . . .	51
3.3	Results . . . . .	54
3.3.1	Efficient data utilization . . . . .	54
3.3.2	Robust reconstruction . . . . .	55
3.3.3	Dirac with common $x/y$ coordinates . . . . .	55
3.3.4	Dirac reconstruction in 3D . . . . .	56
3.4	Conclusion . . . . .	57
3.A	Support size of a 2D annihilating filter . . . . .	58
3.B	Sample complexity of Dirac reconstruction in 2D . . . . .	58
3.C	2D Diracs with shared $x$ or $y$ coordinates . . . . .	58
3.D	Null space dimension of the joint annihilation matrix . . . . .	59
3.E	Equivalent formulation of the joint annihilation (3.2) . . . . .	59
<b>4</b>	<b>Looking beyond Pixels with Continuous-space Estimation of Point Sources</b>	<b>61</b>
4.1	Introduction . . . . .	61
4.2	Interferometric imaging measurement equation . . . . .	63
4.3	Continuous domain sparse recovery with FRI sampling . . . . .	65
4.4	Algorithm . . . . .	65
4.4.1	Estimation of point source locations . . . . .	66
4.4.2	Estimation of point source intensities . . . . .	67
4.4.3	Coherent multiband reconstruction . . . . .	68
4.4.4	Update strategy for the linear mapping $G$ in (4.6) . . . . .	69
4.5	Results . . . . .	72
4.5.1	Data and experiment setup . . . . .	72
4.5.2	Iterative refinement of linear mapping . . . . .	72
4.5.3	Source resolution . . . . .	73
4.5.4	Model order selection . . . . .	76
4.5.5	Actual LOFAR observation: Boötes Field . . . . .	77
4.5.6	Actual LOFAR observation: Toothbrush Cluster . . . . .	80
4.6	Discussion . . . . .	81
4.6.1	Resolution . . . . .	81
4.6.2	Efficient use of visibility measurements . . . . .	81
4.7	Conclusion . . . . .	82
<b>5</b>	<b>Direction of Arrival Estimation with Microphone Arrays</b>	<b>83</b>
5.1	Introduction . . . . .	83
5.2	Signal and measurement models . . . . .	85
5.3	The acoustic direction of arrival problem . . . . .	86
5.3.1	DOA estimation from cross-correlations in 2D . . . . .	86

<b>Contents</b>	<b>xi</b>
<hr/>	
5.3.2 DOA estimation from cross-correlations in 3D . . . . .	87
5.3.3 DOA estimation from raw microphone signals . . . . .	88
5.4 Results . . . . .	89
5.4.1 Influence of noise . . . . .	89
5.4.2 Resolving close sources . . . . .	90
5.4.3 Results with actual speech recordings . . . . .	91
5.5 Conclusion . . . . .	94
<b>6 Image up-sampling via Edge Etching</b>	<b>95</b>
6.1 Introduction . . . . .	95
6.2 Annihilating curves with a mask function . . . . .	97
6.3 Annihilation-driven edge modeling . . . . .	98
6.3.1 Block-wise linear edge model . . . . .	99
6.3.2 A locally adaptive global mask function . . . . .	101
6.4 Application on edge-preserving image up-sampling . . . . .	103
6.4.1 Problem formulation . . . . .	103
6.4.2 Solve (6.13) with ADMM algorithm . . . . .	104
6.5 Experimental results . . . . .	107
6.5.1 Edge modelling . . . . .	108
6.5.2 Single image up-sampling . . . . .	108
6.6 Conclusion . . . . .	113
6.A Efficient implementation of block-wise edge modeling . . . . .	115
6.B Implementation of (6.18) in the DFT domain . . . . .	115
6.C Solve the noisy up-sampling (6.21) with ADMM . . . . .	116
<b>7 Conclusion</b>	<b>119</b>
<b>Bibliography</b>	<b>121</b>
<b>Curriculum Vitæ</b>	<b>133</b>



# Chapter 1

## Introduction

路漫漫其修远兮，吾将上下而求索。

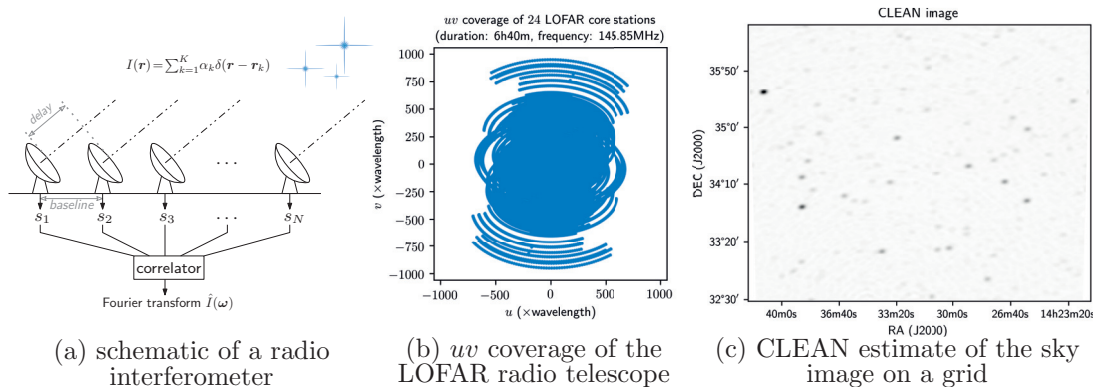
For on and on the road stretches far, where I'll  
search high and low.

---

离骚 *Li Sao*  
屈原 QU, YUAN

### 1.1 Motivation

Sampling plays a pivotal role in establishing the connection between continuous-domain signals and the available discrete measurements. An essential question when dealing with discrete samples is: Are these samples a faithful representation of the original continuous-domain signal that we start with? For bandlimited signals, the answer is given by the classic Shannon sampling theorem [114, 134]. But not all signals are bandlimited. In [3, 15, 128], a generalization of Shannon sampling is proposed for signals, which live in shift-invariant spaces but may not be bandlimited. Finite rate of innovation sampling (FRI) [14, 133] extends the classic sampling framework to signals that are neither bandlimited nor belong to a shift-invariant space as is the case in many practical problems. The focus of this thesis is the further development of FRI sampling on theoretical, algorithmic and applicative aspects. We would like to investigate whether a continuous-domain sparse recovery framework, like FRI, may provide an alternative which surpasses the limitations due to the discretization in conventional approaches. Before diving into the theory and algorithms, let us first give a brief description of a few applications as a motivation.



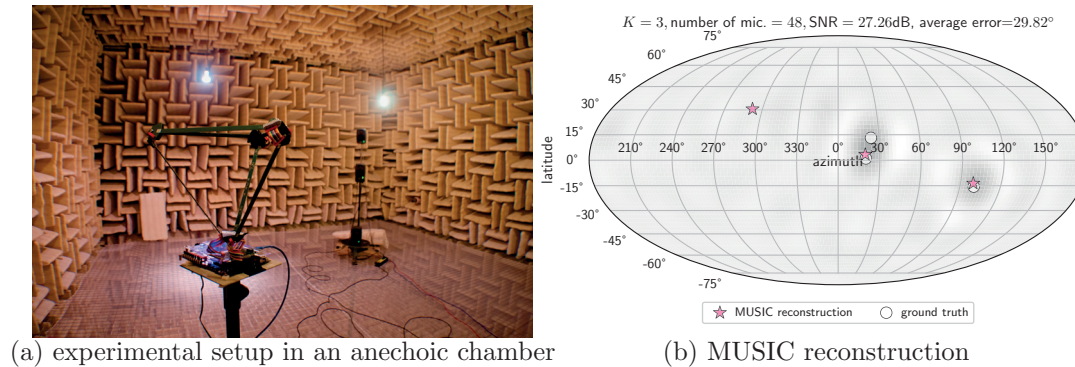
**Figure 1.1:** Source estimation with a radio interferometer. Cross-correlations of the received signals at different antennas corresponds to irregular sampling in the Fourier domain (a.k.a. the  $uv$ -domain in radio astronomy). The challenge is to accurately estimate closely located sources that are otherwise indistinguishable with a conventional technique, e.g., CLEAN.

### 1.1.1 Source estimation in radio astronomy

Source estimation in radio astronomy aims at reconstructing celestial sources from the electromagnetic waves collected with antenna arrays. For example, the Low Frequency Array (LOFAR) consists of around 20 000 dipole antennas, which are located primarily in Netherlands in addition to a few remote sites across Europe. One commonly used source model in radioastronomy is the point source model [66], which represents the sky image as a weighted sum of Diracs. In a far-field context, all sources are assumed to be located on a hypothetical celestial sphere (as the depth information is lost in the far-field setting). Consequently, the electromagnetic waves arrive as parallel wavefronts at the antenna array. The cross-correlations of the received signals at two antennas are determined by the separation of the antenna pair (i.e., the baseline) as well as the point source locations. In narrow field-of-view cases, the source estimation on the sphere can be approximated by projecting the sky image in the neighborhood of the telescope focus to the tangential plane. With such an approximation, it is shown that the cross-correlations sample the Fourier transform of the two-dimensional sky image projection at irregular frequencies (see [126], Chapter 3), which are specified by the radio telescope antenna layout. To summarize, the far-field point source reconstruction within a narrow field-of-view amounts to estimating 2D Dirac locations from irregularly sampled Fourier transform of the sky image.

Conventional methods, like CLEAN (and its many variants), reconstruct the sky on a discrete grid (see e.g., Figure 1.1(c)). An additional blob detection algorithm needs to be applied in order to determine point source locations (and intensities). The blob size in the CLEAN image determines the resolution power and is typically related to the maximum baselines of the radio telescope. For instance, with the configuration in Figure 1.1, the telescope angular resolution is  $3'23.1''$ . The challenge is to estimate point sources, which may be closely located within a small portion of the sky that is beyond the instrument resolution.

Alternatively, one may estimate the point sources directly in the continuous-domain from the cross-correlations of the antenna signals without first reconstructing an intermediate sky image on a grid. Thanks to the continuous-domain formulation, we may surpass the instrument resolution limit and estimate point sources without the grid imposition.

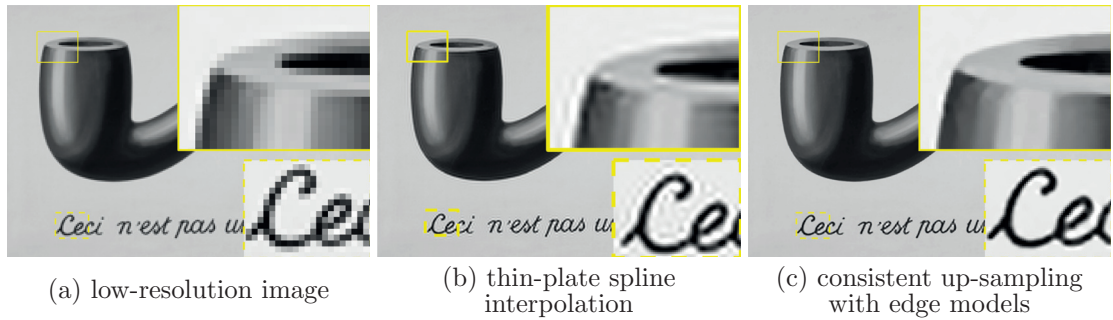


**Figure 1.2:** Direction of arrival estimation of acoustic sources from actual recordings in an anechoic chamber (estimated SNR: 27.26 dB. The noise level is estimated from the silence recording). (a) The Pyramic array used to collect sound waves emitted by loudspeakers located at different locations. (b) The least square estimate of the sound field overlaid with the MUSIC reconstruction. Even with a dense grid (20 000 approximately equidistant grid points on the sphere [71] with median inter-point distance of  $1.53^\circ$ ), MUSIC fails to resolve one of the two closely located sources (great circle distance:  $12.76^\circ$ ).

### 1.1.2 Direction of arrival estimation with microphone arrays

The same principle, which was used in radio astronomy point source reconstruction, applies to the direction of arrival (DOA) estimation in acoustics. The goal is to locate the direction of arrival of sound sources in 3D with microphone arrays. In the far-field context, this boils down to identifying the azimuths and latitudes of the acoustic sources. However, unlike the radio astronomy problem above, which typically considers a small portion of the sky around the telescope focus, the spherical reconstruction can no longer be approximated within a tangential plane for the DOA problem—sources can be located in a much wider range on the sphere. The cross-correlations of the received signals at different microphones are related to the Fourier domain samples at non-uniform frequencies, where the Fourier expansion basis are specified by the spherical harmonics (see e.g., [42, 45, 96]).

Figure 1.2(a) shows a typical setup of the acoustic DOA problem: A tetrahedron-shaped microphone array (the “Pyramic array”) is used to record speech signals from loudspeakers that are located at three different locations in an anechoic chamber. Each side of the tetrahedron consists of 8 MEMS microphones [7]. Conventional methods, e.g., MUSIC [111], then rely on grid-search to localize sources. Similarly to the radio astronomy cases, discretized methods may not always be sufficient to resolve closely located sources even with very dense grid. In Figure 1.2(b), we show such an example of DOA estimation, where two of the three loudspeakers are closely located (great-circle distance:  $12.76^\circ$ ). In the least square estimate of the sound field, both sources are contained within one bright spot. MUSIC fails to resolve one of the two closely located sources even with a dense grid, which consists of 20 000 approximately equidistant grid points on the sphere [71] with median inter-point distance of  $1.53^\circ$ . This motivated us to investigate whether a continuous-domain sparse recovery framework, like FRI, is able to accurately resolve closely located sources.



**Figure 1.3:** Super-resolution from a single image (i.e., image up-sampling). The challenge is to estimate an image with a higher resolution that is consistent with the given measurements and retains sharp edges.

### 1.1.3 Image up-sampling

In addition to the exact enforcement of the sparsity constraint as in the previous two examples, promoting signal sparsity may be used as a regularizer in otherwise severely ill-conditioned inverse problems. One such example is single image “super-resolution” (i.e., image up-sampling). A single discrete image is given and the goal is to estimate a higher resolution one that is consistent with the given measurements. There are infinitely many possible high resolution images that agree with the given low-resolution measurements. One way to better condition the ill-posed problem is to resort to regularization. For instance, one may minimize the  $\ell_2$  norm of the up-sampled image subject to the consistency constraint (which requires that the samples of the high-resolution image match that of the low-resolution one). The solution to the constrained optimization is given by thin-plate spline interpolation. While the method is simple, the resulting high-resolution image suffers from blurred edges and ringing artifacts (see e.g., Figure 1.3(b)). Note that image edges (i.e., the sparse signal to be reconstructed) are relatively localized: if one can estimate an edge model first and regularize the restoration with the model, then the up-sampled image may retain sharper edges (Figure 1.3(c)).

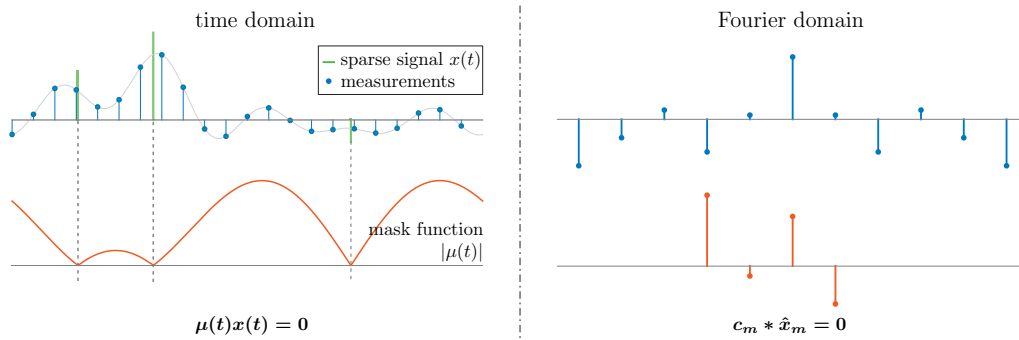
## 1.2 Methods

In this section, we review the FRI-based sparse recovery framework, which is at the core of this thesis. First, we highlight the methodology that FRI adopts: a sparse signal is recovered by finding a complementary mask function that “kills” or annihilates the sparse signal (Section 1.2.1). Then, we summarize a generic algorithm developed in Chapter 2 and Chapter 3 that recovers FRI signals robustly (Section 1.2.2).

### 1.2.1 Sensing sparsity through annihilation

Finite rate of innovation (FRI) sampling is a framework to reconstruct continuous-domain sparse signals, which have finite degrees of freedom per unit of time or space [133]. A typical FRI signal





**Figure 1.4:** Sensing sparsity through annihilation: instead of recovering the sparse signal  $x(t)$  directly, the FRI-based approach estimates a smooth function  $\mu(t)$ , which vanishes at the non-zero locations of the sparse signal. The continuous domain constraint  $\mu(t)x(t) = 0$ , can be equivalently enforced with a discrete convolution equation in Fourier domain:  $c_m * \hat{x}_m = 0$ , where  $c_m$  is known as the “annihilating filter” (see text in Section 1.2.1 for details).

is a  $\tau$ -periodic stream of Diracs:

$$x(t) = \sum_{k' \in \mathbb{Z}} \sum_{k=1}^K \alpha_k \delta(t - t_k - k'\tau). \quad (1.1)$$

The goal is to estimate the Dirac locations  $t_k$  and amplitudes  $\alpha_k$  from the measurements of the sparse signal  $x(t)$ , e.g., its ideally low-pass filtered samples. Instead of directly estimating where the sparse signal is different from zero, the FRI-based approach aims at reconstructing a smooth function  $\mu(t)$  (typically a polynomial), which vanishes at the non-zero locations of the sparse signal (see Figure 1.4). From this perspective,  $\mu(t)$  serves as a mask that *annihilates* the sparse signal  $x(t)$ :

$$\mu(t)x(t) = 0. \quad (1.2)$$

The zero-crossings of the function  $\mu(t)$  then give the Dirac locations. Once the Dirac locations have been reconstructed, estimating the amplitudes of the Diracs is a linear problem, namely a least square minimization.

In general, it is challenging to enforce the continuous domain annihilation constraint (1.2) directly. However, many FRI signals are or can be transformed into a weighted sum of sinusoids, whose frequencies are related to the sparse signal parameters. Thanks to a result known for more than two centuries [103], the continuous domain annihilation constraint can equivalently be enforced with a discrete convolution equation: there exists a finite length discrete filter such that its convolution with the uniform sinusoidal samples are zero (hence the filter is also known as “annihilating filter”). In fact, the polynomial, whose coefficients are specified by the discrete filter (e.g.,  $\mu(t)$  in (1.4)), is the mask function that annihilates the sparse signal.

More concretely, take the periodic stream of Diracs (1.1) as an example. Its Fourier series

coefficients are

$$\hat{x}_m = \frac{1}{\tau} \sum_{k=1}^K \alpha_k e^{-j \frac{2\pi m}{\tau} t_k},$$

which are annihilated (i.e.,  $c_m * \hat{x}_m = 0$ ) by a  $(K + 1)$ -tap filter  $[c_0, \dots, c_{K+1}]$  with  $z$ -transform

$$C(z) = \sum_{k=0}^{K+1} c_k z^{-k} = c_0 \prod_{k=1}^K \left(1 - e^{-j \frac{2\pi}{\tau} t_k} z^{-1}\right). \quad (1.3)$$

Note that  $C(z)$  is a polynomial of degree  $K + 1$  and the mask function  $\mu(t) = C(e^{-j \frac{2\pi}{\tau} t})$  vanishes precisely at the Dirac location  $\mu(t_k) = 0$  for  $k = 1, \dots, K$ :

$$\mu(t)x(t) \xleftrightarrow{\mathcal{F}} c_m * \hat{x}_m. \quad (1.4)$$

The ideally low-pass filtered samples of the sparse signal  $x(t)$  have a one-to-one correspondence with the Fourier series coefficients  $\hat{x}_m$ . Consequently, the annihilating filter coefficients (or equivalently the mask function) are reconstructed by enforcing a discrete convolution equation. The Dirac locations are then obtained by taking the roots of the polynomial (1.3). It has been shown that a stream of  $K$  Dirac deltas can be perfectly recovered from at least  $2K + 1$  ideal (noise-free) samples [133]. We illustrate the annihilating filter based sparse recovery in Figure 1.4.

### 1.2.2 A generic FRI reconstruction algorithm

In Chapter 2, a generic FRI reconstruction problem is formulated as a constrained optimization. The fitting error between the re-synthesized measurements (based on the estimated sparse signal model) and the given measurements is minimized subject to the annihilation constraint:

$$\begin{aligned} \min_{\mathbf{c} \in \mathcal{C}, \mathbf{b}} \quad & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to} \quad & \mathbf{c} * \mathbf{b} = \mathbf{0}. \end{aligned} \quad (1.5)$$

Therefore, the goal is to approximate the measurements optimally (in the least square sense) given that the signal follows the sparsity model. Here

- $\mathbf{a}$  is the vector of the sparse signal measurements to be reconstructed, e.g., the low-pass filtered samples, or the cross-correlations of the received signals with sensor arrays.
- $\mathbf{b}$  is the vector of unknown uniform sinusoidal samples, e.g., the Fourier series coefficients of (1.1); or the Fourier transform of the sky image on a uniform grid as in radio astronomy source estimation. The choice of  $\mathbf{b}$  is problem-dependent. By imposing the discrete convolution equation in (1.5), the estimated signal is guaranteed to follow the sparse signal model.
- $\mathbf{G}$  is the exact or approximate linear mapping from the uniform sinusoidal samples  $\mathbf{b}$  to the measurements  $\mathbf{a}$ , and typically has full column rank. In fact, an inexact mapping is *not* detrimental to the final reconstruction quality [96, 97] although, in general, we should not expect to achieve perfect reconstruction with noiseless measurements, when an approximate linear mapping is used. However, see Chapter 4 (Section 4.4.4) for a strategy to obtain the exact solution by refining the linear mapping based on the previous reconstructions.

- $\mathbf{c}$  is the vector of annihilating filter coefficients. The feasible set  $\mathcal{C}$  specifies a proper normalization in order to avoid a trivial solution  $\mathbf{c} \equiv \mathbf{0}$ . One choice of the feasible set  $\mathcal{C}$ , which leads to robust performance (see Chapter 2), is  $\mathcal{C} = \{\mathbf{c} | \mathbf{c}_0^H \mathbf{c} = 1\}$ , where  $\mathbf{c}_0$  is a random initialization of the annihilating filter.

Because of the annihilation constraint, (1.5) is non-convex with respect to the  $(\mathbf{b}, \mathbf{c})$ -pair. An iterative strategy finds a valid solution, which (i) satisfies the annihilation constraint; and (ii) has a fitting error  $\|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2$  within the noise level. This strategy exploits the “bi-linearity” of the annihilation constraint: for a fixed  $\mathbf{c}$  (respectively  $\mathbf{b}$ ), the annihilation constraint is linear in terms of  $\mathbf{b}$  (respectively  $\mathbf{c}$ ). Hence, with a fixed  $\mathbf{c}$ , we have a closed-form solution  $\mathbf{b}$  as a function of  $\mathbf{c}$ . Then, the original bi-variate optimization (1.5) reduces to a minimization with respect to  $\mathbf{c}$  alone. The simplified formulation inspired the iterative algorithms in Chapter 2 and Chapter 3, which can be efficiently implemented: At each iteration, the updates of  $\mathbf{b}$  and  $\mathbf{c}$  amount to solving two linear systems of equations. Please refer to Chapter 2 and Chapter 3 for detailed discussions.

Any sparse reconstruction problem can be solved with the FRI-based continuous sparse recovery method as soon as two key elements are identified: (i) the (unknown) uniform sinusoidal samples; and (ii) the linear mapping from these sinusoidal samples to the given measurements of the sparse signal. Take the point source reconstruction in radio astronomy (Section 1.1.1) as an example. One possible choice of the uniform samples of sinusoids could be the Fourier transform of the (continuous) sky image on a uniform grid. These uniform samples are related to the cross-correlations of the received antenna signals (i.e., the irregularly sampled Fourier transform of the sky image) via interpolation (see details in Chapter 4).

### 1.3 Thesis outline and main contributions

Motivated by several important applications, our focus in this thesis is to generalize the continuous-domain sparse recovery framework based on finite rate of innovation (FRI) sampling. We extend FRI sampling to cases with *non-uniform* samples, thus making the technique applicable to many practical problems (see Chapter 4 and Chapter 5 for concrete examples). Additionally, we propose an efficient *multi-dimensional* FRI sampling scheme, which has a linear sample complexity in two or higher dimensions. This is a significant improvement over previous methods, which required a minimum number of samples that scales exponentially in the space dimension.

The thesis consists of two integral parts: in the first part (Chapter 2 and Chapter 3) we develop the theory and algorithms for continuous-domain sparse recovery; in the second part (Chapter 4 to Chapter 6), we adapt the generic tool developed in the first part to three practical problems, namely point source reconstruction in radio astronomy, direction of arrival estimation problem in acoustics, and image up-sampling. We summarize each chapter and its respective contributions as follows.

#### Chapter 2 — Continuous Sparse Recovery with Finite Rate of Innovation Sampling

In this chapter, we study the generic continuous-domain sparse recovery problem. This chapter serves as the algorithmic foundation for both the further development of an efficient multidimensional sparse recovery approach in the next chapter as well as the subsequent application chapters.

The earliest occurrence of the classic continuous-domain sparse recovery is probably the estimation of sinusoids in time series using Prony’s method [103]. This is at the root of high

resolution spectral estimation (see [118] Chapter 4). The estimation of continuous-time sparse signals from discrete-time samples is the goal of the sampling theory for finite rate of innovation (FRI) signals [133]. Both spectral estimation and FRI sampling usually assume uniform sampling.

However, not all measurements are obtained uniformly, as exemplified by a concrete radioastronomy problem we set out to solve (see Chapter 4). Thus, we develop the theory and algorithms to reconstruct sparse signals, typically sum of sinusoids, from *non-uniform* samples. We achieve this by identifying a linear transformation that relates the unknown uniform samples of sinusoids to the given measurements. These uniform samples are known to satisfy the annihilation equations. A valid solution is then obtained by solving a constrained minimization such that the reconstructed signal is consistent with the given measurements and satisfies the annihilation constraint.

Thanks to the new approach, we unify a variety of FRI-based methods. We demonstrate the versatility and robustness of the proposed approach with four FRI reconstruction problems, namely Dirac reconstructions with irregular time or Fourier domain samples, FRI curve reconstructions and Dirac reconstructions on the sphere. The proposed algorithm improves substantially over state of the art methods.

### Summary of contributions in Chapter 2

- We unify various FRI-based sparse recovery methods in a simple yet flexible algorithmic framework, which extends the FRI sampling to cases with *non-uniform* samples.
- We propose an efficient and robust reconstruction algorithm that finds a valid solution to the continuous-domain sparse recovery problem such that the reconstruction (i) satisfies the continuous-domain sparsity constraint; and (ii) fits the given measurements (up to the noise level).
- We demonstrate the flexibility and robustness of the algorithmic framework in solving various FRI-based sparse recovery problems.

### Chapter 3 — Multi-dimensional FRI Sampling with Linear Sample Complexity

Thanks to the flexible sparse recovery framework developed in Chapter 2, we can revisit the important multi-dimensional Dirac estimation problem. Continuous-domain sparse recovery in two or higher dimensions is fundamental in imaging, be it for source reconstruction in radio astronomy, direction of arrival estimation in acoustics, or super-resolution in microscopic imaging.

Previous approaches extended one dimensional methods, like the ones based on finite rate of innovation (FRI) sampling, in a separable manner, e.g., along the horizontal and vertical dimensions separately in 2D. The separate estimation leads to a sample complexity of  $\mathcal{O}(K^D)$  for  $K$  Diracs in  $D$  dimensions, despite that the total degrees of freedom only increase linearly with respect to  $D$ .

We propose a new method that enforces the continuous-domain sparsity constraints simultaneously along all dimensions, leading to a reconstruction algorithm with *linear* sample complexity  $\mathcal{O}(K)$ , or a gain of  $\mathcal{O}(K^{D-1})$  over previous methods. The multi-dimensional Dirac locations are subsequently determined by the intersections of hypersurfaces (e.g., curves in 2D), which can be computed algebraically from the common roots of polynomials.

We show results on simulated data of multi-dimensional Diracs with noiseless and noisy measurements. The performance of the new multi-dimensional algorithm will be investigated

in detail for two reconstruction problems with real data in radio astronomy and acoustics in Chapter 4 and Chapter 5, respectively.

### Summary of contributions in Chapter 3

- We extend the generic FRI sampling and reconstruction framework in Chapter 2 to higher dimensions (two dimensions or more)—the resulting sampling scheme has a *linear* sample complexity in signal sparsity. It is a significant improvement over the previous methods, which require a minimum sample size that scales exponentially with respect to the dimension.
- We propose an iterative algorithm to solve the multi-dimensional FRI reconstruction. The algorithm makes it possible to apply the multi-dimensional FRI-based sparse recovery to practical problems, where the measurements are not necessarily taken uniformly, e.g., point source estimation in radio astronomy that we consider in the next chapter.

### Chapter 4 — Looking beyond Pixels with Continuous-space Estimation of Point Sources

Equipped with the robust reconstruction algorithms developed in the previous two chapters, we apply the generic algorithms to real applications in the rest of the thesis. The focus of this chapter is on source estimation in radio astronomy, which aims at estimating point source locations and intensities from measurements with antenna arrays.

Conventional imaging algorithms, e.g., the CLEAN algorithm (and its multiple variants), and compressed-sensing inspired methods, are both discrete in nature: sources are estimated on a regular grid. The resolution power of the traditional CLEAN-based imaging pipeline, is limited by the width of the synthesized beam, which is inversely proportional to the largest baseline.

The continuous formulation allows the FRI recovery to find sources below the perceived tool resolution without an artificial grid imposition. We adapt the multi-dimensional FRI reconstruction algorithm developed in the previous chapter to radio interferometry. The angular resolution performance of the algorithm is assessed under simulation, and with real measurements from the LOFAR telescope.

We show that it is possible to determine accurate off-grid point-source locations and their corresponding intensities. In addition, FRI-based sparse reconstruction requires fewer measurements and smaller baselines to reach a similar reconstruction quality compared to a conventional method. The achieved angular resolution is higher than the perceived instrument resolution, and very close sources can be reliably separated. It is demonstrated that the method is robust to the presence of extended-sources, and that false-positives can be addressed by choosing an adequate model order to match the noise level.

### Summary of contributions in Chapter 4

- We adapt the FRI sparse recovery technique to source estimation in radio interferometry. The proposed algorithm LEAP (Looking beyond Pixels with continuous-space EstimAtion of Point sources) can estimate sources reliably from single-band or multi-band measurements.
- We propose a reconstruction algorithm with cubic complexity in the total number of

estimated uniform Fourier data of the sky image (typically around a few thousand).

- We extensively test the resolving powers of the algorithm in simulations. We show that the algorithm is capable to surpass the instrument angular resolution and reliably distinguish closely located sources.
- The robustness of the approach is validated with real observation from the LOFAR radio telescope from the Boötes field and the Toothbrush cluster.

### Chapter 5 — Direction of Arrival Estimation with Microphone Arrays

Similarly to the radio astronomy point source estimation problem in the previous chapter, it is possible to apply the FRI-based continuous-domain sparse recovery to acoustic DOA estimations. In this chapter, we present FRIDA: an **FRI**-based **D**irection of **A**rrival estimation algorithm for multiple wide-band sound sources. FRIDA combines multi-band measurements coherently and achieves state-of-the-art resolution. FRIDA is applicable to arrays with arbitrary layouts. Unlike various steered-response power and subspace based methods, the proposed algorithm does not require grid search.

We validate the approach in both simulations and experiments with actual recordings of speech signals. The robustness of FRIDA is demonstrated with measurements under different noise conditions. Further, we show the significant improvements of FRIDA in reliably resolving closely located sources.

#### Summary of contributions in Chapter 5

- We adapt the FRI sparse recovery to estimate directions of arrival of acoustic sources in both 2D and 3D cases.
- The robustness of the proposed reconstruction algorithm is validated with both simulations and actual recordings of speech signals.

### Chapter 6 — Image up-sampling via Edge Etching

In the previous two application chapters, we enforce the exact signal sparsity constraint in reconstruction, with a Fourier domain convolution equation (i.e., the annihilation equation). Alternatively, we can promote signal sparsity in the time or spatial domain directly by requiring that the product between a mask function and the sparse signal be small. In this chapter, we adopt the latter approach for image up-sampling, where image edges are modeled as a sparse signal, which satisfies the *approximate* annihilation constraint in the spatial-domain.

The goal of image up-samplings is to reconstruct a high resolution image from a single low resolution one. It is essential to exploit prior knowledge on the reconstructed image in order to better condition this severely ill-conditioned inverse problem. We present a novel edge modeling framework based on a spatial domain interpretation of the annihilation constraint (see Figure 1.4 and [95]). More specifically, we define a continuous-domain *mask function* that vanishes around large gradients, typically around edges. The mask function is reconstructed by minimizing its product with the gradient image. We show that accurate edge models are reconstructed by assuming a simple linear edge model locally. Based on the same idea, we further combine these

local edge models and build a global one, which serves as an edge-preserving constraint in image up-sampling.

We propose an efficient optimization algorithm based on the alternating direction method of multipliers (ADMM) to implement the constrained image up-sampling problem. Experiments with natural images demonstrate the effectiveness of the global edge model in improving the quality of the up-sampled images thus achieving state of the art performance.

#### Summary of contributions in Chapter 6

- We propose an edge modeling framework, which gives a continuous domain representation of image edges, both in a local area and for the whole image. We give an example of one possible application of the reconstructed local edge model with light-field images.
- The estimated edge models are incorporated in the up-sampling process as an edge-friendly constraint. We adapt the alternating direction method of multipliers (ADMM) algorithm to solve the constrained optimization efficiently.





## Chapter 2

# Continuous Sparse Recovery with Finite Rate of Innovation Sampling<sup>\*</sup>

临渊羡鱼，不如退而结网。

Better to go home and weave a net, than to stand by the pond and covet fish.

---

淮南子·说林训 *Huainanzi · Discourse on Forests*

刘安 LIU, AN

### 2.1 Introduction

Consider a classic array signal processing problem in radio interferometry. The electromagnetic (EM) waves emitted by celestial sources in the sky are collected by an array of antennas. The received signals at two antennas differ by a phase shift, which depends on the relative distance of the antennas and the point source locations in the sky (Figure 1.1). It can be shown that the cross-correlation of the received EM waves is related to the Fourier transform of the underlying sky image (see Table 2.1) sampled at *non-uniform* frequencies [126]. The goal is to reconstruct these point sources, which are modeled as a weighted sum of Dirac delta distributions, from the irregularly sampled Fourier measurements in continuous space.

The classic approach in radioastronomy is to assume that the point sources are located on a discrete grid (i.e., gridding). The associated *discretized* sky image is then reconstructed by taking the inverse discrete Fourier transform (see e.g. Figure 2.1(a)) followed by an iterative deconvolution process [66]. Recently, it has been shown that the conventional discretized approach

---

<sup>\*</sup>The material in this chapter is the result of joint work with Thierry Blu and Martin Vetterli [96].

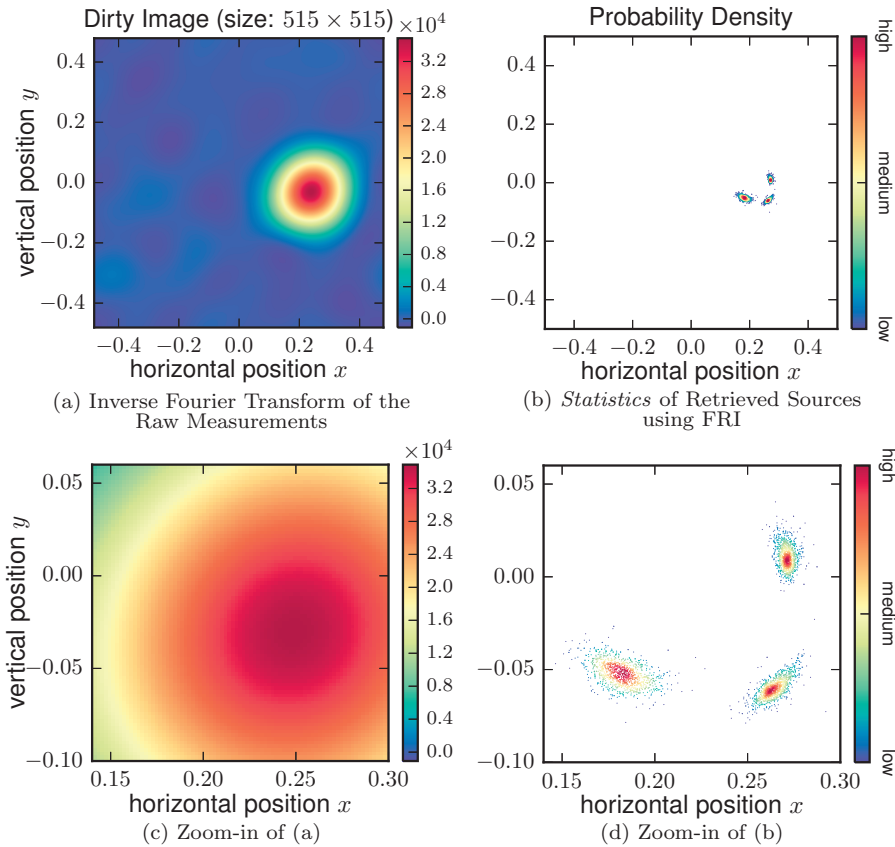
is related to compressed sensing, where the  $\ell_1$  norm of the *discretized* sky image is minimized subject to the data-fidelity constraint [25, 138, 140]. Note that the reconstruction accuracy of the sky image is inherently limited by the resolution of the grid: In order to obtain a more accurate reconstruction, a denser grid has to be used. Additionally, the measurement matrix, which relates the sky image to the Fourier measurements, is determined by the layout of the antenna arrays. It does not necessarily satisfy the restricted isometry property required in standard compressed sensing theory.

Alternatively, we can address the point source reconstruction problem directly in *continuous* space by using algorithms developed for signals with finite rate of innovation [133] (FRI). A common feature of these signals is that they can either be represented as or transformed to a weighted sum of sinusoids, which is also the case for point source reconstruction (see (2.16) in Section 2.5). The key to the reconstruction of FRI-type signals [8, 11, 29, 42, 45, 84, 88, 95, 115, 133] is the annihilating filter method, which is related to Prony’s method [103] in spectrum estimation: we can build a finite length discrete filter (i.e., the annihilating filter) such that its convolution with *uniformly* sampled sinusoids is zero. The point source locations are then given by the roots of a polynomial, whose coefficients are specified by the annihilating filter. However, despite the ability to reconstruct the point sources directly in continuous-domain, the classic FRI approach cannot deal with irregularly sampled data.

In this chapter, we propose a robust reconstruction algorithm that removes the uniform-sampling limitation from the FRI framework. Therefore, it allows to reconstruct the point sources in the continuous space with *irregularly* sampled Fourier measurements. We achieve this by establishing a linear relation between a set of unknown uniform Fourier samples, which can be annihilated by a discrete filter, and the given measurements (Figure 2.2). We recast the point source reconstruction as an approximation problem, where we would like to find a sum of Diracs that is *consistent* with the measurements: the discrepancy between the measured and re-synthesized samples (based on the reconstruction) should stay within the (known or estimated) noise level. A valid solution of the signal approximation problem is obtained with a constrained optimization, where the approximation error is minimized subject to the annihilation constraint (see Section 2.3 for details). Thanks to the new approach, point sources are recovered accurately in continuous space even in severe noise conditions (see Figure 2.1 and Figure 2.10). The proposed approach shows a substantial improvement in both accuracy and resolvability of closely located sources over a state of the art method based on discrete  $\ell_1$  minimization (see Figure 2.10).

It turns out that our contribution is much more general than the specific algorithm to solve the point source reconstruction in radioastronomy: in fact, all FRI reconstruction problems can be formulated concisely within the same algorithmic framework. In the proposed approach, we work directly with the given samples, which themselves may not be annihilated right away. With previous approaches [14, 47, 133], a linear transformation had to be applied to the samples first [48, 129]. We not only simplify the problem formulation but also can address cases that were overlooked and considered very challenging to solve with FRI (e.g., the Dirac reconstruction with non-uniform samples). We demonstrate the versatility and robustness with several examples, including the important Dirac reconstructions with *irregular* time / Fourier domain samples (Section 2.4.1 and 2.4.2), FRI curve reconstructions (Section 2.4.3) and the recovery of Diracs on the sphere (Section 2.4.4).

Our goal is to provide a unified algorithmic framework, which is simple yet flexible so as to cope with various FRI sampling problems. To facilitate future research and the applications of

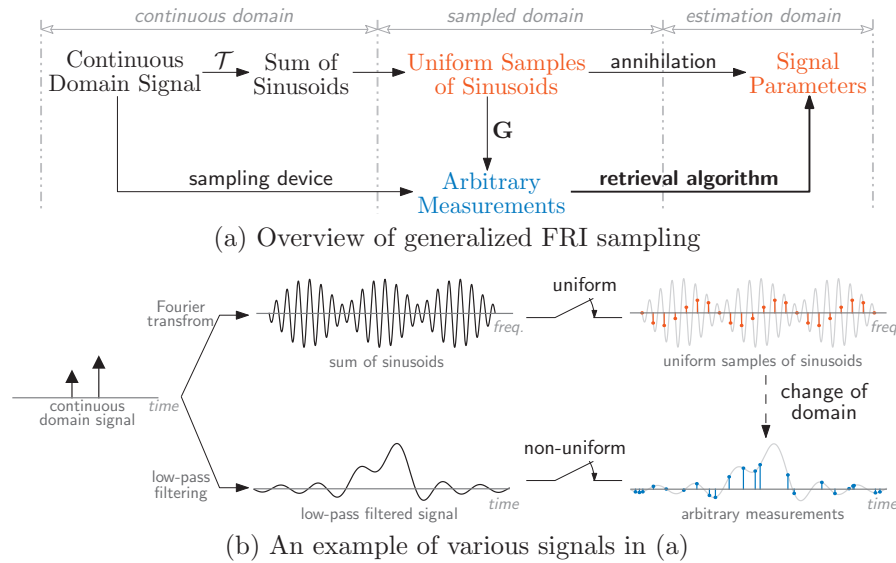


**Figure 2.1:** Accurate reconstruction of point source locations from partial Fourier domain measurements (number of irregular Fourier samples: 8000, SNR = 5 dB). (a) Spatial domain representation (a.k.a. “dirty image” in radioastronomy) associated with the given partial Fourier measurements. (b) Probability density of the reconstructed point source locations with the FRI approach (number of independent noise realizations: 1000; average estimation error of Dirac locations:  $8.07 \times 10^{-3}$ ). For comparison with other methods see Figure 2.10.

the proposed algorithm, a Python implementation of all the examples is made available online<sup>1</sup>.

Before proceeding, we briefly point to literature related to sparse signal recovery. We shall not review various algorithms for the classic spectrum estimation in detail but rather refer to standard textbooks (e.g., [118]) for comprehensive reviews. One notable algorithm in spectrum estimation is iterative quadratic maximum likelihood (IQML) [18] and was an inspiration for our approach. In a classical sparse recovery setting, the signal is recovered by minimizing the  $\ell_1$  norm of the target sparse signal [46]. Recently, it has been extended to continuous domain by minimizing the total variation [22] (or in general an atomic norm [123]) of the continuous sparse signal. It has been shown that the sparse signal can be reconstructed exactly provided that a minimum separation condition is met [6, 22, 79, 122, 123]. Alternatively, the optimal reconstruction can be formulated as a structured low-rank approximation, where the rank of

<sup>1</sup>The codes are available at [https://github.com/hanjiepan/FRI\\_pkg](https://github.com/hanjiepan/FRI_pkg)



**Figure 2.2:** Any continuous domain signal, which can be represented as a sum of sinusoids by applying a certain transformation  $\mathcal{T}$ , is an FRI signal. The classic FRI framework reconstructs the continuous domain signal from a set of *uniform* samples. Our focus in this chapter is on cases where measurements are taken irregularly. We will identify a linear mapping  $\mathbf{G}$  that relates uniform samples to these measurements with a good approximation.

a data matrix (typically of Toeplitz / Hankel structure) is minimized subject to a data fidelity constraint. This non-convex optimization is then solved either heuristically [5, 33, 86] or via convex relaxation [30, 93, 145].

The rest of this chapter is organized as follows. We briefly review the classic sampling and reconstruction framework for signals with finite rate of innovation in Section 2.2. It serves as the basis for Section 2.3, where we propose a novel algorithmic framework. Both the problem formulation and the implementation are developed. Next we demonstrate the versatility and the robustness of the proposed approach by solving four different FRI problems in Section 2.4. Further, we illustrate in detail the application to radio interferometry in Section 2.5 before we conclude in Section 2.6.

## 2.2 The classic FRI reconstruction framework

From an algorithmic point of view, if a continuous domain signal is or can be transformed into a finite sum of sinusoids, then it is a finite rate of innovation (FRI) signal [14, 133]. The FRI sampling problem then boils down to estimating frequencies of the sinusoids from the given measurements. From this perspective, it coincides with the classic harmonic retrieval problem [62], which is encountered in many applications [13, 89, 126].

One such example is point source reconstruction in radioastronomy as discussed briefly earlier. Similar principles are used for target localization in radar systems as well as in acoustic source localization with microphone arrays.

Another example is X-ray crystallography, where the goal is to determine the atom locations

from measured diffraction patterns. The diffraction patterns are the Fourier transform of an autocorrelation function, which is a sum of sinusoids [89]. The distances between atoms are directly related with the frequencies of these sinusoids.

The classic FRI approach (Figure 2.2) for solving the frequency retrieval problem consists in sampling the continuous domain signal uniformly. On the one hand, these uniform samples have a one-to-one correspondence with uniformly sampled sinusoids: by applying an inverse mapping (typically an inverse DFT transformation) to the uniform samples, we get an estimate of the sampled sinusoids. On the other hand, it is easy to show that these sampled sinusoids satisfy a set of annihilation equations [14, 133]: there exists a discrete filter (a.k.a. “annihilating” filter), which depends on the unknown frequencies of the sinusoids, such that its convolution with the sampled sinusoids is zero. The annihilating filter coefficients are reconstructed uniquely (up to a scaling factor) by solving a linear system of equations. The frequencies of the sinusoids are obtained from the roots of a polynomial whose coefficients are specified by the annihilating filter. Once we have recovered the frequencies of the sinusoids, the reconstruction of the sinusoids’ amplitudes amounts to a simple least square minimization, which reverts to solving a linear system of equations.

Various reconstruction algorithms have been proposed to improve the robustness, notably total least square minimization [133], matrix pencil approach [85, 129] and the Cadzow denoising [14, 21]. In the total least square minimization, one obtains the filter that minimizes the  $\ell_2$  norm of the annihilation error; the matrix pencil approach takes advantage of the shift-invariant subspace property in the structured data matrix; the Cadzow denoising method exploits the fact that the convolution matrix associated with the annihilating filter has a Toeplitz structure and is rank deficient. The Cadzow method denoises the data by iterating between a thresholding step (to ensure that the matrix is singular) and a projection step (to make the matrix Toeplitz). Recent works [5, 33] generalize such a strategy by formulating the reconstruction problem explicitly as a structured low-rank approximation.

For the rest of the chapter, we focus on more general cases where the continuous domain signal may not be sampled uniformly. In this case, we can no longer estimate the sampled sinusoids from the given measurements directly. We will develop a generic approach to solve the FRI reconstruction problem in this non-uniformly sampled case.

## 2.3 A generic FRI reconstruction algorithm


In this section, we propose a robust reconstruction algorithm for signals with finite rate of innovation from arbitrary samples. The reconstruction problem is recast as a constrained optimization (Section 2.3.1). We discuss the essential ingredients (Section 2.3.2), the optimization strategy (Section 2.3.3) as well as the implementation details (Section 2.3.3) of the proposed algorithm.

### 2.3.1 FRI reconstruction as a constrained optimization

We reformulate the generic FRI reconstruction question as an approximation problem, where we would like to fit an FRI model to the given measurements, or

*Given a set of measurements, reconstruct an FRI signal that is consistent with the measurements.*

The consistency constraint requires that if we re-synthesize the measurements based on the reconstructed FRI signal parameters, the difference with the given measurements should stay



**Figure 2.3:** Overview of the generic FRI reconstruction problem. The continuous domain FRI signal is reconstructed from the given measurements  $\mathbf{a}$  by enforcing the annihilation constraint on the unknown sampled sinusoids  $\mathbf{b}$ , which are related with  $\mathbf{a}$  via a linear mapping  $\mathbf{G}$ .

---

#### The Generic FRI Reconstruction Problem

---

**Goal:** Reconstruct an FRI signal  $\mathbf{b}$  that is consistent (up to the noise level  $\varepsilon^2$ ) with the given measurements  $\mathbf{a}$ :

$$\begin{aligned} & \text{find } \mathbf{b}, \mathbf{c} \in \mathcal{C} \\ & \text{subject to } \mathbf{b} * \mathbf{c} = \mathbf{0}, && // \mathbf{b} \text{ is FRI} \\ & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \leq \varepsilon^2 && // \text{consistency} \end{aligned}$$

**Key:**

- (i) The linear mapping  $\mathbf{G}$  between  $\mathbf{a}$  and  $\mathbf{b}$ ;
  - (ii) Constrained minimization (P1);
  - (iii) Stopping criterion: compare  $\|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2$  with the noise level  $\varepsilon^2$  (Section 2.3.2);
  - (iv) Random initialization (Section 2.3.2).
- 

within the noise level (or in general within the allowed approximation error). But how can we ensure that the reconstructed signal satisfies our FRI signal model? One key feature of many FRI signals is that they can be transformed into a sum of sinusoids. The uniform samples of sinusoids are known to be annihilated by a filter with a specific structure that is related to the FRI signal parametrization [14, 133]. Therefore, a signal being FRI is algorithmically equivalent to satisfying the annihilation constraint, after some linear transformation (see Figure 2.2).

**Problem 2.1 (Constrained Minimization)**

$$\begin{aligned} \min_{\mathbf{b}, \mathbf{c} \in \mathcal{C}} \quad & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to} \quad & \mathbf{b} * \mathbf{c} = \mathbf{0}, \end{aligned} \tag{P1}$$

where (Figure 2.3)

- $\mathbf{a}$  is the given set of measurements (sampled non-uniformly in general);
- $\mathbf{b}$  is the vector of uniform samples of the sinusoids to be annihilated. The convolution constraint guarantees that  $\mathbf{b}$  is effectively a sum of sinusoids;
- $\mathbf{c}$  is the annihilating filter coefficients, which belongs to a certain feasible set  $\mathcal{C}$  (see a precise discussion in the next subsection);
- $\mathbf{G}$  models the linear mapping<sup>2</sup> between the measurements  $\mathbf{a}$  and the uniform sinusoid samples  $\mathbf{b}$ . In general, in the presence of noise, we need to increase the number of measurements.


Superficially, the data term  $\|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2$  looks similar to the one used in compressed sensing (CS) based approaches. Yet, it arises from a completely different approach to resolution (see Figure 2.4):

- On the one hand, in CS-based sparse recovery the final recoverable resolution is directly related to the step-size of the uniform grid that supports the samples  $\mathbf{b}$  (see e.g., [123] Remark 1.2): perfect reconstruction is obtained for Diracs that are separated by a minimum distance that is proportional to the grid step-size. The matrix  $\mathbf{G}$  then encodes the “non-uniform” down-sampling that provides the known measurements  $\mathbf{a}$ ; for that reason,  $\mathbf{G}$  is necessarily a fat matrix.
- On the other hand, in FRI the final resolution is only related to the noise (or model mismatch) level. This error is typically given by the Cramér-Rao lower bound, which is reached by FRI-based reconstruction algorithms experimentally [14, 33]. The samples  $\mathbf{b}$  have to be taken on a uniform grid in order for the annihilation equation to be satisfied, but the step-size of the grid is flexible and is unrelated to the resolution of the method. In that context, the matrix  $\mathbf{G}$  models a different type of linear relation<sup>3</sup> between arbitrary uniform samples  $\mathbf{b}$  and arbitrary non-uniform known measurements  $\mathbf{a}$  than in compressed sensing; for that reason,  $\mathbf{G}$  can — and will, in this thesis — be a tall matrix.

Note that (P1) is non-convex with respect to  $(\mathbf{b}, \mathbf{c})$  jointly. Despite various attempts to solve similar constrained optimizations [18, 41], a reliable algorithm for finding the optimal solution of (P1) has yet to be discovered. But do we actually need to obtain the optimal solution of (P1)? In many cases, we know (or can estimate accurately) the noise level present in the given measurements. Hence, we may use this additional information in validating whether a solution is feasible or not: we claim that any solution  $\mathbf{b}$  is valid as long as it satisfies both the annihilation and consistency constraint for the given measurements (up to the noise level  $\varepsilon^2$ ), which we formalize as Problem (P2).

<sup>2</sup>For simplicity, we will assume that the linear transformation  $\mathbf{G}$  has full column rank throughout the thesis.

<sup>3</sup>One would expect that a  $\mathbf{G}$  matrix that maps too few uniform samples to the non-uniformly sampled measurements, would lead to worse reconstruction. However, this imbalance seems not to be critical experimentally (see examples in Section 2.4.2 and Section 2.5.3).



**Figure 2.4:** Problem setup differences between compressed sensing and finite rate of innovation (FRI). (a) In compressed sensing based approaches, the recoverable resolution (commonly known as the minimum separation requirement in compressed sensing) is related to the step-size of the underlying grid. (b) In FRI, the reconstruction accuracy is independent of the spacing between the uniform samples (see texts after (P1) in Section 2.3.1).

### Problem 2.2 (*Constrained Approximation with Noise Level*)

$$\begin{aligned}
 & \text{find } \mathbf{b}, \mathbf{c} \in \mathcal{C} \\
 & \text{subject to } \mathbf{b} * \mathbf{c} = \mathbf{0}, \\
 & \quad \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \leq \varepsilon^2.
 \end{aligned} \tag{P2}$$

One way to find a valid solution of (P2) is to resort to the constrained minimization (P1). However, we should keep in mind that it is not the optimal solution of (P1) that we seek but rather a valid solution that satisfies the constraints in (P2). Indeed, for any non-zero  $\varepsilon$ , (P2) has infinitely many solutions, among which one is the optimal solution of (P1). From an approximation point of view, all these solutions are valid because, for each of them, the reconstructed parametric signal explains the given measurements up to the noise level<sup>4</sup>. This subtle difference is important, since it allows to develop a reliable algorithm in the rest of this section.

### 2.3.2 Essential ingredients

Before we present the algorithm that solves (P2), we want to highlight five key elements of the proposed constrained formulation.

#### “Bilinearity” of the annihilation constraint

The annihilation equation is nonlinear with respect to the  $(\mathbf{b}, \mathbf{c})$ -pair. However, if we fix one variable (e.g.,  $\mathbf{c}$ ), then the annihilation constraint reduces to a linear constraint with respect to the other variable (e.g.,  $\mathbf{b}$  here). Motivated by the bilinearity of the annihilation constraint, we define a right dual operator.

<sup>4</sup>Notice that  $\varepsilon$  controls the approximation error with the ground truth: from the triangle inequality, the difference (2-norm) between the re-synthesized measurements with the ground truth signal and any one of the valid solution of (P2) is at most  $2\varepsilon$ . This corresponds to a maximum  $2\varepsilon/\sqrt{\lambda_{\min}(\mathbf{G}^H\mathbf{G})}$  difference between the reconstructed and the ground truth  $\mathbf{b}$ , where  $\lambda_{\min}(\cdot)$  is the smallest eigenvalue of a matrix.



**Definition 2.1**

Denote the annihilation constraint in (P2) as  $\mathbf{T}(\mathbf{b})\mathbf{c} = \mathbf{0}$ , then the right dual of  $\mathbf{T}(\cdot)$  is an operator  $\mathbf{R}(\cdot)$  such that  $\mathbf{R}(\mathbf{c})\mathbf{b} = \mathbf{T}(\mathbf{b})\mathbf{c}$  for all  $\mathbf{b}, \mathbf{c}$ .

In many FRI reconstruction problems, the annihilation equations are convolutions, which implies that  $\mathbf{T}(\mathbf{b})$  and  $\mathbf{R}(\mathbf{c})$  are Toeplitz-structured convolution matrices. We can justify the right dual definition from the commutativity of the convolution:  $\mathbf{b} * \mathbf{c} = \mathbf{c} * \mathbf{b}$ .

Thanks to the bilinearity of the annihilation constraint, it can be shown that the bivariate optimization (P1) is equivalent to a constrained optimization with respect to  $\mathbf{c}$  alone. This equivalent formulation provides an iterative strategy for finding a valid solution of (P2) (see Section 2.3.3).

**Forward mapping**

Unlike most annihilating filter based reconstruction algorithms [14, 47, 133], we deal with the measurements directly *without* pre-processing the given measurements first (e.g., a truncated DFT transformation). The linear mapping  $\mathbf{G}$ , which links the measurements to a sequence that can be annihilated, is integrated in the reconstruction algorithm. Thanks to the new approach, we are not only able to extend the FRI framework to cases with irregularly measured samples but also streamline otherwise rather complicated FRI reconstructions (see Section 2.3.3).

**Stopping criteria**

Because of the non-convexity of (P1), we should not expect the algorithm to always find the global optimal solution in general. In fact, it is *not* the optimal solution of (P1) that we should seek but rather *a* solution that (i) satisfies the annihilation constraint and (ii) has a fitting error  $\|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2$  below the noise level [58]. After all, our goal is to use the constrained minimization *as a tool* to find a valid solution of (P2)—any solution that meets the two criteria is a valid one for the FRI reconstruction.

The criteria are constructive: if we can guarantee that the reconstructed signal always satisfies the annihilation constraint, which is the case with the proposed algorithm (see details in Section 2.3.3), then we only need to check the fitting error in order to decide whether to terminate the algorithm or not.

**Random initialization**

Because of the non-convexity of (P1), a commonly used strategy is to initialize the algorithm with a “good” candidate solution, which is hopefully close to the ground truth, e.g., the total least square reconstruction. In our proposed algorithm, we choose to initialize the annihilating filter coefficients  $\mathbf{c}$  with a random vector instead.

The randomness of the initialization actually gives the algorithm the flexibility to have fresh restarts to increase the likelihood of meeting the stopping criteria—if the algorithm fails to find a solution that meets the aforementioned stopping criteria (Section 2.3.2), we can always reinitialize the algorithm with a different annihilating filter. The random initialization strategy has been shown to result in a valid solution within a finite number of initializations (typically less than 15) in extensive tests [59].

### Feasible set $\mathcal{C}$ of the annihilation filter coefficients

Observe that (P1) is scale invariant with respect to the annihilating filter coefficients  $\mathbf{c}$ . Without any normalization, we have a trivial solution  $\mathbf{c} \equiv \mathbf{0}$ . Experimentally, we have observed that the most robust performance is achieved by restricting<sup>5</sup>  $\mathbf{c}_0^H \mathbf{c} = 1$ , where  $\mathbf{c}_0$  is the random initialization for the algorithm (see [47] as well).

### 2.3.3 An iterative algorithm to solve (P2)

#### Inspiration from (P1)

For a given  $\mathbf{c}$ , (P1) is a constrained quadratic minimization with respect to  $\mathbf{b}$ . By substituting the solution  $\mathbf{b}$  (in a function of  $\mathbf{c}$ ) to (P1), we end up with an optimization over the annihilating filter coefficients  $\mathbf{c}$  alone. Denote the annihilation constraint in (P1) as a matrix vector product:  $\mathbf{T}(\mathbf{b})\mathbf{c} = \mathbf{0}$ . It can be shown that (P1) is equivalent to (see Appendix 2.A):

$$\begin{aligned} \min_{\mathbf{c}} \quad & \mathbf{c}^H \mathbf{T}^H(\boldsymbol{\beta}) (\mathbf{R}(\mathbf{c}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}))^{-1} \mathbf{T}(\boldsymbol{\beta}) \mathbf{c} \\ \text{subject to} \quad & \mathbf{c}_0^H \mathbf{c} = 1, \end{aligned} \quad (2.1)$$

where  $\boldsymbol{\beta} = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \mathbf{a}$ , and  $\mathbf{R}(\cdot)$  is the right dual of  $\mathbf{T}(\cdot)$  defined in Definition 2.1.

Then, the reconstructed FRI signal can be expressed as a function of  $\mathbf{c}$  as:

$$\mathbf{b} = \boldsymbol{\beta} - (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}) (\mathbf{R}(\mathbf{c}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}))^{-1} \mathbf{R}(\mathbf{c}) \boldsymbol{\beta}. \quad (2.2)$$

In general, it is very difficult to solve (2.1). But we can draw inspiration from this equivalent formulation and devise an iterative algorithm for finding a valid solution of (P2).

More specifically, our strategy amounts to minimizing the objective function in (2.1) iteratively: at each iteration, we build the matrix  $(\mathbf{R}(\mathbf{c}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}))^{-1}$  with  $\mathbf{c} = \mathbf{c}_{n-1}$ , the filter coefficients from the *previous* iteration. The updated  $\mathbf{c}$  is then obtained by solving a quadratic minimization:

$$\begin{aligned} \min_{\mathbf{c}} \quad & \mathbf{c}^H \mathbf{T}^H(\boldsymbol{\beta}) (\mathbf{R}(\mathbf{c}_{n-1}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}_{n-1}))^{-1} \mathbf{T}(\boldsymbol{\beta}) \mathbf{c} \\ \text{subject to} \quad & \mathbf{c}_0^H \mathbf{c} = 1, \end{aligned} \quad (2.3)$$

which has a simple closed form solution. For considerations of numerical stability, we are not going to implement the solution of (2.3) directly but revert to solving an equivalent linear system of equations instead (see Section 2.3.3). The uniform sinusoidal samples  $\mathbf{b}_n$  is updated based on (2.2) with the reconstructed  $\mathbf{c}_n$  at the *current* iteration.

Note that since (2.2) is obtained by solving (P1) for a given  $\mathbf{c}$ , by construction,  $\mathbf{b}_n$  obtained this way will always satisfy the annihilation constraint:  $\mathbf{R}(\mathbf{c}_n) \mathbf{b}_n = \mathbf{0}$ . Hence, we only need to compute the approximation error  $\|\mathbf{a} - \mathbf{G} \mathbf{b}_n\|_2^2$  and check whether it is below the noise level in order to terminate the iteration (see comments in Section 2.3.2).

The proposed approach may be judged similar to the *iterative quadratic maximum likelihood* [18] method in the spectrum estimation community, with the important difference: we only use the constrained optimization as a way to find a valid FRI reconstruction (P2). Hence, it

<sup>5</sup>Other (natural, but less successful) normalization strategies [14, 133] include a quadratic constraint  $\|\mathbf{c}\|_2^2 = 1$ ; or a linear constraint on one component of  $\mathbf{c}$ , e.g.,  $\mathbf{e}_0^H \mathbf{c} = 1$ , where  $\mathbf{e}_0 = [1, 0, \dots, 0]^T$ .

is not the convergent solution of (P1) that matters (as in [18]) but rather any solution that meets the stopping criteria in Section 2.3.2. The randomness in the initialization and the linear constraint give the algorithm more flexibility and has been shown [59] to achieve more robust reconstruction results for FRI problems.

### Efficient implementation

A direct implementation of (2.3) for the update of the annihilating filter coefficients involves several nested matrix inverses and would not only be inefficient (compared with solving linear system of equations) but also numerically unstable even with double-precision accuracy.

We can obtain the solution of (2.3) by solving a larger (compared with the dimension of  $\mathbf{c}$ ) linear system of equations with a simple trick: we introduce an auxiliary variable as a substitute of a matrix inverse applied to the input vector. Extra equations, which only involve multiplication of the matrix (instead of its inverse), are subsequently added to ensure that the resultant problem is equivalent to the original one.

#### Proposition 2.1

The solution of (2.3) is given by solving a linear system of equations

$$\begin{bmatrix} \mathbf{0} & \mathbf{T}^H(\boldsymbol{\beta}) & \mathbf{0} & \mathbf{c}_0 \\ \mathbf{T}(\boldsymbol{\beta}) & \mathbf{0} & -\mathbf{R}(\mathbf{c}_{n-1}) & \mathbf{0} \\ \mathbf{0} & -\mathbf{R}^H(\mathbf{c}_{n-1}) & \mathbf{G}^H\mathbf{G} & \mathbf{0} \\ \mathbf{c}_0^H & \mathbf{0} & \mathbf{0} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{c}_n \\ \ell \\ \mathbf{v} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ 1 \end{bmatrix}, \quad (2.4)$$

where  $\ell$ ,  $\mathbf{v}$ , and  $\lambda$  are newly introduced auxiliary variables. Similarly, the reconstructed FRI signal  $\mathbf{b}_n$  is updated as the solution of

$$\begin{bmatrix} \mathbf{G}^H\mathbf{G} & \mathbf{R}^H(\mathbf{c}_n) \\ \mathbf{R}(\mathbf{c}_n) & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{b}_n \\ \ell \end{bmatrix} = \begin{bmatrix} \mathbf{G}^H\mathbf{a} \\ \mathbf{0} \end{bmatrix}, \quad (2.5)$$

where  $\ell$  is the auxiliary variable.

*Proof.*

See Appendix 2.B. □

Instead of calculating several matrix inverses, we only need to solve a larger linear system of equations and extract the corresponding components of the solution in order to update  $\mathbf{c}$  and  $\mathbf{b}$  from iteration to iteration. For a vector  $\mathbf{b}$  of size  $L$ , the computational complexity at each iteration is  $\mathcal{O}(L^3)$  (see [61] Chapter 3).

In the actual implementation, we randomly initialize the algorithm with a maximum number of iterations, e.g., 50. At each iteration, we compute the approximation error  $\|\mathbf{a} - \mathbf{G}\mathbf{b}_n\|_2^2$  with the current reconstruction  $\mathbf{b}_n$  and compare it with the noise level. If the error is below the noise level, then the iteration is terminated. In the case where the algorithm fails to find such a solution after the maximum number of iterations is reached, we reinitialize the algorithm with a different random vector. We summarize the proposed algorithm in Algorithm 1. In [59], this strategy (50 inner iterations and 15 random initializations) is shown to succeed in 99.9% cases. Alternatively, we can always run the algorithm with a (fixed) maximum number of random initializations and return the reconstructed  $(\mathbf{b}, \mathbf{c})$ -pair that has the smallest fitting error. This strategy is useful for

**Algorithm 1:** Robust FRI Signal Reconstruction

---

**Input** : Measurements  $\mathbf{a}$  of the FRI signal, transformation matrix  $\mathbf{G}$ , noise level  $\varepsilon^2$   
**Output:** Uniform sinusoid samples  $\mathbf{b}$ , annihilating filter coefficients  $\mathbf{c}$   
**for**  $loop \leftarrow 1$  **to** *max. initializations* **do**  
1 | Initialize  $\mathbf{c}$  with a random vector  $\mathbf{c}_0$ ;  
| **for**  $n \leftarrow 1$  **to** *max. iterations* **do**  
2 | | Build the augmented matrix with  $\mathbf{c}_{n-1}$  and update  $\mathbf{c}_n$  by solving (2.4);  
3 | | Build the augmented matrix with  $\mathbf{c}_n$  and update  $\mathbf{b}_n$  by solving (2.5);  
| | **if**  $\|\mathbf{a} - \mathbf{G}\mathbf{b}_n\|_2^2 \leq \varepsilon^2$  **then**  
4 | | | Terminate both loops;  
| | **end**  
| **end**  
**end**  
5  $\mathbf{b} \leftarrow \mathbf{b}_n, \mathbf{c} \leftarrow \mathbf{c}_n$ .

---

cases where we do not know (or do not have a good estimate of) the noise level a priori, albeit less efficient than the approach used in Algorithm 1: with a given (or estimated) noise level, Algorithm 1 usually terminates much earlier before reaching the maximum number of random initializations.

## 2.4 Four algorithmic examples for various types of FRI signals

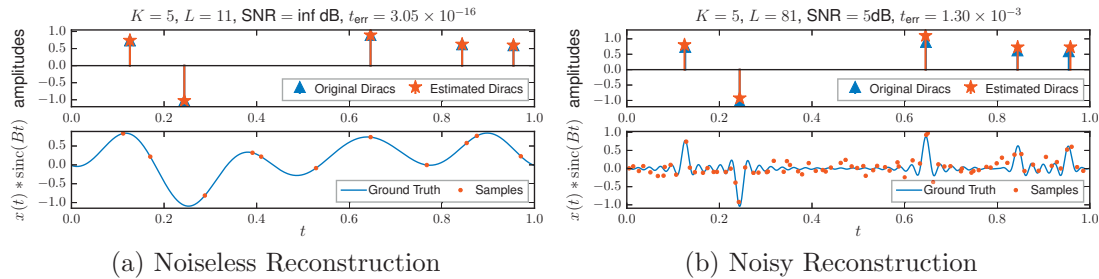
In the previous section, we presented a generic reconstruction algorithm for signals with finite rate of innovation. To demonstrate the versatility of the proposed algorithm, we showcase several FRI reconstruction problems, including the Dirac reconstruction with *non-uniform* time / Fourier domain samples (Section 2.4.1 and 2.4.2), FRI curve reconstruction (Section 2.4.3) and the reconstruction of Diracs on the sphere (Section 2.4.4).

A common misconception on annihilating-filter based FRI reconstruction algorithms is that they can only deal with *uniformly* taken samples, which have one to one correspondence with uniform samples of a sum of sinusoids, e.g., the DFT coefficients in the reconstruction of the Dirac stream (2.6). Such an artificial limitation is waived as soon as we are able to identify the forward mapping (or an approximation of it) that links the sequence to be annihilated and the given measurements, i.e.,  $\mathbf{G}$  in (P2). We show that we can either find the exact mapping (Section 2.4.1) or approximate it by interpolation (Section 2.4.2). The new formulation is flexible in the choice of the objective function or a proper constraint: instead of being restricted by the reconstruction algorithm, e.g., the Cadzow denoising method, we can use a different formulation, which is simpler and gives more robust reconstruction results (Section 2.4.3 and Section 2.4.4).

### 2.4.1 Stream of Diracs with irregular time domain samples

Consider a  $\tau$ -periodic stream of Diracs:

$$x(t) = \sum_{k' \in \mathbb{Z}} \sum_{k=1}^K \alpha_k \delta(t - t_k - k'\tau), \quad (2.6)$$



**Figure 2.5:** Reconstruction of a stream of Diracs (2.6) from ideally low-pass filtered samples taken at irregular time instances (2.7). (a) Exact reconstruction in the noiseless case (filter bandwidth  $B = 11$ , number of samples  $L = 11$ ). (b) Robust reconstruction in the noisy case (SNR = 5 dB, filter bandwidth  $B = 81$ , number of samples  $L = 81$ , average reconstruction error for  $t_k$ :  $1.30 \times 10^{-3}$ ).

where  $\alpha_k$  and  $t_k$  are unknown signal parameters. The goal is to reconstruct these parameters from a set of ideally low-pass filtered samples that are taken at *irregular* (but known) time instances  $t'_1, \dots, t'_L$ :

$$y_\ell = \sum_{k=1}^K \alpha_k \varphi(t'_\ell - t_k) \quad \text{for } \ell = 1, \dots, L. \quad (2.7)$$

Here  $\varphi$  is the Dirichlet kernel  $\varphi(t) \stackrel{\text{def}}{=} \frac{\sin(\pi B t)}{B \tau \sin(\pi t / \tau)}$  and  $B$  is the bandwidth of the ideal lowpass filter.

**Uniform samples of sinusoids** Observe that  $x(t)$  is a linear combination of the same function with different shifts—if we transform the signal to the Fourier domain, the spectrum is a sum of sinusoids:

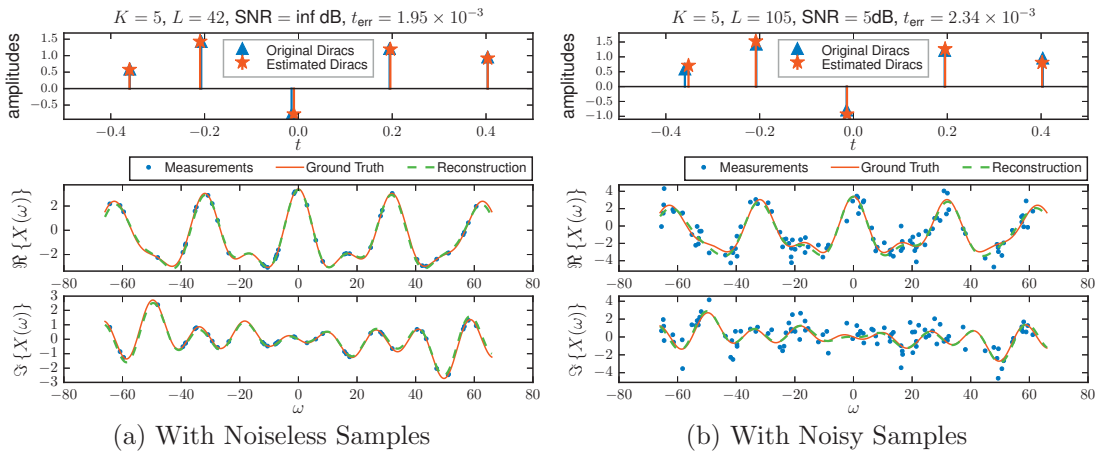
$$\hat{x}_m = \frac{1}{\tau} \sum_{k=1}^K \alpha_k e^{-j \frac{2\pi m}{\tau} t_k},$$

where  $\hat{x}_m$  is the Fourier series coefficients of the *periodic* signal  $x(t)$ . Since  $\hat{x}_m$  is a uniformly sampled sum of sinusoids, we know that it can be annihilated by a discrete filter.

**Relation with the given measurements** It is easy to show that the given measurements are linearly related with  $\hat{x}_m$  via a truncated inverse DFT transformation:

$$y_\ell = \frac{1}{B} \sum_{|m| \leq \lfloor \frac{\tau B}{2} \rfloor} \hat{x}_m e^{j \frac{2\pi m}{\tau} t'_\ell}. \quad (2.8)$$

In terms of the reconstruction algorithm, we can rearrange the samples  $y_\ell$  and  $\hat{x}_m$  as column vectors  $\mathbf{a}$  and  $\mathbf{b}$ , respectively; the linear mapping (2.8) between  $\mathbf{b}$  and  $\mathbf{a}$  is denoted by the matrix  $\mathbf{G}$ . Then the “denoised” FRI signal as well as the associated annihilating filter  $\mathbf{c} = [c_0, \dots, c_K]^T$  are given by Algorithm 1. The Dirac locations are reconstructed by taking the roots of the polynomial whose coefficients are specified by  $\mathbf{c}$ ; while the amplitudes are reconstructed with least square minimization [14, 133].



**Figure 2.6:** Reconstruction of weighted Diracs (2.9) from non-uniform Fourier samples (2.10). The FRI framework makes use of the *piecewise linear interpolation* (2.11) with 21 uniform knots. (a) Reconstruction with noiseless Fourier domain samples (number of samples  $L = 42$ , average reconstruction error for  $t_k$ :  $1.95 \times 10^{-3}$ ). (b) Robust reconstruction with noisy Fourier measurements (number of samples  $L = 105$ , SNR = 5 dB, average reconstruction error for  $t_k$ :  $2.34 \times 10^{-3}$ ).

We summarize the reconstruction results in Figure 2.5 for both the noiseless and noisy cases, where Gaussian white noise is added to the lowpass filtered samples. Note that the irregular sampling scheme does not change the minimum number of samples required in order to recover the original signal (Figure 2.5(a)): with at least  $2K + 1$  samples, the exact reconstruction (up to numerical accuracy) is obtained. In the presence of noise, we need to over-sample the signal. The proposed algorithm is robust enough to give a reliable reconstruction in the presence of severe noise (SNR = 5 dB).

## 2.4.2 Weighted sum of Diracs with irregular Fourier domain samples

In this example, we consider a slightly different Dirac reconstruction problem than that in the previous section. In particular, consider a sparse signal that consists of  $K$  weighted Diracs:

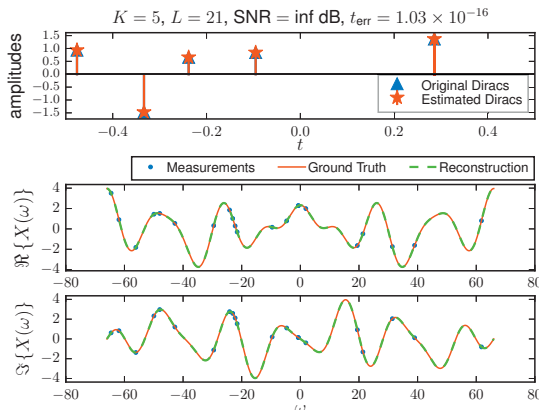
$$x(t) = \sum_{k=1}^K \alpha_k \delta(t - t_k), \quad (2.9)$$

with limited time support between  $-\tau/2$  to  $\tau/2$ , i.e.,  $|t_k| \leq \tau/2$ . Instead of taking the time domain samples (as in the previous example), the Fourier transform

$$X(\omega) = \sum_{k=1}^K \alpha_k e^{-j\omega t_k}, \quad (2.10)$$

is measured at some frequencies  $\omega_\ell$  for  $\ell = 1, \dots, L$ .

The question at hand is: can we recover the original signal (2.9) from non-uniform Fourier samples  $X(\omega_\ell)$ ? In many applications, e.g., magnetic resonance imaging [75], holography [56],



**Figure 2.7:** Exact reconstruction of weighted Diracs (2.9) with Dirichlet interpolation kernel from irregular Fourier domain measurements (period of the spectrum:  $2\pi \times 21$ ; number of samples  $L = 21$ ).

crystallography [89] and radio interferometry [101], direct Fourier domain measurements are available thus making the sparse reconstruction problem of particular interest.

**Uniform samples of sinusoids** Since the Fourier transform (2.10) is a weighted sum of sinusoids, the uniformly sampled Fourier transform on a grid,  $X(2\pi m/\tau)$  for  $m \in \mathbb{Z}$ , can be annihilated.

**Relation with the given measurements** In general, the given Fourier measurements are taken non-uniformly. Hence, we cannot apply the annihilating filter method directly. However, not everything is lost: we may interpolate the Fourier transform over a finite interval, e.g.,  $\omega \in [-M\pi, M\pi]$ :

$$X(\omega) \approx \sum_{|m| \leq \lfloor \frac{M\tau}{2} \rfloor} X\left(\frac{2\pi m}{\tau}\right) \psi\left(\frac{\omega}{2\pi/\tau} - m\right), \quad (2.11)$$

where  $\psi(\cdot)$  is a certain interpolation kernel, e.g., a spline function.

By evaluating (2.11) at  $\omega_\ell$ , we establish a linear mapping (i.e.,  $\mathbf{G}$  in the reconstruction algorithm) between the given Fourier measurements  $\mathbf{a} : X(\omega_\ell)$  and the unknown sampled sinusoids  $\mathbf{b} : X(2\pi m/\tau)$ . Provided that we have sufficiently many measurements, i.e.,  $L \geq M\tau$ , then we can reconstruct (2.9) with Algorithm 1 (see [5] for a similar strategy in spectral estimation).

We may justify such an approach by considering a specific case, where the Fourier transform  $X(\omega)$  is periodic with period  $2\pi M$  for some  $M$  such that  $M\tau$  is an odd number. It is proved that we can represent  $X(\omega)$  exactly by interpolating with the Dirichlet kernel  $\psi(\omega) = \frac{\sin(\pi\omega)}{M\tau \sin(\frac{\pi\omega}{M\tau})}$  in this case (see Appendix 2.C and Figure 2.7).

Note that rather than enforcing the interpolation equation (2.11) as a hard constraint on the reconstructed signal, we only use it to derive a data-fidelity metric in (P2) that measures the approximation quality. The tradeoff is that we can no longer reconstruct the signal exactly in general—We do not have the actual mapping  $\mathbf{G}_0$  (which depends on the unknown signal parameters  $\alpha_k$  and  $t_k$ ) but only its approximation from the interpolation (2.11):  $\mathbf{G} = \mathbf{G}_0 + \mathbf{G}_\varepsilon$ .

Consequently, we will have a model mismatch even in the noiseless cases<sup>6</sup> (Figure 2.6(a)). As we have mentioned in Section 2.3.3, one possible way to circumvent the difficulty in choosing  $\varepsilon$  in (P2) is to run the algorithm with fixed random initializations. The solution that gives the minimum fitting error is taken as the reconstruction.

We demonstrate the effectiveness of the interpolation strategy (2.11) in Figure 2.6, where the interpolation kernel is the first order B-spline:

$$\psi(\omega) = \begin{cases} 1 - \omega & \text{if } \omega \in [0, 1), \\ 1 + \omega & \text{if } \omega \in [-1, 0), \\ 0 & \text{otherwise.} \end{cases}$$

We have chosen 21 interpolation knots located uniformly on the interval  $[-21\pi, 21\pi]$ , where the Fourier transform  $X(\omega)$  is approximated. Complex-valued Gaussian white noise is added to the Fourier samples in the noisy case (SNR = 5 dB). Even with such a coarse approximation, we still obtain robust and accurate reconstruction of Diracs in the presence of noise (Figure 2.6).

### 2.4.3 FRI curves

As we mentioned in the introduction, the Cadzow denoising algorithm [14] tries to find a structured matrix (typically Toeplitz / block-Toeplitz) that satisfies the rank constraint while being as close as possible to the noisy data matrix. With the Cadzow denoising method, we are restricted to work directly with a sequence that can be annihilated (so that we can enforce the rank constraint on the matrix). In comparison, we have more freedom with the proposed algorithm in defining what is the unknown data  $\mathbf{b}$  other than the obvious choice as the sampled sinusoids. We demonstrate this flexibility with an example of curves with finite rate of innovation [95].

Consider an interior indicator image associated with a curve:

$$I_C(x, y) = \begin{cases} 1 & \text{if } (x, y) \in \mathring{C} \\ 0 & \text{otherwise,} \end{cases}$$

where  $\mathring{C}$  denotes the interior of the curve  $C$ . Our goal is to reconstruct the curve locations in the *continuous* domain from a set of ideally lowpass filtered samples of the binary image  $I_C(x, y)$ .

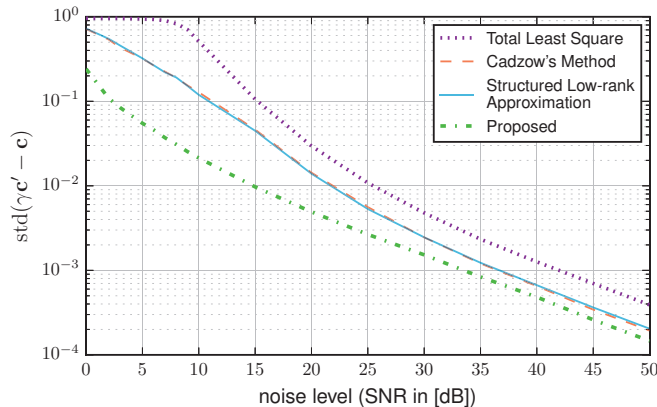
**Uniform samples of sinusoids** We may treat the derivative of the indicator image as an infinite sum (in fact a line integration) of Diracs along the curve. In the Fourier domain, the Wirtinger derivative (i.e.,  $\partial \stackrel{\text{def}}{=} \frac{\partial}{\partial x} + j \frac{\partial}{\partial y}$ ) is therefore a sum of sinusoids. Consequently, we know that the Fourier transform of the derivative image on a uniform grid,

$$\widehat{\partial I}_{k,l} = \left( \frac{2\pi k}{\tau_1} + j \frac{2\pi l}{\tau_2} \right) \hat{I}_C \left( \frac{2\pi k}{\tau_1}, \frac{2\pi l}{\tau_2} \right), \quad (2.12)$$

satisfies the annihilation equations:  $c_{k,l} * \widehat{\partial I}_{k,l} = 0$ . The curve locations are specified by the roots of a 2D polynomial with coefficients  $c_{k,l}$ . See [95] for detailed discussions on FRI curves.

<sup>6</sup>Equivalently, we can view the noiseless measurements  $\mathbf{a} = \mathbf{G}_0 \mathbf{b}$  as being “noisy” with respect to  $\mathbf{G}$ , which is used in the data-fidelity constraint in (P2):  $\mathbf{a} = \mathbf{G} \mathbf{b} + \text{noise}$ , with  $\text{noise} = -\mathbf{G}_\varepsilon \mathbf{b}$ .





**Figure 2.8:** FRI curve coefficient error against different noise levels (curve coefficients  $c_{k,l}$  size:  $3 \times 3$ , sample size:  $45 \times 45$ , bandwidth  $B_1 = B_2 = 25$ , periods  $\tau_1 = \tau_2 = 1$ ). The results are averaged over 500 independent noise realizations.

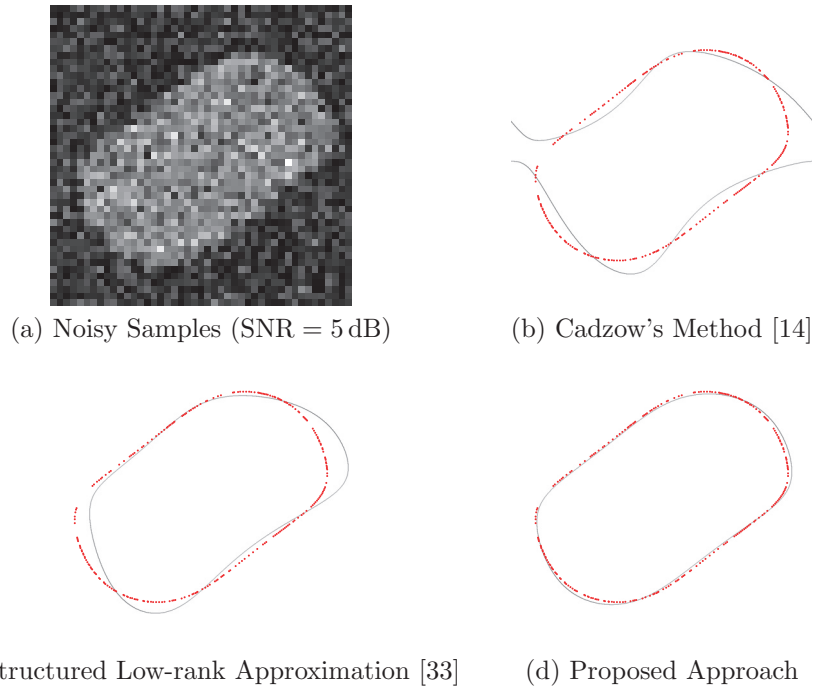
**Relation with the given measurements** Similar to the 1D case in Section 2.4.1, the Fourier transform of the indicator image on a uniform grid is related with the ideally lowpass filtered samples<sup>7</sup> via a truncated inverse DFT. Combined with (2.12), we have a linear mapping from the unknown sampled sinusoids  $\widehat{\partial I}_{k,l}$  to the measured spatial domain samples.

In our original approach [95], we first obtain  $\hat{I}_C\left(\frac{2\pi k}{\tau_1}, \frac{2\pi l}{\tau_2}\right)$  by applying a truncated DFT transformation to the given samples. Then we apply Cadzow's method to denoise  $\widehat{\partial I}_{k,l}$ , since it is  $\widehat{\partial I}_{k,l}$  that satisfies the annihilation. Unfortunately, doing so inevitably amplifies the high frequency noise components, which explains the relatively limited performance of Cadzow's method for FRI curve reconstructions (Figure 2.8).

In our new approach, we do not have such a restriction anymore: in (P2), we can choose  $\mathbf{a}$  directly as the pixel values  $I_{m,n}$ , the unknown  $\mathbf{b}$  as the Fourier transform  $\hat{I}_C\left(\frac{2\pi k}{\tau_1}, \frac{2\pi l}{\tau_2}\right)$ , and the linear mapping  $\mathbf{G}$  as the truncated inverse DFT transformation. The right dual matrix  $\mathbf{R}(\cdot)$  in this case is no longer the convolution matrix associated with the filter  $c_{k,l}$  alone—we should right multiply the convolution matrix by a diagonal matrix whose entries are specified by the corresponding frequencies  $\left(\frac{2\pi k}{\tau_1}, \frac{2\pi l}{\tau_2}\right)$  in (2.12).

We summarize the reconstruction results obtained with total least square minimization [14, 133], Cadzow's method [14], structured low-rank approximation [33] and the proposed method in Figure 2.8, where different levels of Gaussian white noise is added to the ideally lowpass filtered samples. Since the annihilating filter coefficients are invariant with respect to any non-zero scaling, we measure the reconstruction error with a scale-invariant standard deviation of the error between the ground truth  $\mathbf{c}$  and the reconstructed coefficients  $\mathbf{c}'$ :  $\text{std}(\gamma\mathbf{c}' - \mathbf{c})$ . Here the scalar  $\gamma$  is chosen in such a way that  $\|\gamma\mathbf{c}' - \mathbf{c}\|_2^2$  is minimized. An example at noise level  $\text{SNR} = 5$  dB is also included for visual comparison (Figure 2.9). The proposed algorithm is robust even under such severe noise condition.

<sup>7</sup>Note that with the same argument as in Section 2.4.1, we can deal with ideally lowpass filtered samples that are taken non-uniformly.



**Figure 2.9:** Visual comparisons of the reconstructed curves with Cadzow's method, structured low-rank approximation [33] and the proposed approach (noise level: 5 dB, curve coefficients  $c_{k,l}$  size:  $3 \times 3$ , sample size:  $45 \times 45$ , periods  $\tau_1 = \tau_2 = 1$ ). The solid black line is the reconstructed curve; while the dotted red line is the ground truth.

#### 2.4.4 Diracs on the sphere

In most cases, the given measurements of an FRI signal cannot be annihilated directly. With previous approaches, we had to apply an inverse transformation in order to obtain a sequence that can be annihilated. However, such an inverse transformation is not always easy to identify. Additionally, it makes the reconstruction problem unnecessarily complicated.

We demonstrate how the proposed algorithmic framework helps simplify the formulation of FRI problems. Consider an example of Diracs defined on the 2-sphere  $\mathbb{S}^2 \stackrel{\text{def}}{=} \{\mathbf{r} \in \mathbb{R}^3 | \mathbf{r}^T \mathbf{r} = 1\}$ :

$$I(\theta, \varphi) = \sum_{k=1}^K \alpha_k \delta(\cos \theta - \cos \theta_k) \delta(\varphi - \varphi_k),$$

for  $0 \leq \theta, \theta_k < \pi$  and  $0 \leq \varphi, \varphi_k < 2\pi$ . Here  $\theta_k$  and  $\varphi_k$  are the unknown angles of the colatitude and azimuth of the  $k$ -th Dirac and  $\alpha_k$  is its amplitude.

Similar to the planar case, if we apply a Fourier-like transformation to Diracs, we will get a sum of sinusoids. More specifically, the Fourier-like transformation on  $\mathbb{S}^2$  involves spherical harmonics. It can be shown (see Appendix 2.D) that the coefficients of the spherical harmonic

decomposition is

$$\hat{I}_{l,m} = N_{l,m} \sum_{n=0}^{l-|m|} p_{n,|m|} \sum_{k=1}^K \alpha_k (\cos \theta_k)^n (\sin \theta_k)^{|m|} e^{-jm\varphi_k}, \quad (2.13)$$

with  $|m| \leq l$  and  $0 \leq l \leq L_0$ , for some *fixed* coefficients  $p_{n,|m|}$  that can be precomputed (see Appendix 2.E). Here  $N_{l,m}$  is a normalization factor associated with the spherical harmonics of degree  $l$  and order  $m$ .

The FRI reconstruction problem that we want to solve is as follows: reconstruct  $(\alpha_k, \theta_k, \varphi_k)$  for  $k = 1, \dots, K$  from a given set of spherical harmonic coefficients (2.13). In an actual setup, the measurements are spatial domain samples (on  $\mathbb{S}^2$ ), which have a linear relationship [42, 45] with the spherical harmonic coefficients (2.13). With the proposed algorithm, we can reconstruct the signal from the spatial domain samples directly. But the complexity is beyond the scope of this section and hence is omitted.

**Uniform samples of sinusoids** Complicated as it may appear, (2.13) is a linear combination of uniformly sampled sinusoids:  $b_{n,m} \stackrel{\text{def}}{=} \sum_{k=1}^K \alpha_k (\cos \theta_k)^n (\sin \theta_k)^{|m|} e^{-jm\varphi_k}$ , then

$$b_{n,m} = \sum_{k=1}^K \tilde{\alpha}_{n,k}^{(1)} u_k^m = \sum_{k=1}^K \tilde{\alpha}_{m,k}^{(2)} v_k^n, \quad (2.14)$$

where  $\tilde{\alpha}_{n,k}^{(1)} = \alpha_k (\cos \theta_k)^n$ ,  $u_k = \sin \theta_k e^{-j\varphi_k}$  for  $m \geq 0$  and  $u_k = \frac{1}{\sin \theta_k} e^{-j\varphi_k}$  for  $m < 0$ ;  $\tilde{\alpha}_{m,k}^{(2)} = \alpha_k (\sin \theta_k)^{|m|}$  and  $v_k = \cos \theta_k$ . Consequently, we know that there exist two discrete filters such that

$$b_{n,m} *_{m} c_m^{(1)} = 0 \quad \forall n \quad \text{and} \quad b_{n,m} *_{n} c_n^{(2)} = 0 \quad \forall m.$$

**Relation with the given measurements** The expression (2.13) is nothing but a linear mapping from the sampled sinusoids  $b_{n,m}$  to the given measurements (i.e., the spherical harmonics  $\hat{I}_{l,m}$ ). Thanks to this analysis, we are now ready to formulate the Dirac reconstruction problem on the sphere as two (for  $u_k$  and  $v_k$  each) constrained approximations of the form (P2). Take the reconstruction of  $v_k$  as an example:

- **a** is the given spherical harmonics  $\hat{I}_{l,m}$ ;
- **b** is  $b_{n,m}$  as defined in (2.14);
- **c** is the annihilating filter coefficients  $c_k^{(2)}$ ;
- **G** is the linear relation between the given measurements  $\hat{I}_{l,m}$  and  $b_{n,m}$  in (2.13).

By reconstructing  $u_k$  and  $v_k$  from the given spherical harmonics, the angles  $\varphi_k$  and  $\theta_k$  are uniquely specified. The Dirac amplitudes  $\alpha_k$  can be easily obtained using least square minimization once we have reconstructed the values of  $\varphi_k$  and  $\theta_k$ .

One major challenge in the earlier work [45] was to find the correct inverse transformation that should be applied to  $\hat{I}_{l,m}$ . In comparison, such an inverse mapping is no longer required with the proposed framework, which leads to a significantly more simplified formulation.

Table 2.1: Summary of Radio Astronomy Terms

<i>Term</i>	<i>Meaning</i>
<i>sky image</i>	Brightness distribution of the sky.
<i>point sources</i>	Celestial sources that can be modeled as Dirac delta distributions.
<i>visibility</i>	Cross-correlations of the received signals at different antennas. It is related to the Fourier transform of the sky image.
<i>uv-coverage</i>	Fourier domain coverage. A radio interferometer can only cover part of the Fourier domain at some irregular frequencies.
<i>dirty image</i>	Inverse Fourier transform of the irregularly sampled Fourier transform of the sky image.

## 2.5 Application to radio interferometry

In this section, we apply the proposed reconstruction algorithm to a simplified radio interferometry problem. Cases with realistic settings will be considered in Chapter 4, where the FRI-based sparse recovery method is applied to real observation from the LOFAR radio telescope.

### 2.5.1 Data acquisition and signal model

A radio interferometer consists of an array of antennas that collect the electromagnetic (EM) waves emitted by celestial sources in the sky. In a far field context, we can assume that these sources are located on a hypothetical celestial sphere and that the signals arriving at each antenna follow parallel lines (Figure 1.1). Consequently, the received signals at two different antennas differ by a time delay, which is determined by the relative locations of the antennas with respect to the celestial sources. It can be shown that under the assumption of a narrow field-of-view, the cross-correlation between the received EM waves at two different antennas (a.k.a. visibility in radioastronomy) is related to the Fourier transform of the underlying sky image  $I(x, y)$  at a certain frequency (see [126] Chapter 3, equation (3.10)). Since there is a finite number of antennas with fixed locations, the radio interferometer will only have a *partial* Fourier domain coverage.

The conventional approach reconstructs the point sources in the discrete space by deconvolving the dirty image iteratively, which is the inverse discrete Fourier transform of the irregularly sampled Fourier measurements. Alternatively, as we demonstrate in the next section, we can directly address the reconstruction problem in the continuous-domain. In particular, our focus in this section is on the reconstruction of a sky image, which consists of *point sources* within the field of view:

$$I(x, y) = \sum_{k=1}^K \alpha_k \delta(x - x_k, y - y_k). \quad (2.15)$$

Here  $\delta(\cdot, \cdot)$  is the Dirac delta distribution or generalized function,  $(x_k, y_k)$  is the location of the  $k$ -th point source, and  $\alpha_k \geq 0$  is its intensity.

To summarize, the point source reconstruction problem in radio interferometry is as follows: how can we reconstruct the  $K$  Diracs on a 2D plane (2.15) from a given set of Fourier domain

measurements at irregular frequencies:

$$\hat{I}(\omega_1^{(\ell)}, \omega_2^{(\ell)}) = \sum_{k=1}^K \alpha_k e^{-j\omega_1^{(\ell)} x_k - j\omega_2^{(\ell)} y_k}, \quad (2.16)$$

for  $\ell = 1, \dots, L$ ? Note that in a realistic setting, these frequencies  $(\omega_1^{(\ell)}, \omega_2^{(\ell)})$  should be based on the layout of the radio telescope. We have considered a simplified experimental setup here so as to be as close as possible to the algorithmic examples in the previous section. We leave the extra complications encountered in practice for detailed analysis of real data acquired with a radio telescope in Chapter 4.

### 2.5.2 Reconstruction of point sources

Note that we have considered a similar 1D Dirac reconstruction problem in Section 2.4.2. Hence, we adopt the same strategy and approximate the Fourier transform over a finite area, e.g.,  $\omega_1 \times \omega_2 \in [-M\pi, M\pi] \times [-N\pi, N\pi]$ . For practical considerations, we have chosen the Dirichlet interpolation kernel:

$$\psi(\omega_1, \omega_2) = \frac{\sin(\pi\omega_1) \sin(\pi\omega_2)}{MN\tau_1\tau_2 \sin\left(\frac{\pi\omega_1}{M\tau_1}\right) \sin\left(\frac{\pi\omega_2}{N\tau_2}\right)},$$

for some  $M$  and  $N$  such that  $M\tau_1$  and  $N\tau_2$  are odd numbers. The interpolation equation provides the link between the given Fourier measurements  $\hat{I}(\omega_1^{(\ell)}, \omega_2^{(\ell)})$  and the sampled sinusoids  $\hat{I}(2\pi m/\tau_1, 2\pi n/\tau_2)$  for  $|m| \leq \lfloor \frac{M\tau_1}{2} \rfloor$  and  $|n| \leq \lfloor \frac{N\tau_2}{2} \rfloor$ , which can be annihilated by a discrete 2D filter.

In general, the solution that satisfies the 2D annihilation equations is a curve instead of a few isolated Diracs [95]. In fact, any Dirac that is located on the curve will satisfy the same set of 2D annihilation equations. One way to overcome such a difficulty is to reconstruct the Dirac's  $x$  and  $y$  locations *separately* by enforcing the annihilation constraint along each direction<sup>8</sup>. Specifically, we would like to find two annihilating filters, whose  $z$ -transforms are

$$C^{(1)}(z_1) = \sum_{k=0}^K c_k^{(1)} z_1^{-k} = c_0^{(1)} \prod_{k=1}^K (1 - u_k z_1^{-1}) \quad \text{and} \quad C^{(2)}(z_2) = \sum_{k=0}^K c_k^{(2)} z_2^{-k} = c_0^{(2)} \prod_{k=1}^K (1 - v_k z_2^{-1})$$

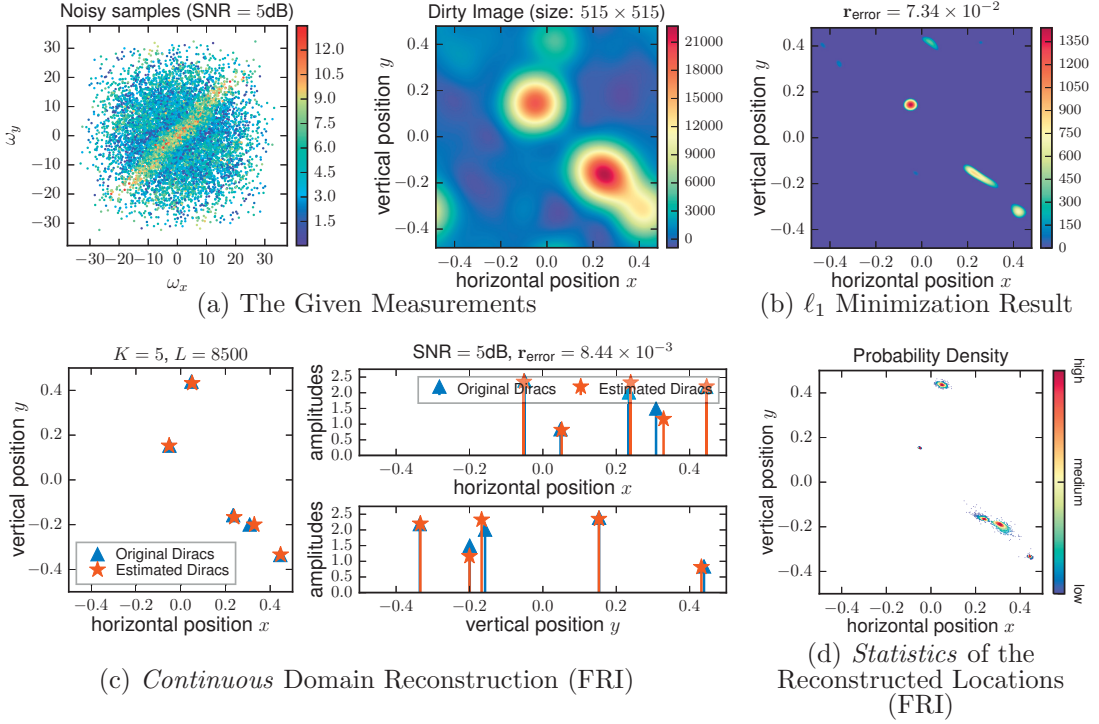
with  $u_k = e^{-j\frac{2\pi}{\tau_1} x_k}$  and  $v_k = e^{-j\frac{2\pi}{\tau_2} y_k}$ . The rows and columns of the Fourier transforms  $\hat{I}(2\pi m/\tau_1, 2\pi n/\tau_2)$  are annihilated by the filter  $[c_0^{(1)}, \dots, c_K^{(1)}]$  and  $[c_0^{(2)}, \dots, c_K^{(2)}]$ , respectively:

$$c_m^{(1)} * \hat{I}\left(\frac{2\pi m}{\tau_1}, \frac{2\pi n}{\tau_2}\right) = 0 \quad \text{and} \quad c_n^{(2)} * \hat{I}\left(\frac{2\pi m}{\tau_1}, \frac{2\pi n}{\tau_2}\right) = 0.$$

The Dirac locations are then reconstructed by solving two constrained approximation problems (P2). For the sake of brevity, we detail the exact formulation for the reconstruction of the Dirac vertical locations  $y_k$  only. The formulation for the reconstruction of  $x_k$  can be derived similarly.

Denote the annihilating filter coefficients  $c_k$  and the given Fourier measurements  $\hat{I}(\omega_1^{(\ell)}, \omega_2^{(\ell)})$

<sup>8</sup>A more efficient multi-dimensional FRI reconstruction, which enforces the annihilation constraints along all dimensions simultaneously, is the focus of the next chapter.



**Figure 2.10:** Reconstruction of point sources from irregular Fourier measurements (SNR = 5 dB, number of Fourier measurements:  $L = 8500$ ). (a) The given noisy Fourier samples and their spatial domain representation via inverse FFT (a.k.a., the the dirty image in radioastronomy). (b) The compressed sensing result by minimizing the  $\ell_1$  norm of the sky image (estimation error for point sources' locations:  $7.34 \times 10^{-2}$ ). (c) The reconstructed point sources with FRI (estimation error for point sources' locations:  $8.44 \times 10^{-3}$ ). (d) Probability density of the estimated point sources' locations with FRI approach (number of independent noise realizations: 1000; the average estimation error:  $1.09 \times 10^{-2}$ ).

as column vectors  $\mathbf{c} = [c_0^{(2)}, \dots, c_K^{(2)}]^T$  and  $\mathbf{a}$ , respectively. We rearrange the unknown Fourier transform values  $\hat{I}(2\pi m/\tau_1, 2\pi n/\tau_2)$  column by column as a vector  $\mathbf{b}$ . Then the constrained minimization that we would like to solve is:

$$\begin{aligned} \min_{\mathbf{b}, \mathbf{c}} \quad & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to} \quad & \mathbf{R}(\mathbf{c})\mathbf{b} = \mathbf{0}, \\ & \mathbf{c}_0^H \mathbf{c} = 1, \end{aligned}$$

where  $\mathbf{G} = [\tilde{\mathbf{G}}^{(-\lfloor M\tau_1/2 \rfloor)} \dots \tilde{\mathbf{G}}^{\lfloor M\tau_1/2 \rfloor}]$  with  $[\tilde{\mathbf{G}}^{(m)}]_{\ell, n} = \psi\left(\frac{\omega_1^{(\ell)}}{2\pi/\tau_1} - m, \frac{\omega_2^{(\ell)}}{2\pi/\tau_2} - n\right)$ ; and<sup>9</sup>

$$\mathbf{R}(\mathbf{c}) = \overbrace{\begin{bmatrix} \tilde{\mathbf{R}}(\mathbf{c}) & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{R}}(\mathbf{c}) & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \tilde{\mathbf{R}}(\mathbf{c}) \end{bmatrix}}^{2\lfloor M\tau_1/2 \rfloor \text{ blocks}},$$

with  $\tilde{\mathbf{R}}(\mathbf{c})$  the convolution matrix associated with the filter  $\mathbf{c}$ .

Once we have reconstructed the Dirac vertical and horizontal locations with the proposed algorithm in Section 2.3, we still need to identify the correct associations—in principle, the 2D Diracs can be located on any one of the  $K!$  possible combinations. The naive way would be the exhaustive search, where we try all the possible combinations and reconstruct the Dirac amplitudes  $\alpha_k$  with the least square minimization. If the re-synthesized Fourier samples (2.16) based on the reconstructed parameters  $(\alpha_k, x_k, y_k)$  are within the noise level, then we have correctly identified the Dirac locations.

However, such a straightforward approach is only computationally feasible for cases with few Diracs. Experimentally, we observed that we can find the correct associations of vertical and horizontal locations with a simple trick: we first reconstruct the amplitudes of the Diracs with the least square minimization by pretending there were  $K^2$  Diracs (i.e., all the intersections of the reconstructed  $\{x_k\}_{k=1}^K$  and  $\{y_k\}_{k=1}^K$ ). Among these  $K^2$  possible locations, we select  $K$  of them with the largest amplitudes. We should reconstruct  $\alpha_k$  by solving the least square minimization once more with the correctly identified  $K$  Dirac locations on the 2D plane.

### 2.5.3 Simulation results

We apply the algorithm to reconstruct 2D Diracs, which are located randomly in  $[-0.5, 0.5] \times [-0.5, 0.5]$ . The Dirac amplitudes follow a log-normal distribution:  $\alpha_k \sim \log \mathcal{N}(\log 2, 0.5)$ . The Fourier transform is sampled at  $L = 8500$  irregular frequencies, which are randomly generated with higher concentrations around low frequencies. This is because low frequency Fourier measurements correspond to the cross-correlations between antennas that are close to each other, a case that is more convenient in practice. We have chosen  $15 \times 15$  interpolation knots located uniformly on the area  $[-12\pi, 12\pi] \times [-12\pi, 12\pi]$ , where the Fourier transform is approximated. Complex-valued Gaussian white noise is added to the Fourier transforms so that the signal-to-noise ratio is 5 dB. The algorithm is able to reconstruct the Diracs correctly even in the presence of severe noise (Figure 2.10). With our current Python implementation (which can be further optimized), it takes 42 seconds for the reconstruction on average with a MacBook Pro laptop. Following [57], we also include the classical sparse recovery result obtained when the  $\ell_1$ -norm of the discretized sky image is minimized. As evidenced in Figure 2.10, this approach is not only less accurate than the FRI method, but it also fails to resolve Diracs that are closely located.

## 2.6 Conclusion

Motivated by the point source reconstruction problem in radio interferometry, we have developed a robust algorithmic framework for FRI reconstruction with arbitrary measurements, including

<sup>9</sup>It is possible to write  $\mathbf{R}(\mathbf{c})$  compactly as the Kronecker product between  $\tilde{\mathbf{R}}(\mathbf{c})$  and an identity matrix of size  $2\lfloor M\tau_1/2 \rfloor$ :  $\mathbf{R}(\mathbf{c}) \otimes \mathbf{I}_{2\lfloor M\tau_1/2 \rfloor}$ .

non-uniform sampling cases. We have unified all FRI-based methods concisely with a constrained formulation by establishing a linear relation between the given measurements and a set of unknown uniform samples of sinusoids. We have demonstrated the versatility of the proposed approach with various FRI signal recoveries in addition to an application to radio interferometry. The algorithm outperforms state of the art methods and is able to recover point sources accurately even in severe noise conditions. For future work, it would be interesting to consider an alternative (convex) formulation for FRI reconstructions, where a properly chosen atomic norm (see e.g., [123]) is minimized subject to a data-fitting constraint.

## 2.A Derivation of the equivalent formulation of (P1)

The constrained optimization for the reconstruction of FRI signals is

$$\begin{aligned} \min_{\mathbf{b}, \mathbf{c}} \quad & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to} \quad & \mathbf{R}(\mathbf{c})\mathbf{b} = \mathbf{0} \\ & \mathbf{c}_0^H \mathbf{c} = 1, \end{aligned} \tag{2.17}$$

where  $\mathbf{c}_0$  is a random initialization for the annihilating filter coefficients.

For a fixed  $\mathbf{c}$ , (2.17) is a constrained quadratic minimization with respect to  $\mathbf{b}$ . The associated Lagrangian is:

$$\mathcal{L}(\mathbf{b}, \ell) = \frac{1}{2} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 + \ell^H \mathbf{R}(\mathbf{c})\mathbf{b},$$

where  $\ell$  is the Lagrange multiplier. From the optimality conditions, we have:

$$\begin{cases} \mathbf{G}^H(\mathbf{G}\mathbf{b} - \mathbf{a}) + \mathbf{R}^H(\mathbf{c})\ell = \mathbf{0}, \\ \mathbf{R}(\mathbf{c})\mathbf{b} = \mathbf{0}. \end{cases} \tag{2.18}$$

Since  $\mathbf{G}$  has full column rank (see footnote 2), from (2.18) we have

$$\mathbf{b} = \boldsymbol{\beta} - (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}) \left( \mathbf{R}(\mathbf{c}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}) \right)^{-1} \mathbf{R}(\mathbf{c}) \boldsymbol{\beta}, \tag{2.19}$$

where  $\boldsymbol{\beta} = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \mathbf{a}$ .

We can substitute (2.19) into the objective function:

$$\begin{aligned} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 &= \boldsymbol{\beta}^H \mathbf{R}^H(\mathbf{c}) \left( \mathbf{R}(\mathbf{c}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}) \right)^{-1} \mathbf{R}(\mathbf{c}) \boldsymbol{\beta} \\ &\quad + \|\mathbf{a} - \mathbf{G}\boldsymbol{\beta}\|_2^2 \\ &\stackrel{(a)}{=} \mathbf{c}^H \mathbf{T}^H(\boldsymbol{\beta}) \left( \mathbf{R}(\mathbf{c}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}) \right)^{-1} \mathbf{T}(\boldsymbol{\beta}) \mathbf{c} \\ &\quad + \text{terms independent of } \mathbf{c}, \end{aligned}$$

where (a) results from the definition of the right dual matrix in Definition 2.1.



## 2.B Equivalent form for the solution of (2.3)

The Lagrangian associated with the constrained optimization (2.3) is

$$\mathcal{L}(\mathbf{c}, \lambda) = \frac{1}{2} \mathbf{c}^H \mathbf{T}^H(\boldsymbol{\beta}) (\mathbf{R}(\mathbf{c}_{n-1}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}_{n-1}))^{-1} \mathbf{T}(\boldsymbol{\beta}) \mathbf{c} + \lambda (\mathbf{c}_0^H \mathbf{c} - 1).$$

From the optimality conditions, we have

$$\begin{cases} \mathbf{T}^H(\boldsymbol{\beta}) (\mathbf{R}(\mathbf{c}_{n-1}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}_{n-1}))^{-1} \mathbf{T}(\boldsymbol{\beta}) \mathbf{c} + \lambda \mathbf{c}_0 = \mathbf{0}, \\ \mathbf{c}_0^H \mathbf{c} = 1. \end{cases} \quad (2.20)$$

Denote an auxiliary variable

$$\boldsymbol{\ell} = (\mathbf{R}(\mathbf{c}_{n-1}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}_{n-1}))^{-1} \mathbf{T}(\boldsymbol{\beta}) \mathbf{c},$$

then (2.20) is equivalent to

$$\begin{cases} \mathbf{T}^H(\boldsymbol{\beta}) \boldsymbol{\ell} + \lambda \mathbf{c}_0 = \mathbf{0}, \\ \mathbf{R}(\mathbf{c}_{n-1}) (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}_{n-1}) \boldsymbol{\ell} = \mathbf{T}(\boldsymbol{\beta}) \mathbf{c}, \\ \mathbf{c}_0^H \mathbf{c} = 1. \end{cases} \quad (2.21)$$

We can apply the same manipulation again by introducing another auxiliary variable  $\mathbf{v} = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}^H(\mathbf{c}_{n-1}) \boldsymbol{\ell}$ , then (2.21) is equivalent to:

$$\begin{cases} \mathbf{T}^H(\boldsymbol{\beta}) \boldsymbol{\ell} + \lambda \mathbf{c}_0 = \mathbf{0}, \\ \mathbf{R}(\mathbf{c}_{n-1}) \mathbf{v} = \mathbf{T}(\boldsymbol{\beta}) \mathbf{c}, \\ \mathbf{G}^H \mathbf{G} \mathbf{v} = \mathbf{R}^H(\mathbf{c}_{n-1}) \boldsymbol{\ell}, \\ \mathbf{c}_0^H \mathbf{c} = 1. \end{cases} \quad (2.22)$$

If we rearrange (2.22) in a matrix / vector form, we have (2.4).

Once we have the updated annihilating filter coefficients  $\mathbf{c}_n$ , (2.5) is obtained directly by rewriting the optimality conditions (2.18) with  $\mathbf{c} = \mathbf{c}_n$  as a linear system.

## 2.C Exact interpolation with Dirichlet kernel

Since  $x(t)$  has finite time support between  $-\frac{\tau}{2}$  and  $\frac{\tau}{2}$ , we can rewrite  $x(t)$  with its periodized version multiplied by a rectangular window:  $x(t) = \text{rect}(t/\tau) \sum_{n \in \mathbb{Z}} x(t - n\tau)$ . Hence, the Fourier

transform of  $x(t)$  is

$$\begin{aligned} X(\omega) &= \int_{-\infty}^{\infty} \text{rect}(t/\tau) \sum_{n \in \mathbb{Z}} x(t - n\tau) e^{-j\omega t} dt \stackrel{(a)}{=} \int_{-\infty}^{\infty} \text{rect}(t/\tau) e^{-j\omega t} \cdot \frac{1}{\tau} \sum_{m \in \mathbb{Z}} X\left(\frac{2\pi m}{\tau}\right) e^{j\frac{2\pi m}{\tau}t} dt \\ &= \sum_{m \in \mathbb{Z}} X\left(\frac{2\pi m}{\tau}\right) \int_{-\infty}^{\infty} \frac{1}{\tau} \text{rect}\left(\frac{t}{\tau}\right) e^{-j\left(\omega - \frac{2\pi m}{\tau}\right)t} dt = \sum_{m \in \mathbb{Z}} X\left(\frac{2\pi m}{\tau}\right) \text{sinc}\left(\frac{\tau}{2}\left(\omega - \frac{2\pi m}{\tau}\right)\right), \end{aligned} \quad (2.23)$$

where (a) is from the Poisson sum formula.

Further, from the periodicity of  $X(\omega)$ , we can rewrite the infinite summation in (2.23) as

$$\begin{aligned} X(\omega) &= \sum_{n \in \mathbb{Z}} \sum_{|m| \leq \lfloor \frac{M\tau}{2} \rfloor} X\left(\frac{2\pi m}{\tau}\right) \text{sinc}\left(\frac{\tau}{2}\left(\omega - \frac{2\pi(m + nM\tau)}{\tau}\right)\right) \\ &= \sum_{|m| \leq \lfloor \frac{M\tau}{2} \rfloor} X\left(\frac{2\pi m}{\tau}\right) \sum_{n \in \mathbb{Z}} \text{sinc}\left(\frac{\tau}{2}\left(\omega - \frac{2\pi(m + nM\tau)}{\tau}\right)\right) \\ &= \sum_{|m| \leq \lfloor \frac{M\tau}{2} \rfloor} X\left(\frac{2\pi m}{\tau}\right) \frac{\sin\left(\frac{\tau\omega - 2\pi m}{2}\right)}{M\tau \sin\left(\frac{\tau\omega - 2\pi m}{2M\tau}\right)}. \end{aligned}$$

## 2.D Spherical harmonics of Diracs on the sphere

Conventionally, spherical harmonics of degree  $l$  and order  $m$  is defined as  $Y_l^m(\theta, \varphi) = N_{l,m} P_l^{|m|}(\cos \theta) e^{jm\varphi}$ . Here the normalization factor  $N_{l,m} = (-1)^{(m+|m|)/2} \sqrt{\frac{2l+1}{4\pi} \frac{(l-|m|)!}{(l+|m|)!}}$ ; and  $P_l^{|m|}$  is the Legendre polynomial of degree  $l$  and order  $m$  ( $|m| \leq l$ ):

$$P_l^{|m|}(t) = (-1)^{|m|} (1-t^2)^{|m|/2} \frac{d^{|m|}}{dt^{|m|}} P_l(t),$$

where  $P_l(t) \stackrel{\text{def}}{=} \frac{1}{2^l l!} \frac{d^l}{dt^l} (t^2 - 1)^l$ . The spherical harmonic coefficient is given by the inner product between the signal and the spherical harmonics basis on  $\mathbb{S}^2$ :

$$\begin{aligned} \hat{I}_{l,m} &= \langle I(\theta, \varphi), Y_l^m(\theta, \varphi) \rangle \\ &= \int_0^{2\pi} \int_0^\pi I(\theta, \varphi) N_{l,m} P_l^{|m|}(\cos \theta) e^{-jm\varphi} \sin \theta d\theta d\varphi \\ &= N_{l,m} \sum_{k=1}^K \alpha_k P_l^{|m|}(\cos \theta_k) e^{-jm\varphi_k}. \end{aligned}$$

Note that the  $|m|$ -th order derivative of the Legendre polynomial  $P_l(t)$  is a polynomial of degree  $l - |m|$ . Hence, we may rewrite  $P_l^{|m|}$  in terms of canonical polynomial bases as:

$$P_l^{|m|}(t) = (-1)^{|m|} (1-t^2)^{|m|/2} \sum_{n=0}^{l-|m|} p_{l,n,|m|} t^n,$$

for some coefficients  $p_{n,|m|}$ , which are independent of where the polynomial is evaluated, and can be precomputed. Consequently the spherical harmonic coefficient of the Diracs is:

$$\hat{I}_{l,m} = N_{l,m} \sum_{n=0}^{l-|m|} p_{l,n,|m|} \sum_{k=1}^K \alpha_k (\cos \theta_k)^n (\sin \theta_k)^{|m|} e^{-jm\varphi_k}.$$

## 2.E Coefficients of the canonical basis in (2.13)

First we derive the Legendre polynomial coefficients with the canonical basis. From the definition

$$P_l(t) \stackrel{\text{def}}{=} \frac{1}{2^l l!} \frac{d^l}{dt^l} (t^2 - 1)^l = \sum_{n=0}^l \tilde{p}_{l,n} t^n,$$

we know that the term  $t^n$  can only come from the term with order  $l+n$  in the expansion of  $(t^2 - 1)^l$ . Also, the expansion will be in terms of  $t^2$ . Hence, we only have non-zero coefficients for terms such that  $l+n$  is *even*.

- if  $l$  is even, then take  $n = 2n'$  for  $0 \leq n' \leq \frac{l}{2}$ :

$$P_l(t) = \sum_{n'=0}^{l/2} \tilde{p}_{l,2n'} t^{2n'}.$$

The coefficient for  $t^{l+2n'}$  in the expansion of  $(t^2 - 1)^l$  is:

$$(-1)^{\frac{l}{2}-n'} \binom{l}{\frac{l}{2}+n'} = (-1)^{\frac{l}{2}-n'} \binom{l}{\frac{l}{2}-n'},$$

because we can choose  $t^2$   $(l+2n')/2 = l/2 + n'$  times among the  $l$  multiplying terms. Equivalently, we are choosing the constant term (i.e.,  $-1$ )  $l - (l/2 + n')$  times.

Then the  $l$ -th derivative will introduce another factor:

$$(l+2n')(l+2n'-1) \cdots (2n'+1) = \frac{(l+2n')!}{(2n')!}.$$

Therefore, the coefficients for the term  $t^{2n'}$  is:

$$\begin{aligned} \tilde{p}_{l,2n'} &= \frac{(-1)^{\frac{l}{2}-n'}}{2^l l!} \binom{l}{\frac{l}{2}-n'} \frac{(l+2n')!}{(2n')!} = \frac{(-1)^{\frac{l}{2}-n'}}{2^l} \binom{l}{\frac{l}{2}-n'} \frac{(l+2n')!}{(2n')! l!} \\ &= \frac{(-1)^{\frac{l}{2}-n'}}{2^l} \binom{l}{\frac{l}{2}-n'} \binom{l+2n'}{l} \quad \text{for } n' = 0, \dots, \frac{l}{2}. \end{aligned}$$

- if  $l$  is odd, then take  $n = 2n' + 1$  for  $0 \leq n' \leq \frac{l-1}{2}$ :

$$P_l(t) = \sum_{n'=0}^{l/2-1} \tilde{p}_{l,2n'+1} t^{2n'+1}.$$

The coefficients for  $t^{l+2n'+1}$  in the expansion of  $(t^2 - 1)^l$  is:

$$(-1)^{\frac{l-1}{2}-n'} \binom{l}{\frac{l+1}{2} + n'} = (-1)^{\frac{l-1}{2}-n'} \binom{l}{\frac{l-1}{2} - n'}.$$

From the  $l$ -th order derivative, we have an extra factor:

$$(l + 2n' + 1)(l + 2n') \cdots (2n' + 2) = \frac{(l + 2n' + 1)!}{(2n' + 1)!}.$$

The coefficients for the term  $t^{2n'+1}$  is:

$$\begin{aligned} \tilde{p}_{l,2n'+1} &= \frac{(-1)^{\frac{l-1}{2}-n'}}{2^l l!} \binom{l}{\frac{l-1}{2} - n'} \frac{(l + 2n' + 1)!}{(2n' + 1)!} = \frac{(-1)^{\frac{l-1}{2}-n'}}{2^l} \binom{l}{\frac{l-1}{2} - n'} \frac{(l + 2n' + 1)!}{(2n' + 1)! l!} \\ &= \frac{(-1)^{\frac{l-1}{2}-n'}}{2^l} \binom{l}{\frac{l-1}{2} - n'} \binom{l + 2n' + 1}{l} \quad \text{for } n' = 0, \dots, \frac{l-1}{2}. \end{aligned}$$

Observe that

$$\frac{d^{|m|}}{dt^{|m|}} P_l(t) = \sum_{n=0}^{\ell} \tilde{p}_{l,n} \frac{d^{|m|}}{dt^{|m|}} t^n = \sum_{n=|m|}^l \tilde{p}_{l,n} \underbrace{n(n-1) \cdots (n-|m|+1)}_{=\frac{n!}{(n-|m|)!}} t^{n-|m|}.$$

With change of variable:  $n \rightarrow n - |m|$ , we have

$$\frac{d^{|m|}}{dt^{|m|}} P_l(t) = \sum_{n=0}^{l-|m|} \underbrace{\tilde{p}_{l,n+|m|}}_{p_{l,n,|m|}} \frac{(n+|m|)!}{n!} t^n.$$

## Chapter 3

# Multi-dimensional FRI Sampling with Linear Sample Complexity\*

Plurality should not be posited without  
necessity.

---

WILLIAM OF OCKHAM

### 3.1 Introduction

Continuous-domain sparse recovery is a classic problem at the heart of various applications, like point source estimation in radio astronomy [101], microscopic image super-resolution [108], and direction of arrival estimation (DOA) with sensor arrays [118]. Typically, a low-pass filtered version of the sparse signal (or equivalently its low frequency spectrum) is observed in the measurement process. It is often the case that only a limited number of measurements is available—either due to the bandwidth of the instruments (as in microscopic imaging) or because of the finite number of sensors (as in array signal processing).

Therefore, it is essential to devise an efficient algorithm that estimates the sparse signal robustly from a set of discrete samples. The sampling framework for signals with finite rate of innovation (FRI) [133] is a natural candidate for this task. For instance, it has been shown that a periodic stream of Diracs in 1D, which consists of  $K$  Diracs within each period, is reconstructed exactly from  $2K+1$  ideal samples of the Dirac stream. The key to the FRI-based reconstruction is the annihilating filter method [103, 133], which allows the enforcement of the continuous domain sparsity constraint with a discrete convolution equation (see Section 1.2.1). In the previous

---

\*The material in this chapter is the result of joint work with Thierry Blu and Martin Vetterli [99].

**Table 3.1:** Sample complexities of separate and joint estimation approaches to reconstruct  $K$   $D$ -dimensional Diracs.

Approach	Separate estimation (Section 3.2.1)	Joint estimation with hypercubic sample shape (Corollary 3.1 and 3.4)	Joint estimation with flexible sample shape (Corollary 3.2 and 3.5)
minimum number of samples	$(K + 1)^D$	$2^D K$	$(D + 1)K$
sample complexity	$\mathcal{O}(K^D)$	$\mathcal{O}(K)$	$\mathcal{O}(K)$
degrees of freedom		$(D + 1)K$	

chapter, the FRI framework has been generalized to cope with *non-uniform* measurements [96], making the framework applicable to many practical problems, such as radio astronomy [98] and DOA estimation for arbitrary array layouts [97]. Previous attempts to extend the FRI-based sparse recovery to two dimensions led to various sampling schemes for polygons [29, 115], and curves [52, 90, 95].

The focus in this chapter is on sampling and reconstruction of high dimensional Diracs, which is directly related to point source estimation in various fields. Previous FRI-based approaches for two-dimensional Diracs [84] recasts the problem as two sub-problems, which estimate the Dirac locations along horizontal and vertical directions separately. The separate reconstruction approach requires an additional pairing step to combine the estimated Dirac coordinates. Additionally, the sample complexity is quadratic in 2D [84, 96].

Alternatively, as we propose in this chapter, we reconstruct multi-dimensional Diracs by enforcing the annihilation constraints in all dimensions simultaneously. The Dirac locations are then given by the intersections of hypersurfaces, which are the zero-crossings of multivariate polynomials with coefficients specified by the reconstructed annihilating filter coefficients. For instance, in 2D, the estimation of Dirac locations amounts to finding the intersections of two curves, which can be computed algebraically by finding the common roots of two polynomials (see Figure 3.1 and Section 3.2.1). Our main contributions in this chapter are:

- (i) The extension of the generic FRI sampling and reconstruction framework [96] to higher dimensions (two dimensions and above) — the resulting sampling scheme has *linear* sample complexity in signal sparsity (see e.g., Corollary 3.1). The reduced sample complexity is a significant improvement over our previous approach, where the multi-dimensional sparse recovery was reduced to several one-dimensional sub-problems and where, for a fixed number of Diracs, the minimum required sample size increases exponentially with respect to the dimension. We summarize sample complexities of the previous separate estimation and the new joint estimation approaches in Table 3.1.
- (ii) An iterative algorithm (similar to the one in [96]) to solve the multi-dimensional FRI reconstruction<sup>1</sup> (Section 3.2.2). The algorithm makes it possible to apply the multi-dimensional FRI-based sparse recovery to practical problems, where the measurements are not necessarily taken uniformly (see Chapter 4 for a concrete example in radio astronomy).

Before proceeding, we point to literature on sparse recovery in high-dimensions. Prony’s method [103], which is at the root of the FRI-based sparse recovery, also plays an essential

<sup>1</sup>The Python code will be available at <http://1cav.epfl.ch/people/hanjie.pan>

role in high-resolution spectrum estimation. Various attempts have been made to generalize the approach to high dimensions such as matrix enhancement and matrix pencil (MEMP) [67] and ESPRIT-like subspace methods [70, 105, 106, 109]. One drawback with MEMP is the necessity to have an additional pairing step in order to obtain the final multi-dimensional reconstruction. The sample complexity is quadratic in terms of signal sparsity in 2D [67]. In comparison, ESPRIT-like approaches exploit the shift-invariance property of several sub-matrices, which are jointly diagonalized, and achieve linear sample complexity [70, 109] (albeit the requirement of having at least 4 samples along each dimension). Recent efforts have been made to extend Prony's method to multi-dimensional settings [74, 102]. However, the sample complexity with the proposed approach therein scales exponentially in terms of the signal dimension [74]. Finally, recent work on spectral compressed sensing (e.g., [30]) is not directly related with the multi-dimensional Dirac reconstruction. Instead, the focus is on recovering the Fourier spectrum of the sparse signal, which is partially sub-sampled on a grid: the sparse signal parameters are retrieved with standard spectrum estimation algorithm, such as MEMP [67], once the full spectrum has been recovered.

The rest of the chapter is organized as follows. First, we adapt the generic FRI reconstruction framework developed in the previous chapter to address the two-dimensional sparse recovery in Section 3.2.1. Sample complexity as well as the reconstruction algorithm details are discussed in Section 3.2.2. Then, the FRI-based 2D sparse recovery is further generalized to higher dimensions in Section 3.2.3. Next, we validate the proposed approach with simulations in Section 3.3. The significance of the new approach will be further investigated in two applications on point source reconstruction in radio astronomy and direction of arrival estimation with microphone arrays in the following two chapters.

## 3.2 Methods

### 3.2.1 FRI reconstruction in 2D

In our previous approach [96], the multi-dimensional problem was reduced to several sub-problems by enforcing the annihilation constraints along each dimension separately (Section 3.2.1). We discuss the limitations associated with the separate estimation approach and propose an alternative formulation, where the annihilation constraints are enforced along all dimensions simultaneously (Section 3.2.1).

#### Separate estimation formulation

In [96], the multi-dimensional FRI reconstruction problem was recast as separate annihilation problems along each dimension. Take the reconstruction of  $K$  two-dimensional Diracs as an example. The Fourier transform on a uniform grid can be annihilated by two discrete filters of size  $K + 1$  along rows and columns, respectively:

$$\begin{aligned} \min_{\mathbf{c}_1 \in \mathcal{C}_{1,\mathbf{b}}} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 & \quad \text{and} \quad \min_{\mathbf{c}_2 \in \mathcal{C}_{2,\mathbf{b}}} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to } \mathbf{c}_1 * \mathbf{b} = \mathbf{0} & \quad \text{subject to } \mathbf{c}_2 * \mathbf{b} = \mathbf{0} \end{aligned} \quad (3.1)$$

Here

- $\mathbf{a}$  is the vector of the given measurements of the 2D Diracs, e.g., the ideally low-pass filtered

samples.

- $\mathbf{b}$  is the vector of the (unknown) uniformly sampled Fourier transform, e.g., the vectorized Fourier series coefficients of the 2D Diracs in a column-by-column order.
- $\mathbf{G}$  is the linear transformation that links the sinusoidal samples  $\mathbf{b}$  to the ideally low-pass filtered samples  $\mathbf{a}$ , e.g., the inverse DFT transformation.
- $\mathbf{c}_1$  and  $\mathbf{c}_2$  are the annihilating filters along the horizontal and vertical directions, respectively. They each belong to a certain space, e.g.,  $\mathcal{C}_1 = \left\{ \mathbf{c}_1 \in \mathbb{C}^{K+1} \mid (\mathbf{c}_1^{(0)})^H \mathbf{c}_1 = 1 \right\}$  where  $\mathbf{c}_1^{(0)}$  is a random initialization of  $\mathbf{c}_1$  in the iterative algorithm (see [96] for details).  $\mathcal{C}_2$  is similarly defined.

The  $x$ -coordinates (respectively  $y$ -coordinates) are then given by the roots of a univariate polynomial whose coefficients are specified by  $\mathbf{c}_1$  (respectively  $\mathbf{c}_2$ ).

All Diracs that are located on the  $K^2$  intersections of the reconstructed  $x$  and  $y$  coordinates (see Figure 3.1(a)), satisfy the annihilation constraints in both problems in (3.1). Therefore, it is necessary to find the correct correspondence among the possible  $x$  and  $y$  locations in order to reconstruct the 2D Diracs. The exhaustive approach amounts to considering  $\binom{K^2}{K}$  possible combinations, which becomes computationally infeasible with a large number of Diracs. In [96], we showed experimentally that we can identify the  $(x, y)$ -pairs by reconstructing amplitudes of Diracs, which are located at  $K^2$  possible  $x$  and  $y$  intersections, all at once with least square minimization. Then, the coordinates of the  $K$  Diracs (among the  $K^2$  possibilities) that have the largest amplitudes give the Dirac locations in 2D. Finally, we reconstruct the Dirac amplitudes with the  $K$  correctly identified 2D Dirac locations.

Despite the empirical success of such a simple strategy, there are two noticeable drawbacks:

- The restrictive shape of the 1D filters used in the separate estimation approach, requires that the uniform sinusoidal samples  $\mathbf{b}$  should be at least  $(K+1) \times (K+1)$ , which increases *quadratically* with respect to the number of Diracs. However, an optimal algorithm should only use  $\mathcal{O}(K)$  number of measurements. More generally, in  $D$  dimensions, the sample complexity of the separate estimation approach is  $\mathcal{O}(K^D)$ , i.e., a penalty of  $\mathcal{O}(K^{D-1})$  over the optimal linear sample complexity.
- The two 1D FRI reconstructions in (3.1) are solved independently, without enforcing the estimated uniform sinusoidal samples  $\mathbf{b}$  to be the same. But we know that there is always a feasible solution (i.e., the noiseless  $\mathbf{b}$ ) that satisfies both annihilation constraints in the two 1D problems and fits the given measurements up to the noise level. Additionally, since  $\mathbf{b}$  should be the same in both problems, any linear combination of the two filters  $\mathbf{c}_1$  and  $\mathbf{c}_2$  is also a valid solution. Therefore, instead of estimating two annihilating filters separately, what we should reconstruct is a set of (two) linear independent vectors that spans the null space of the convolution matrix associated with  $\mathbf{b}$ .

Because of the limitations above, we now propose a new joint estimation approach.



### Joint estimation formulation

Instead of solving two 1D Dirac estimations separately, a more natural choice for the 2D FRI reconstruction is:

$$\begin{aligned} & \min_{\substack{\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}, \\ \mathbf{b}}} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ & \text{subject to } \mathbf{c}_1 * \mathbf{b} = \mathbf{0} \\ & \qquad \qquad \mathbf{c}_2 * \mathbf{b} = \mathbf{0}. \end{aligned} \quad (3.2)$$

The goal is to estimate a set of uniform sinusoidal samples  $\mathbf{b}$ , which fits the noisy measurements (up to the noise level) and two annihilating filters  $\mathbf{c}_1$ ,  $\mathbf{c}_2$  that form bases of the null space of the convolution matrix specified by  $\mathbf{b}$ . Note that here  $\mathbf{c}_1$  and  $\mathbf{c}_2$  are, in general, 2D non-separable filters. Here,  $\mathcal{C}$  specifies a few<sup>2</sup> linear constraints such that the solution to (3.2) corresponds to one (out of many possible) set of annihilating filters that spans the null space of the convolution matrix specified by  $\mathbf{b}$ . One possible choice is  $\mathcal{C} = \{\mathbf{C}|\mathbf{C}_0^T\mathbf{C} + \mathbf{C}^T\mathbf{C}_0 = 2\mathbf{I}\}$ , where  $\mathbf{C}$  and  $\mathbf{C}_0$  are 2-column matrices, with columns specified by the two annihilating filters  $\mathbf{c}_1$ ,  $\mathbf{c}_2$  and their corresponding initializations  $\mathbf{c}_1^{(0)}$ ,  $\mathbf{c}_2^{(0)}$ , respectively.

A necessary condition in order to satisfy the annihilation equations is that the DTFT of each of the two annihilating filters  $\mathbf{c}_1$  and  $\mathbf{c}_2$

$$\mu(x, y) = \sum_{k,l} c_{k,l} (e^{-j2\pi x/\tau_1})^k (e^{-j2\pi y/\tau_2})^l, \quad (3.3)$$

should vanish at the Dirac locations  $(x_k, y_k)$  for  $k = 1, \dots, K$ :

$$\mu(x_k, y_k) = 0.$$

Here  $c_{k,l} = [\mathbf{c}_i]_{k,l}$  for  $i = 1$  or  $2$  and  $\tau_1$ ,  $\tau_2$  are the periods of the 2D Dirac stream along  $x$  and  $y$  directions, respectively.

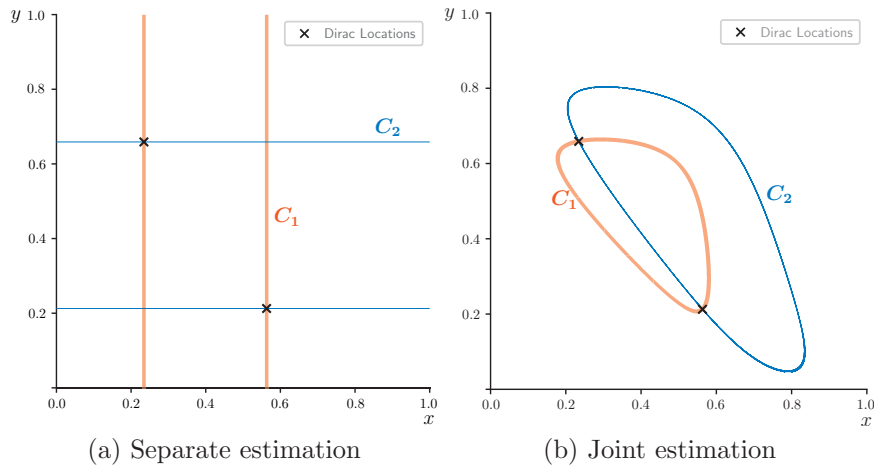
The corresponding FRI signal (i.e., the zero-crossings of a bivariate polynomial (3.3) with respect to  $e^{-j2\pi x/\tau_1}$  and  $e^{-j2\pi y/\tau_2}$ ) is a curve [95]. Any Dirac located on this curve satisfies the annihilation constraints in (3.2). The 2D Dirac locations are then obtained from the intersections of the two curves (Figure 3.1).

A straight-forward way to determine the intersections of curves, is to evaluate both bivariate polynomials (with coefficients specified by  $\mathbf{c}_1$  and  $\mathbf{c}_2$ ) on a grid: with a finer grid, more accurate Dirac locations can be determined. Alternatively, we will discuss in details an algebraic approach in Section 3.2.2, which treats the curve intersections as the common roots of two polynomials. This approach allows us to determine the curve intersections reliably without resorting to grid search.

#### 3.2.2 Algorithm and implementation

We give algorithmic details related to the joint estimation introduced in the previous section. We first derive the sample complexity with respect to the number of Diracs in Section 3.2.2. Next, an algebraic approach is outlined in Section 3.2.2, which finds curve intersections based on the Bézout resultant of a polynomial system. Finally, an iterative algorithm to solve the joint annihilation problem (3.2) is presented in Section 3.2.2.

<sup>2</sup>In general,  $(D+1)D/2$  constraints are needed to uniquely specify a set of orthogonal vectors that spans a vector space of dimension  $D$ .



**Figure 3.1:** Joint estimation of 2D Diracs based on curve intersections. The Dirac locations are given by the intersections of two curves  $C_1$  and  $C_2$  (see text in Section 3.2.1). (a) The zero-crossings of two univariate polynomials specified by two 1D annihilating filters in the separate estimation. (b) The zero-crossings of two (bivariate) polynomials specified by general shaped 2D annihilating filters in the joint estimation.

### Sample complexity of 2D Dirac estimation

Conventionally, a  $(K + 1)$ -tap annihilating filter is used in order to recover  $K$  Diracs in 1D [133]. In higher dimensions, e.g., 2D, it is not always possible to have a general 2D filter with a size matched to the number of Diracs  $K$  (except in the degenerated cases where the 2D filter reduces to a 1D filter). We then need to specify the support size (i.e., the number of non-zero entries) for a given choice of a 2D filter shape.

#### Proposition 3.1 (*Filter support size*)

Suppose the two annihilating filters in (3.2) have the same support, then the minimum support size is  $K + 2$  for the estimation of  $K$  Diracs.

*Proof.*

See Appendix 3.A. □

Because the annihilating filter is scale-invariant with respect to any non-zero scalar, each filter has one degree of freedom less than the total number of non-zero elements. Consequently, at least  $2(K + 1)$  annihilation equations are needed to reconstruct  $K$  Diracs.

In a typical setup, an equal number of samples is available along each direction. Under this assumption, a filter shape, which is as square as possible, e.g.,  $\lceil \sqrt{K + 2} \rceil \times \lceil \sqrt{K + 2} \rceil$ , allows the maximum number of annihilation equations to be built<sup>3</sup>. Since the ideally low-pass filtered samples have a direct correspondence with the uniform sinusoidal samples, we can derive the sample complexity by comparing the total number of equations (from the annihilation constraint

<sup>3</sup>For cases where the sample sizes differ along each direction, we may choose the filter shapes accordingly such that the ratio of the filter sizes along each direction is the same with that of the sample sizes.

and the normalization of the filter coefficients  $\mathcal{C}_1$  and  $\mathcal{C}_2$ ) with the unknowns (i.e., the filter coefficients).

**Corollary 3.1 (*Sample complexity*)**

Assuming that the ideally low-pass filtered sample size is  $M \times M$ , where  $M$  is an odd number<sup>4</sup>. In order to reconstruct  $K$  2D Diracs, the sample size should satisfy

$$M \geq \sqrt{K+2} + \sqrt{K} - 1. \quad (3.4)$$

*Proof.*

See Appendix 3.B. □

From (3.4), the minimum number of samples  $M^2$  has a *linear* complexity  $\mathcal{O}(K)$  with respect to the number of Diracs. In comparison, the sample complexity is  $\mathcal{O}(K^2)$  in the previous approach, where the annihilation constraints were enforced separately along each direction (see Section 3.2.1). For  $K$  2D Diracs, the degrees of freedom are  $3K$ :  $2K$  for the Dirac locations and  $K$  for the amplitudes. The total degrees of freedom give the lowest possible sample complexity of *any* reconstruction algorithm. It is possible to reach this lower bound with the proposed joint estimation if the sample sizes are not required to be the same along each dimension. We summarize the result with a simplified setup, where uniformly sampled Fourier transform of the Diracs are directly available.

**Corollary 3.2**

For simplicity, let the number of Diracs  $K$  be an even number, then  $K$  2D Diracs can be reconstructed from  $2 \times (K/2 + K)$  uniform Fourier transform samples of the Diracs.

*Proof.*

The proof follows similarly to that of Corollary 3.1 by choosing annihilating filters of shape  $2 \times (K/2 + 1)$ . □


We show the achievable sample complexity with the separate and the proposed joint estimation approaches against the total degrees of freedom in Figure 3.2.

The merits of an algorithm with lower sample complexity are two-fold: Not only can more Diracs be reconstructed for a given sample size, but also a more robust reconstruction may be expected — Because of more efficient data utilization, more annihilation equations are available to further constrain the reconstruction. We demonstrate both improvements with simulations in Section 3.3.1 and Section 3.3.2.

Finally, we point out that the proposed joint annihilation approach can also cope with the singular cases where some Diracs have the same  $x$  or  $y$  coordinates. In these cases, the minimum size of the annihilating filter is further constrained by the number of Diracs that have common  $x$  or  $y$  locations. Please refer to Appendix 3.C for detailed discussions and Section 3.3.3 for an example.

## Polynomial common-root finding

Instead of resorting to grid search for finding the curve intersections, an alternative approach is to re-cast the geometric problem as finding the common roots of two polynomials. Since the



**Figure 3.2:** Sample complexities (ideal low-pass filtering with equal sizes along each dimension) of the proposed joint and the previous separate estimation approaches. The joint estimation approach is able to estimate  $K$  Diracs from  $3K$  measurements if the sample sizes are not restricted to be the same along each dimension (see Corollary 3.1 and Corollary 3.2).

mask function (3.3) is a bivariate polynomial with respect to  $u = e^{-j2\pi x/\tau_1}$  and  $v = e^{-j2\pi y/\tau_2}$ , the curve intersections are the solutions of:

$$\begin{cases} 0 = \mu_1(u, v) = \sum_{k,l} [\mathbf{c}_1]_{k,l} u^k v^l, \\ 0 = \mu_2(u, v) = \sum_{k,l} [\mathbf{c}_2]_{k,l} u^k v^l. \end{cases} \quad (3.5)$$

One way to determine the polynomial common roots is to compute the Bézout resultant [26, 120] of the polynomial system of equations. It is based on the following two observations:

- (i) The bivariate polynomials in (3.5) can be rearranged with respect to one variable, e.g.,  $u$ , where the polynomial coefficients are functions of the other variable, e.g.,  $v$ :

$$\mu_1(u, v) = \sum_{k=0}^{K_1-1} p_k(v) u^k \quad \text{and} \quad \mu_2(u, v) = \sum_{k=0}^{K_2-1} q_k(v) u^k.$$

Here  $p_k(\cdot)$  and  $q_k(\cdot)$  are some polynomials of  $v$ , and the filters  $\mathbf{c}_1$  and  $\mathbf{c}_2$  in (3.5) are of size  $L_1 \times K_1$  and  $L_2 \times K_2$ , respectively.

- (ii) For any  $n \geq 0$ , we have

$$u^n \mu_1(u, v) = u^n \mu_2(u, v) = 0.$$

From (i) and (ii), we can build a linear system of equations in terms of a vector  $[u^0, \dots, u^{K_1+K_2-3}]^T$ :

$$\underbrace{\begin{bmatrix} p_0 & \cdots & p_{K_1-1} & 0 & \cdots \\ 0 & p_0 & \cdots & p_{K_1-1} & \ddots \\ \vdots & \ddots & \ddots & & \ddots \\ q_0 & \cdots & q_{K_2-1} & 0 & \cdots \\ 0 & q_0 & \cdots & q_{K_2-1} & \ddots \\ \vdots & \ddots & \ddots & & \ddots \end{bmatrix}}_{\mathbf{A}(v)} \begin{bmatrix} u^0 \\ \vdots \\ u^{K_1+K_2-3} \end{bmatrix} = \mathbf{0}. \quad (3.6)$$

A necessary and sufficient condition for the polynomial systems (3.5) to have common roots is that the Bézout resultant should vanish (see [36] Theorem 7):

$$\det(\mathbf{A}(v)) = 0. \quad (3.7)$$

Since (3.7) is a polynomial equation of  $v$  with constant coefficients, the common roots  $v_k \stackrel{\text{def}}{=} e^{-j\frac{2\pi}{\tau_2}y_k}$  are then given by the roots of this univariate polynomial. Once we have computed  $v_k$ ,  $u_k \stackrel{\text{def}}{=} e^{-j\frac{2\pi}{\tau_1}x_k}$  is obtained from (3.5) with  $v = v_k$ . Unlike the separate estimation cases, an extra pairing step (see text in Section 3.2.1) is no longer needed: indeed, from the common root finding procedure, the Dirac locations in 2D are directly reconstructed with the correct association.

From (3.6), the degree of the polynomial in (3.7) is:

$$(K_2 - 1) \max_k (\deg(p_k(v))) + (K_1 - 1) \max_k (\deg(q_k(v))). \quad (3.8)$$

Consequently, there are at least<sup>5</sup> a number of common roots given by (3.8). In general, this may be larger than the number of Diracs  $K$ . The reason is that in the root-finding procedure, we proceed as if the common roots were arbitrary complex numbers. However, since our goal is to reconstruct multi-dimensional Diracs (which have real-valued coordinates, i.e.,  $x_k, y_k \in \mathbb{R}$ ), the roots are necessarily located on the unit circle<sup>6</sup>. One way to eliminate invalid solutions for the Dirac reconstruction (i.e., the  $u_k, v_k$  that leads to complex-valued  $x_k, y_k$ ) is based on the moduli of the common roots. However, it is not always obvious how to choose an adequate threshold level on the moduli of the roots in noisy scenarios. Alternatively, an approach that works empirically is as follows: all roots are first projected onto the unit circle  $u_k/|u_k|, v_k/|v_k|$ , from which the Dirac locations are reconstructed. Then, amplitudes for all Diracs are estimated by minimizing the discrepancies between the re-synthesized and given measurements in the least square sense. The valid Dirac locations then correspond to the ones that have the  $K$  largest amplitudes. We then solve the least square minimization once more with the extracted  $K$  Dirac locations (see [96] for a similar treatment in the separate annihilation of 2D Dirac cases).

<sup>5</sup>There could be multiple roots for  $u$  with a given  $v = v_k$ .

<sup>6</sup>With the strategy used here, any damping factors are eliminated, which could be important in another related (but different) spectrum estimation problem [118]. Further investigation is required to apply the proposed approach to general spectrum estimation in future work.

### Iterative algorithm for the joint annihilation

Following the notation in the previous chapter, let  $\mathbf{T}(\cdot)$  be the operator that builds the Toeplitz matrix from the input data and the associated right dual  $\mathbf{R}(\cdot)$  such that  $\mathbf{R}(\mathbf{c})\mathbf{b} = \mathbf{T}(\mathbf{b})\mathbf{c}$ ,  $\forall \mathbf{c}, \mathbf{b}$ . Then the joint annihilation constraints in (3.2) amount to vertically stacking two right dual matrices:

$$\begin{bmatrix} \mathbf{R}(\mathbf{c}_1) \\ \mathbf{R}(\mathbf{c}_2) \end{bmatrix} \mathbf{b} = \mathbf{0}. \quad (3.9)$$

For a set of uniform sinusoidal samples  $\mathbf{b}$  with a sufficiently large size (compared with that of the filters), it may appear that we could have *more* equations than unknowns. For instance, if the shape of both filters  $\mathbf{c}_1$  and  $\mathbf{c}_2$  is  $2 \times 2$  and that of the sinusoidal samples  $\mathbf{b}$  is  $5 \times 5$ , then the total number of annihilation equations is 32 (i.e.,  $4 \cdot 4 + 4 \cdot 4$ ), which is larger than the size of  $\mathbf{b}$  (i.e., 25 in this example). If we solve (3.2) with the annihilation constraint (3.9) for any fixed  $\mathbf{c}_1$  and  $\mathbf{c}_2$  different from zero, then it appears that  $\mathbf{b}$  would be forced to be identically zero. But as we will show next, the vertically stacked convolution matrices in (3.9) do not have full rank.

#### Corollary 3.3 (*Null space dimension*)

Define the joint annihilation matrix as

$$\mathbf{R}_{\text{joint}} \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{R}(\mathbf{c}_1) \\ \mathbf{R}(\mathbf{c}_2) \end{bmatrix},$$

then the null space dimension of  $\mathbf{R}_{\text{joint}}$  is at least  $2(K_0 - 1)(L_0 - 1)$  for filters  $\mathbf{c}_1$  and  $\mathbf{c}_2$  of size  $L_0 \times K_0$ .

*Proof.*

See Appendix 3.D. □

In terms of implementation, this requires the extraction of the independent rows of  $\mathbf{R}_{\text{joint}}$  when (3.2) is minimized with respect to  $\mathbf{b}$  (for a certain choice of  $\mathbf{c}_1$  and  $\mathbf{c}_2$ ) at each iteration. One possibility is to resort to the QR-decomposition<sup>7</sup> of the joint annihilation matrix. Let  $\mathbf{R}_{\text{joint}} = \mathbf{Q}\mathbf{U}$ , where  $\mathbf{Q}$  is an orthogonal matrix and  $\mathbf{U}$  is an upper triangle matrix, and  $\mathbf{Q}_{\text{sub}}$  extracts all but the last  $(K_2 - 1)(L_1 - 1) + (K_1 - 1)(L_2 - 1)$  columns of  $\mathbf{Q}$ . Then, the joint annihilation matrix  $\mathbf{R}_{\text{joint}}$  is replaced with  $\mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}}$ , which is an orthogonal projection onto the row space of  $\mathbf{R}_{\text{joint}}$ . Similarly, we can express the two *linearly independent* annihilating filters with their *effective* degrees of freedom, which are less than their total sizes. Specifically, the iterative algorithm (see details in Appendix 3.E) involves the convolution between  $\boldsymbol{\beta} \stackrel{\text{def}}{=} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \mathbf{a}$  and the two annihilating filters. Let the QR-decomposition of  $\mathbf{T}(\boldsymbol{\beta})^H$  be  $\tilde{\mathbf{Q}}\tilde{\mathbf{U}}$ , then the convolutions can be equivalently written as:

$$\begin{cases} \mathbf{T}(\boldsymbol{\beta})\mathbf{c}_1 = \mathbf{0} \\ \mathbf{T}(\boldsymbol{\beta})\mathbf{c}_2 = \mathbf{0} \end{cases} \Leftrightarrow \begin{cases} \mathbf{T}(\boldsymbol{\beta})\tilde{\mathbf{Q}}_1\boldsymbol{\gamma}_1 = \mathbf{0} \\ \mathbf{T}(\boldsymbol{\beta})\tilde{\mathbf{Q}}_2\boldsymbol{\gamma}_2 = \mathbf{0} \end{cases},$$

for two linearly independent vectors<sup>8</sup>  $\mathbf{c}_1 = \tilde{\mathbf{Q}}_1\boldsymbol{\gamma}_1$  and  $\mathbf{c}_2 = \tilde{\mathbf{Q}}_2\boldsymbol{\gamma}_2$ . Here  $\tilde{\mathbf{Q}}_1 = \tilde{\mathbf{Q}}$  and  $\tilde{\mathbf{Q}}_2$  is a

<sup>7</sup>Alternatively, we can use the singular value decomposition, which is computationally more demanding.

<sup>8</sup>In  $D$  dimensions,  $D$  linearly independent annihilating filters are represented as  $\mathbf{c}_i = \tilde{\mathbf{Q}}_i\boldsymbol{\gamma}_i$  for  $i = 1, \dots, D$ , where  $\tilde{\mathbf{Q}}_i$  is a sub-matrix built from all but  $i - 1$  columns of  $\tilde{\mathbf{Q}}$ .

sub-matrix of  $\tilde{\mathbf{Q}}$  that consists of all but one column of  $\tilde{\mathbf{Q}}$ . In the algorithm, only  $\gamma_1$  and  $\gamma_2$  are reconstructed.

An equivalent formulation of (3.2) (expressed in terms of  $\gamma_1$  and  $\gamma_2$ ) is derived in Appendix 3.E, based on which an iterative algorithm (similar to that in [96]) is proposed. For numerical stability, we may introduce auxiliary variables such that no nested matrix inverses are involved at each iteration. It can be shown<sup>9</sup> that at each iteration, the updated  $\gamma_1$  and  $\gamma_2$  are solutions of an extended linear system of equations:

$$\begin{bmatrix} \mathbf{0} & \mathbf{T}_{\text{joint}}(\beta)^{\text{H}} \mathbf{Q}_{\text{sub}} & \mathbf{0} & \mathbf{\Gamma}_0^* \\ \mathbf{Q}_{\text{sub}}^{\text{H}} \mathbf{T}_{\text{joint}}(\beta) & \mathbf{0} & -\mathbf{Q}_{\text{sub}}^{\text{H}} \mathbf{R}_{\text{joint}} & \mathbf{0} \\ \mathbf{0} & -\mathbf{R}_{\text{joint}}^{\text{H}} \mathbf{Q}_{\text{sub}} & \mathbf{G}^{\text{H}} \mathbf{G} & \mathbf{0} \\ \mathbf{\Gamma}_0^{\text{T}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\gamma} \\ \boldsymbol{\ell} \\ \mathbf{v} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \boldsymbol{\rho} \end{bmatrix}, \quad (3.10)$$

where  $\boldsymbol{\ell}$ ,  $\mathbf{v}$ ,  $\boldsymbol{\lambda}$  are auxiliary variables,  $\boldsymbol{\gamma} = [\gamma_1^{\text{T}}, \gamma_2^{\text{T}}]^{\text{T}}$  and the constant vector  $\boldsymbol{\rho} = [1, 1, 0]^{\text{T}}$ . Here  $\mathbf{\Gamma}_0$  is a 3-column matrix specified by initializations of the two annihilating filters:

$$\mathbf{\Gamma}_0 = \begin{bmatrix} \tilde{\mathbf{Q}}_1^{\text{T}} \mathbf{c}_1^{(0)} & \mathbf{0} & \tilde{\mathbf{Q}}_1^{\text{T}} \mathbf{c}_2^{(0)} \\ \mathbf{0} & \tilde{\mathbf{Q}}_2^{\text{T}} \mathbf{c}_2^{(0)} & \tilde{\mathbf{Q}}_2^{\text{T}} \mathbf{c}_1^{(0)} \end{bmatrix},$$

and

$$\mathbf{T}_{\text{joint}}(\beta) = \begin{bmatrix} \mathbf{T}(\beta) \tilde{\mathbf{Q}}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{T}(\beta) \tilde{\mathbf{Q}}_2 \end{bmatrix}.$$

The uniform sinusoidal samples  $\mathbf{b}$  are updated as the solution of

$$\begin{bmatrix} \mathbf{G}^{\text{H}} \mathbf{G} & \mathbf{R}_{\text{joint}}^{\text{H}} \mathbf{Q}_{\text{sub}} \\ \mathbf{Q}_{\text{sub}}^{\text{H}} \mathbf{R}_{\text{joint}} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ \boldsymbol{\ell} \end{bmatrix} = \begin{bmatrix} \mathbf{G}^{\text{H}} \mathbf{a} \\ \mathbf{0} \end{bmatrix}, \quad (3.11)$$

where  $\boldsymbol{\ell}$  is an auxiliary variable. We summarize the joint estimation algorithm in Algorithm 2.

### 3.2.3 Generalization to higher dimensions

In the previous two sections, we detailed the problem formulation and reconstruction algorithm for the joint annihilation problem in 2D. The approach can be generalized to higher dimensions following the same line of reasoning. Specifically, for Dirac reconstruction in  $D$ -dimensions, the joint estimation amounts to solving

$$\begin{aligned} & \min_{\substack{\mathbf{c}_1, \dots, \mathbf{c}_D \in \mathcal{C}, \\ \mathbf{b}}} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ & \text{subject to } \mathbf{c}_i * \mathbf{b} = \mathbf{0}, \text{ for } i = 1, \dots, D, \end{aligned} \quad (3.12)$$

where  $\mathcal{C}$  is similarly defined as in the 2D cases. The  $D$ -dimensional Dirac locations are the common roots of  $D$  polynomials whose coefficients are specified by the annihilating filters  $\mathbf{c}_i$  for  $i = 1, \dots, D$ . Geometrically, the Dirac locations are the intersections of  $D$  hypersurfaces, which are the zero-crossings of these  $D$ -variate polynomials.

One way to compute the common roots of multiple polynomials is to collect all polynomials with respect to one variable and compute the Bézout resultants of polynomial pairs, which

<sup>9</sup>We omit the derivation for brevity, which are very similar to the ones in Appendix 2.A and Appendix 2.B in the previous chapter.

**Algorithm 2:** Joint estimation of 2D Diracs

---

**Input** : Measurements of the 2D Diracs  $\mathbf{a}$ , the linear transformation matrix  $\mathbf{G}$ ,  
*(optional)* noise level  $\varepsilon$ .

**Output:** annihilating filters  $\mathbf{c}_1, \mathbf{c}_2$ , uniform sinusoidal samples  $\mathbf{b}$ .

**if**  $\varepsilon$  not given **then**  
|  $\varepsilon = 0$   
**end**

**for**  $loop \leftarrow 1$  **to** *max. initializations* **do**

1 | Initialize  $\mathbf{c}_1$  with  $\mathbf{c}_1^{(0)}$ ,  $\mathbf{c}_2$  with  $\mathbf{c}_2^{(0)}$ ;  
| **for**  $n \leftarrow 1$  **to** *max. iterations* **do**

2 |     Build  $\mathbf{R}_{\text{joint}}$  from  $\mathbf{c}_1^{(n-1)}$  and  $\mathbf{c}_2^{(n-1)}$  and computes its QR-decomposition;

3 |     Build the linear system (3.10) and update  $\mathbf{c}_1^{(n)}$  and  $\mathbf{c}_2^{(n)}$  as  $\tilde{\mathbf{Q}}_1 \gamma_1$  and  $\tilde{\mathbf{Q}}_2 \gamma_2$ ;

4 |     Build the linear system (3.11) and update  $\mathbf{b}^{(n)}$  from its solution;

5 |     **if**  $\|\mathbf{a} - \mathbf{G}\mathbf{b}^{(n)}\|_2^2 \leq \varepsilon^2$  **then**  
|     | Terminate both loops;  
|     **end**

**end**

**end**

6  $\mathbf{b} \leftarrow \mathbf{b}^{(n)}$  and  $\mathbf{c} \leftarrow \mathbf{c}^{(n)}$ .

---


themselves are polynomials of the other variables. The elimination process is then repeated by computing the resultant of resultants (which depends on one less variable) until the resultant reduces to an univariate polynomial, whose roots can be computed directly.

Concretely, take the 3D case as an example. Let  $\mu_1(x, y, z)$ ,  $\mu_2(x, y, z)$ ,  $\mu_3(x, y, z)$  be the polynomials specified by the annihilating filters and  $\text{res}_s(\cdot, \cdot)$  be the resultant of two polynomials with respect to  $s$ . Then  $\text{res}_z(\mu_1, \mu_2)$  and  $\text{res}_z(\mu_2, \mu_3)$  are bivariate polynomials of  $x$  and  $y$ . The common roots in  $x$  is given by the roots of the univariate polynomial  $\text{res}_y(\text{res}_z(\mu_1, \mu_2), \text{res}_z(\mu_2, \mu_3))$ . Such a recursive approach may not be numerically stable, especially for polynomials of higher degree (i.e., more Diracs  $K$ ). Further work is needed to find alternatives in determining the common roots of multiple polynomials. One possible idea could be based on [91], where the canonical polynomial basis is replaced with Chebyshev polynomials.

For simplicity, if we assume that the sample size is the same along each dimension, then by choosing annihilating filters of size<sup>10</sup>  $\lceil (K + D)^{1/D} \rceil$  along each dimension, the maximum number of annihilation equations can be built from the given samples. The sample complexity compared with the number of Diracs is summarized as follows.

<sup>10</sup>This is because the filter support size is  $K + D$  in  $D$ -dimension Dirac reconstruction. See Appendix 3.A for details.





**Figure 3.3:** Illustration of sample and annihilating filter shapes in the critical sampling cases in 3D (see Corollary 3.5).

**Corollary 3.4 (*D*-dimensional sample complexity)**

Assuming that the ideally low-pass filtered sample size is  $\underbrace{M \times \cdots \times M}_{D \text{ terms}}$ , then the sample size  $M$ , which is an odd number, should satisfy

$$M \geq (K + D)^{1/D} + K^{1/D} - 1, \quad (3.13)$$

in order to reconstruct  $K$   $D$ -dimensional Diracs.

The minimum total number of samples required  $M^D$  is *linear* in terms of the number of Diracs  $K$ . For  $K \gg D$ , the sample complexity can be approximated as  $M^D \geq 2^D K$ . Even though the sample complexity is still proportional to the number of Diracs, the proportion factor grows exponentially with respect to the dimension. In comparison, the total degrees of freedom of  $K$   $D$ -dimensional Diracs are  $(D + 1)K$ . The gap is bridged as soon as sample sizes are not required to be the same along all dimensions. We summarize the result as follows.

**Corollary 3.5**

For simplicity, let the number of Diracs  $K$  be an integer multiple of the dimension  $D$ , then  $K$   $D$ -dimensional Diracs can be reconstructed from  $(D + 1)K$  uniformly sampled Fourier transform of the Diracs, which are supported at indices  $\mathbf{k} = (k_1, \dots, k_D)$  such that

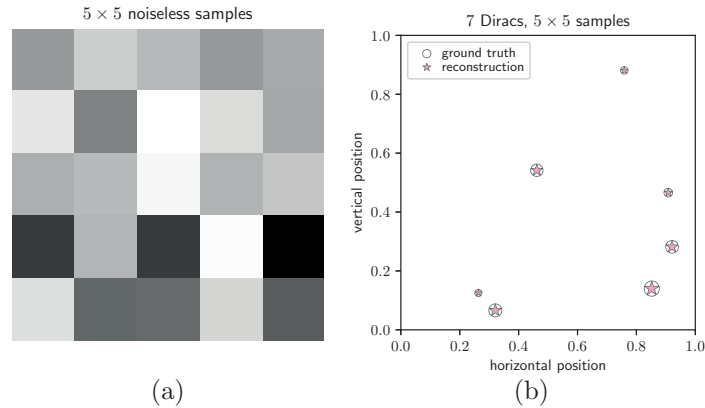
- $k_1 = 0, \dots, K/D + K - 1$ ;
- $k_j = 0, 1$  for  $j \neq 1$  and  $\sum_{j=2}^D k_j \leq 1$  (i.e., at most one  $j$  such that  $k_j$  can be non-zero).

*Proof.*

The proof is straightforward by choosing a filter that has the same support with the Fourier samples  $k_j$  for  $j = 2, \dots, D$  and  $k_1 = 0, \dots, K/D$ .  $\square$

Pictorially, the samples and filters consist of  $D$  parallel “lines” along dimension  $k_1$ . We illustrate this with an example in 3D (Figure 3.3). Further investigation is required to design a sampling scheme such that these uniform Fourier transforms are accessible from a given set of spatial domain samples of the same size.

The joint estimation approach in higher dimensions is exemplified with a 3D Dirac reconstruction example with noiseless and noisy measurements in Section 3.3.4.



**Figure 3.4:** Exact reconstruction of 7 Diracs on a 2D plane from  $5 \times 5$  ideally low-pass filtered samples. (a) The noiseless ideally low-pass filtered samples. (b) The exact reconstruction with the proposed joint estimation. The joint annihilation approach utilizes the given measurements efficiently. In comparison, the separate annihilation approach fails to reconstruct any Dirac, because it requires a minimum sample size  $8 \times 8$  in this case.

### 3.3 Results


In this section, we validate the proposed joint annihilation algorithm with synthetic experiments for the reconstruction of multi-dimensional Diracs. Diracs are ideally low-pass filtered and uniformly sampled with equal sample sizes along each dimension. We first demonstrate the improvements of the joint estimation approach in 2D (Section 3.3.1 to Section 3.3.3). The generality of the proposed method is further exemplified with a 3D simulation in Section 3.3.4.

In all simulations, 50 random initializations and 20 maximum iterations per initialization are used in the joint estimation algorithm (Algorithm 2).

#### 3.3.1 Efficient data utilization

In our previous approach, where the annihilation constraints were enforced separately, the sample size along each dimension should at least be the size of the 1D annihilating filter. The restrictive filter shape limits the applicability of the separate estimation approach. Such an artificial constraint is lifted with the new approach: as long as the sample size satisfies (3.4) in 2D or more generally (3.13) in  $D$  dimensions, Diracs can be successfully estimated.

We demonstrate the efficiency of the joint estimation approach with a periodic stream of Diracs in 2D. Within each period  $[0, 1) \times [0, 1)$ , 7 Diracs of random amplitudes and locations are ideally low-pass filtered and uniformly sampled. Even with sample size as small as  $5 \times 5$ , the joint annihilation approach manages to reconstruct Diracs exactly (Figure 3.4). Both annihilating filters used here are of size  $3 \times 3$ . In comparison, the separate estimation (3.1) has to use two 1D filters of length 8 and hence requires a minimum sample size  $8 \times 8$ . Because of the restriction on the filter shapes, the previous separate approach fails to reconstruct any Dirac from such a limited number of samples.



**Figure 3.5:** Average reconstruction error of 2D Diracs against different noise levels with the joint and separate estimation approach, respectively. The reconstruction error at each noise level is averaged over 3000 different noise realizations (number of Diracs: 7, ideally low-pass filtered sample size:  $9 \times 9$ ).

### 3.3.2 Robust reconstruction

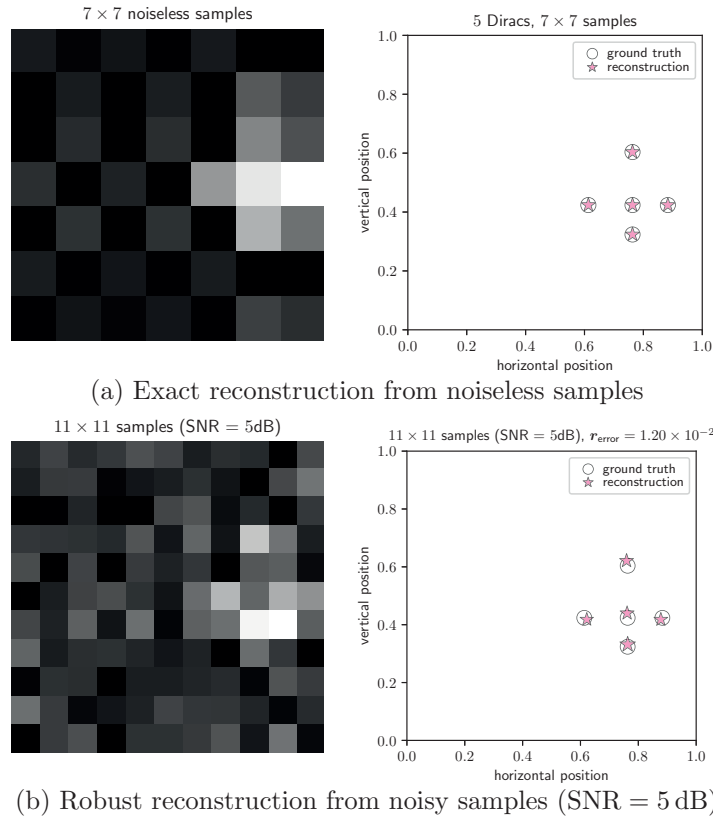
Another improvement, as a consequence of the flexible annihilating filter shapes, is that more annihilation equations can be built from the given measurements. These extra equations further constrain the reconstruction, leading to a more robust estimation in the presence of noise.

To quantify these improvements experimentally, we consider a periodic stream of Diracs, which consists of 7 Diracs with random locations and amplitudes within each period. The ideally low-pass filtered and uniformly sampled measurements (sample size:  $9 \times 9$ ) are contaminated with Gaussian white noise such that the signal to noise ratio (SNR) varies from  $-10$  dB to  $40$  dB. We measure the reconstruction quality with the average distance<sup>11</sup> between the estimated and ground truth Dirac locations. The performances of both the joint and separate estimation approaches are averaged over 3000 different noise realizations at each noise level. The joint estimation approach reaches the Cramér-Rao lower bound down to around  $12$  dB. In comparison, the separate estimation does not reach the lower bound even in very high SNR cases (Figure 3.5).

### 3.3.3 Dirac with common $x$ / $y$ coordinates

In specific cases where some Diracs share common  $x$  or  $y$  coordinates, the annihilating filter shapes are further constrained as a consequence of the fundamental theorem of algebra (see details in Appendix 3.C). Similar to examples in the previous two sections, we consider 2D periodic stream of Diracs with unit period. Within each period, 5 Diracs of unitary amplitudes are simulated such that they are arranged in a dagger-shape (the right column in Figure 3.6). From our analysis in Section 3.2.2, the minimum annihilating filter shape is  $\lceil \sqrt{7} \rceil \times \lceil \sqrt{7} \rceil$  (i.e.,  $3 \times 3$ ) in the generic cases where no Diracs share common coordinates along any direction. However, in this specific case, there could be up to 3 Diracs that share the same  $x$  (or  $y$ )

<sup>11</sup>The pairing between the reconstruction and the ground truth is identified by permuting the Dirac locations such that the distance between the permuted reconstruction and the ground truth Dirac locations is minimized.



**Figure 3.6:** Reconstruction of 2D Diracs, some of which have common  $x$  or  $y$  coordinates, with the joint estimation (3.2). (a) The exact reconstruction from the noiseless ideally low-pass filtered samples (sample size:  $7 \times 7$ ). (b) Robust estimation of 2D Diracs from noisy samples (sample size:  $11 \times 11$ , SNR = 5 dB, reconstruction error on the Dirac locations:  $1.2 \times 10^{-2}$ ).

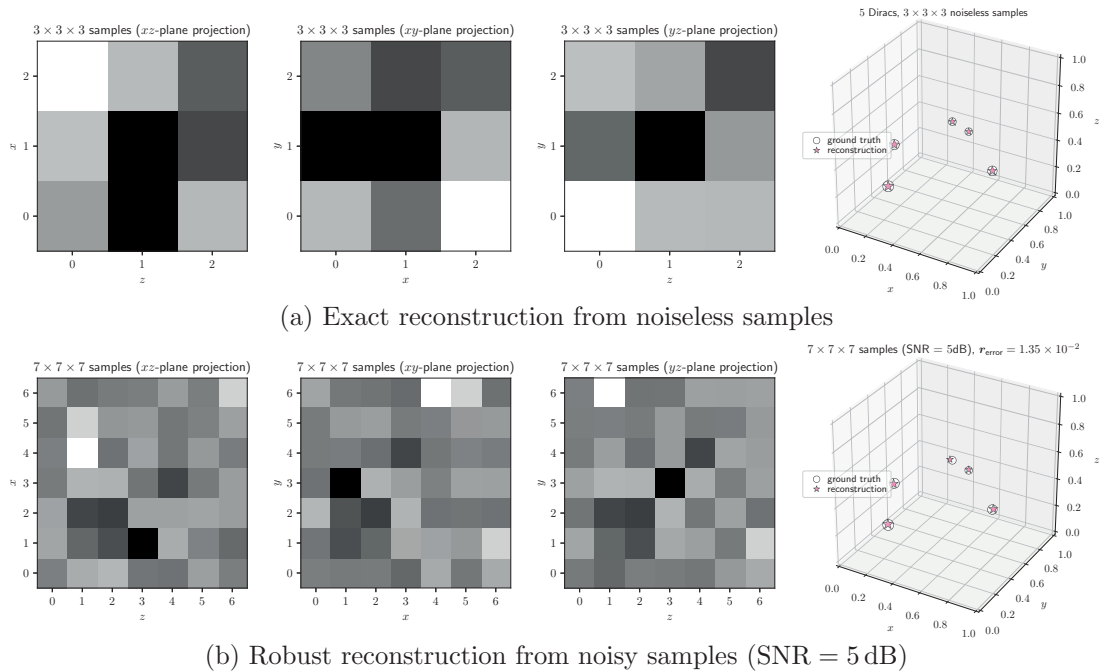
coordinates. From our analysis in Appendix 3.C, the annihilating filter shape, therefore, should be at least  $4 \times 4$ . The joint estimation of the 2D Diracs from  $7 \times 7$  noiseless ideally low-pass filtered samples, are shown in Figure 3.6(a). Additionally, we also considered the case with noisy samples (size:  $11 \times 11$ ). Gaussian white noise is added to the noiseless samples such that the signal to noise ratio is 5 dB. The reconstruction error on the Dirac locations is  $1.2 \times 10^{-2}$ .

### 3.3.4 Dirac reconstruction in 3D

To demonstrate the generality of the joint estimation approach to higher dimensions (see Section 3.2.3), we show an example of a periodic stream of Diracs in 3D. Within each period (1 along each dimension), 5 Diracs with random locations and amplitudes are simulated.

In the first case, the Diracs are ideally low-pass filtered and sampled. The joint estimation reconstructs 5 Diracs exactly from the minimum number of noiseless samples based on Corollary 3.4, i.e.,  $3 \times 3 \times 3$  (Figure 3.7(a)).

Next, we consider the same 3D Diracs but the samples are contaminated by Gaussian white noise such that the signal to noise ratio is 5 dB. In this noisy scenario, the Diracs are over-



**Figure 3.7:** Reconstruction of 3D periodic stream of Diracs from ideally low-pass filtered samples. (a) Exact reconstruction from the noiseless samples (sample size:  $3 \times 3 \times 3$ ). (b) Robust reconstruction from noisy samples (sample size:  $7 \times 7 \times 7$ , SNR = 5 dB, average reconstruction error on the Dirac locations:  $1.35 \times 10^{-2}$ ). Maximum intensity projection is used to project the 3D samples onto relevant planes for display.

sampled (compared with the minimum sample sizes). The average reconstruction error on the Dirac locations from  $7 \times 7 \times 7$  noisy samples is  $1.35 \times 10^{-2}$  (Figure 3.7(b)).

### 3.4 Conclusion

We have proposed an efficient FRI-based sampling scheme for multi-dimensional Dirac estimation. The minimum sample size required, in order to have the exact reconstruction, scales linearly with the number of Diracs. The linear sample complexity is a substantial improvement over previous separate estimation method, which has an exponential scaling in the space dimension. We have adapted the algorithmic framework for the reconstruction of FRI signals [96] to estimate multi-dimensional Diracs reliably. The Dirac locations are determined by the intersections of hypersurfaces, which are the zero-crossings of the reconstructed mask functions. The efficiency and robustness of the proposed approach has been validated in simulations.

Because of more efficient data utilization as well as a more robust performance, the joint estimation approach of FRI signals may have significant implications for many real applications. In the following two chapters, we adapt the continuous-domain sparse recovery techniques to the challenging point source estimation problem in radio astronomy (Chapter 4) and direction of arrival estimation problem in acoustics (Chapter 5).

### 3.A Support size of a 2D annihilating filter

For brevity of derivations, we use multi-index notation, where  $\mathbf{z} = (z_1, z_2) = (e^{-j2\pi x/\tau_1}, e^{-j2\pi y/\tau_2})$ ,  $\mathbf{k} = (k, l)$ , and  $\mathbf{z}^{\mathbf{k}} = (z_1^k, z_2^l)$ . Let the support size of a 2D annihilating filter  $\mathbf{c}$  be  $S$  and the non-zero indices in the filter be  $\mathbf{k}_s = (k_s, l_s)$  for  $s = 1, \dots, S$ .

From the annihilation constraint, we have that the mask function should vanish at each one of the  $K$  Dirac locations:

$$\mu(\mathbf{z}^{(i)}) = \sum_{s=1}^S (\mathbf{z}^{(i)})^{\mathbf{k}_s} \mathbf{c}_{\mathbf{k}_s} = 0,$$

for  $i = 1, \dots, K$ , i.e.,

$$\underbrace{\begin{bmatrix} (\mathbf{z}^{(1)})^{\mathbf{k}_1} & \dots & (\mathbf{z}^{(1)})^{\mathbf{k}_S} \\ \vdots & \dots & \vdots \\ (\mathbf{z}^{(K)})^{\mathbf{k}_1} & \dots & (\mathbf{z}^{(K)})^{\mathbf{k}_S} \end{bmatrix}}_{\text{a } K \times S \text{ matrix } \mathbf{Z}} \begin{bmatrix} \mathbf{c}_{\mathbf{k}_1} \\ \vdots \\ \mathbf{c}_{\mathbf{k}_S} \end{bmatrix} = \mathbf{0}.$$

In general when  $S \geq K$ , the matrix  $\mathbf{Z}$  is of rank  $K$  (i.e., the values  $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(K)}$  for which this does not happen is of zero measure). Therefore, a necessary (and sufficient) condition to have two linearly independent solutions  $\mathbf{c}_1$  and  $\mathbf{c}_2$  with the same support in (3.2) is  $S \geq K + 2$ .

The derivation generalizes to higher dimensions following the same reasoning in a straightforward manner. Specifically, the support size of a  $D$ -dimensional annihilating filter in a joint estimation formation (3.12) is at least  $K + D$ .

### 3.B Sample complexity of Dirac reconstruction in 2D

On the one hand,  $2(M - \sqrt{K + 2} + 1)^2$  annihilation equations can be built from  $M \times M$  ideally low-pass filtered samples in addition to 3 (because of symmetry) linear constraints from the normalization  $\mathcal{C} = \{\mathbf{C} | \mathbf{C}_0^T \mathbf{C} + \mathbf{C}^T \mathbf{C}_0 = 2\mathbf{I}\}$ . On the other hand, the number of unknowns is  $2(K + 2) - 1$ , where<sup>12</sup>  $-1$  is a consequence of linear independence of  $\mathbf{c}_1$  and  $\mathbf{c}_2$  (see Section 3.2.2 for implementation details). The minimum number of samples is then obtained from

$$2(M - \sqrt{K + 2} + 1)^2 + 3 \geq 2(K + 2) - 1.$$

### 3.C 2D Diracs with shared $x$ or $y$ coordinates

In the particular cases, where some of the Diracs share the same  $x$  or  $y$  coordinates, the minimum size of the annihilating filter shape is further restricted by the number of Diracs with common coordinates along each direction. This can be understood from the fundamental theorem of algebra.

Suppose there are  $K_{\text{same}}$  Diracs that have the same  $x$  coordinate  $x_0$  and the shape of the filter is  $P \times Q$ . From the annihilation constraint, the DTFT of the annihilating filter should

<sup>12</sup>More generally in  $D$  dimensions, the degrees of freedom of  $D$  vectors are reduced by  $D(D - 1)/2$  because of linear independence.

vanish at the Dirac locations:

$$\mu(x_0, y) = \sum_{l=0}^{P-1} \left( \sum_{k=0}^{Q-1} c_{k,l} (e^{-j2\pi x_0/\tau_1})^k \right) (e^{-j2\pi y/\tau_2})^l = 0. \quad (3.14)$$

Observe that (3.14) is a polynomial in  $e^{-j2\pi y/\tau_2}$  of degree  $P - 1$ . In order to have  $K_{\text{same}}$  such  $y$ -s within the period  $[0, \tau_2)$  that satisfy the annihilation constraint,  $P - 1 \geq K_{\text{same}}$ . Similarly, if there are  $L_{\text{same}}$  Diracs that have the same  $y$  coordinate, then  $Q - 1 \geq L_{\text{same}}$ .

In general, this leads to annihilating filters with different shapes than  $\lceil \sqrt{K+2} \rceil \times \lceil \sqrt{K+2} \rceil$  in the generic cases discussed in Section 3.2.2. Consequently, the minimum number of samples required in order to reconstruct  $K$  Diracs, where some of them have shared coordinates along one direction, is also larger than that in the generic cases in Corollary 3.1.

### 3.D Null space dimension of the joint annihilation matrix

In order to show the null space dimension of the joint annihilation matrix, we construct explicitly a non-zero vector  $\mathbf{b}$  such that  $\mathbf{R}_{\text{joint}} \mathbf{b} = \mathbf{0}$ . The number of linearly independent  $\mathbf{b}$ -s gives a lower bound of the null space dimension.

Construct  $b_{k,l} = u^{-k}v^{-l}$ , where  $u, v \in \mathbb{C}$  and are different from zero ( $\mathbf{b}$  is the vectorization of the 2D data block in e.g., a column-by-column order), then

$$[c_{k,l} * b_{k,l}]_{m,n} = \sum_{k,l} c_{k,l} b_{m-k, n-l} = u^{-m}v^{-n} \sum_{k,l} c_{k,l} u^k v^l.$$

Therefore,  $b_{k,l}$  that satisfies the joint annihilation constraints

$$\begin{cases} [\mathbf{c}_1]_{k,l} * b_{k,l} = 0 \\ [\mathbf{c}_2]_{k,l} * b_{k,l} = 0 \end{cases},$$

can be treated as the common roots of the two polynomial linear systems of equations (3.5). From (3.8), we know that there are at least  $2(K_0 - 1)(L_0 - 1)$  such  $\mathbf{b}$ -s, for filters  $\mathbf{c}_1, \mathbf{c}_2$  of size  $L_0 \times K_0$ . Consequently, we have proved the minimum null space dimension of the joint annihilation matrix in Corollary 3.3.

### 3.E Equivalent formulation of the joint annihilation (3.2)

From our analysis in Section 3.2.2, for fixed  $\mathbf{c}_1, \mathbf{c}_2$  at each iteration, (3.2) is equivalent to

$$\begin{aligned} \min_{\mathbf{b}} \quad & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to} \quad & \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} \mathbf{b} = \mathbf{0}. \end{aligned}$$

The associated Lagrangian is

$$\mathcal{L}(\mathbf{b}, \ell) = \frac{1}{2} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 + \ell^H \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} \mathbf{b}.$$

From the optimality conditions, we have

$$\begin{cases} \mathbf{G}^H(\mathbf{G}\mathbf{b} - \mathbf{a}) + \mathbf{R}_{\text{joint}}^H \mathbf{Q}_{\text{sub}} \ell = \mathbf{0}, \\ \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} \mathbf{b} = \mathbf{0}. \end{cases}$$

Hence,

$$\mathbf{b} = \boldsymbol{\beta} - (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}_{\text{joint}}^H \mathbf{Q}_{\text{sub}} \cdot \left( \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}_{\text{joint}}^H \mathbf{Q}_{\text{sub}} \right)^{-1} \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} \boldsymbol{\beta},$$

where  $\boldsymbol{\beta} = (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H \mathbf{a}$ . Consequently, the objective function reduces to

$$\begin{aligned} \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 &= \boldsymbol{\beta}^H \mathbf{R}_{\text{joint}}^H \mathbf{Q}_{\text{sub}} \left( \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}_{\text{joint}}^H \mathbf{Q}_{\text{sub}} \right)^{-1} \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} \boldsymbol{\beta} \\ &\quad + \text{terms independent of } \mathbf{c}_1 \text{ and } \mathbf{c}_2. \end{aligned}$$

Additionally, from the ‘‘bi-linearity’’ of the annihilation constraint, we have:

$$\begin{bmatrix} \mathbf{R}(\mathbf{c}_1) \\ \mathbf{R}(\mathbf{c}_2) \end{bmatrix} \boldsymbol{\beta} = \begin{bmatrix} \mathbf{T}(\boldsymbol{\beta}) & \mathbf{0} \\ \mathbf{0} & \mathbf{T}(\boldsymbol{\beta}) \end{bmatrix} \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{T}(\boldsymbol{\beta}) \tilde{\mathbf{Q}}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{T}(\boldsymbol{\beta}) \tilde{\mathbf{Q}}_2 \end{bmatrix}}_{\stackrel{\text{def}}{=} \mathbf{T}_{\text{joint}}(\boldsymbol{\beta})} \begin{bmatrix} \boldsymbol{\gamma}_1 \\ \boldsymbol{\gamma}_2 \end{bmatrix}.$$

Therefore, an equivalent form of the joint annihilation (3.2) is

$$\begin{aligned} \min_{\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2} & \begin{bmatrix} \boldsymbol{\gamma}_1 \\ \boldsymbol{\gamma}_2 \end{bmatrix}^H \boldsymbol{\Lambda}(\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2) \begin{bmatrix} \boldsymbol{\gamma}_1 \\ \boldsymbol{\gamma}_2 \end{bmatrix}, \\ \text{subject to} & \quad \mathbf{c}_1 = \tilde{\mathbf{Q}}_1 \boldsymbol{\gamma}_1 \in \mathcal{C}, \\ & \quad \mathbf{c}_2 = \tilde{\mathbf{Q}}_2 \boldsymbol{\gamma}_2 \in \mathcal{C}, \end{aligned}$$

where

$$\boldsymbol{\Lambda}(\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2) = \mathbf{T}_{\text{joint}}^H(\boldsymbol{\beta}) \mathbf{Q}_{\text{sub}} \left( \mathbf{Q}_{\text{sub}}^H \mathbf{R}_{\text{joint}} (\mathbf{G}^H \mathbf{G})^{-1} \mathbf{R}_{\text{joint}}^H \mathbf{Q}_{\text{sub}} \right)^{-1} \mathbf{Q}_{\text{sub}}^H \mathbf{T}_{\text{joint}}(\boldsymbol{\beta}).$$

The iterative strategy proposed in [96] then amounts to building  $\boldsymbol{\Lambda}(\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2)$  from the *previous* reconstruction and updating  $\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2$  from the quadratic minimization.

Since  $\boldsymbol{\Lambda}(\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2)$  contains nested matrix inverses, the direct implementation may not be numerically stable. By introducing auxiliary variables as in [96], it can be shown that the update of  $\boldsymbol{\gamma}_1$  and  $\boldsymbol{\gamma}_2$  amounts to solving an extended linear system of equation (3.10) at each iteration.



## Chapter 4

# Looking beyond Pixels with Continuous-space Estimation of Point Sources<sup>\*</sup>

Astronomy compels the soul to look upwards  
and leads us from this world to another.

---

*The Republic*

PLATO


### 4.1 Introduction

Existing radio interferometric imaging algorithms are discrete in nature, e.g., CLEAN [66] and its numerous variants [12, 37] or the compressed sensing inspired methods proposed in [25, 39, 117, 141]. As such, they estimate the locations and intensities of celestial sources on a uniform grid that is artificially imposed over the field of view [112].

Sources do not line up so conveniently in reality. Instead, they are located instead in-between pixels of the pre-defined grid. This leads to inaccurate source position and intensity estimation, with contributions from closely located sources being merged together into a single pixel (see Figure 4.1 for an illustration). Depending on the ultimate goal (e.g. calibration), it may be desired to have more accurate location estimates (and thus distances between objects) than achievable on a grid.

---

<sup>\*</sup>The material in this chapter is the result of joint work with Matthieu Simeoni, Paul Hurley, Thierry Blu and Martin Vetterli [98].



**Figure 4.1:** Existing imaging algorithms estimate the locations and intensities of celestial sources on a uniform grid. In practice, sources do not line up so conveniently, and can fall off-grid. Gridding hence results in a less accurate estimation of the location estimations as well as a potential overestimation of intensities due to multiple sources being merged to the same pixel.

The starting point for this chapter was hence to see if we could accurately determine the intensities and locations of sources directly from visibility data without a grid imposition in an intermediate image domain. The framework commonly referred to as *finite rate of innovation (FRI)* sampling is a natural candidate for this task. Introduced first in the signal processing community, FRI sampling generalizes the Shannon sampling theorem to *sparse* non-bandlimited signals. For example, [133] proposed a sampling scheme permitting the exact recovery of a stream of Dirac from a number of Fourier series coefficients proportional to the number of Diracs. The framework has since been applied successfully in other fields, and extended to 2D signals as well as to noisy measurements [85, 94, 95, 115]. Having been originally designed to work only with equally spaced Fourier samples as input, [96] extended the FRI framework to cases with non-uniform samples (as is the case in radio interferometry). It thus becomes possible to envisage an FRI-based approach in radio astronomy, albeit with the substantial remaining challenge of recovering a large number of sources given very weak signals and massive data sets.

A continuously defined framework such as FRI, allows the significance of the notion of achievable angular resolution to be revisited. Indeed, for the traditional CLEAN-based imaging pipeline, the *resolution power* of the tool, or the ability to distinguish neighboring sources, is limited by the width of the synthesized beam, whose width is inversely proportional to the longest baseline of the interferometer. Sources closer than this critical beam width are indistinguishable from one another.

In comparison, the FRI-based sparse recovery allows sources separated by a distance smaller

than this apparent bound to be distinguished. The continuous-domain formulation makes the performance only dependent on the number of measurements and sources in the sky in addition to the noise level in the measurements, but not on the number of pixels from an arbitrarily imposed, and potentially very large, grid. We name the proposed FRI-based approach as *Looking beyond pixels with continuous-space Estimation of Point sources* (LEAP).

Compressed sensing [e.g., 117], while it surpasses the instrument resolution limit as well, does rely on a grid. LEAP differs in that the estimation of source locations is decoupled from the estimation of their intensities. Hence, it is possible to exploit the consistency in source locations among different frequency sub-bands and have a coherent reconstruction in a multi-band setting (see Section 4.4.3).

The present chapter quantifies how LEAP can be successfully applied to recover point sources in realistic radio astronomy conditions. Our experiments, carried out through simulation and actual interferometric measurements from LOFAR [131], show that the reconstruction is more accurate and requires fewer measurements, reaching a comparable source estimate to CLEAN from much less integration time and smaller baselines. The achievable position accuracy goes beyond the perceived angular resolution, which allows closely located sources to be reliably distinguished. To confirm that these super-resolved sources were indeed actual sources, we showed that CLEAN could also recover them given longer baselines (and hence sharpening its PSF). Finally, we showed that LEAP could leverage together the information from multiple frequency bands in a coherent fashion to improve point source estimation.

The rest of this chapter is organized as follows. After a briefly review of the radio interferometer measurement equation in Section 4.2, we propose to adapt the sparse recovery framework based on FRI sampling (Section 4.3) to source estimations in radioastronomy in Section 4.4. The algorithmic details for the reconstruction of the source locations and intensities are discussed in Section 4.4.1 and Section 4.4.2. Further, we present the multi-band formulation in Section 4.4.3. The method is validated with both synthetic experiments (Section 4.5.2 to Section 4.5.4), actual LOFAR observations from the Boötes field (Section 4.5.5), and the “Toothbrush” cluster (Section 4.5.6), respectively. We discuss the advantages and limitations in Section 4.6 before concluding the work with a few possible future directions in Section 4.7.

## 4.2 Interferometric imaging measurement equation

A typical radio interferometer consists of an array of antennas that collect the electromagnetic waves emitted by celestial sources. In the far field context, these sources are assumed to be located on a hypothetical celestial sphere and the emitted electromagnetic waves arrive at each antenna in parallel. Consequently, the signals received at two antennas differ only by a geometric time delay, which is determined by the baseline of the antenna pair and the observation frequency. When the field-of-view is sufficiently narrow, the celestial sphere can be approximated locally by a tangential plane. It can be shown that the visibility measurements  $V_{ij}$ , given by the cross-correlations of antenna pairs  $(i, j)$ , then correspond to a 2D Fourier domain (conventionally referred to as  $(u, v)$ -domain) sampling of the sky brightness<sup>1</sup> distribution  $I$  [116, 125, 126]:

$$V_{ij} = \iint_{\mathbb{R}^2} I(\mathbf{r}) e^{-j2\pi\langle \mathbf{r}, \Delta \mathbf{p}_{ij} \rangle} d^2\mathbf{r}, \quad \forall i, j = 1, \dots, L. \quad (4.1)$$

<sup>1</sup>In this chapter, we use the term brightness and intensity interchangeably.

Here,  $L$  is the total number of antennas forming the interferometer;  $\mathbf{r} = (l, m)$  are the spatial coordinates of the sky image; and  $\Delta\mathbf{p} = (\mathbf{p}_i - \mathbf{p}_j)/\lambda := (u_{ij}/2\pi, v_{ij}/2\pi)$  is the projection onto the tangent plane of the baseline between antenna  $i$  and  $j$ , normalized by the wavelength  $\lambda$  of the received electromagnetic waves. For simplicity, we assumed antennas have uniform gains and omni-directional primary beams in (4.1). The  $w$ -term [see 37] is considered as a constant for all baselines in a sufficiently small field-of-view. However, the proposed algorithm can straightforwardly be extended to more complex data models such as the ones considered in [116].

The measurement equation (4.1) is known as the van Cittert-Zernike theorem [126, Chapter 3]. It establishes an approximate Fourier relationship between interferometric measurements and the sky brightness distribution: the visibilities  $V_{ij}$  are *samples* of the Fourier transform of the sky image at discrete frequencies  $(u_{ij}, v_{ij})$ . For a given antenna layout, a radio interferometer has finite number of possible baselines. Hence, it can only have a partial Fourier domain coverage. By exploiting the earth rotation, a wider  $uv$  coverage can be achieved, which sharpens the point-spread-function, and hence improves the resolution, of the various reconstruction algorithms.

In a modern radio telescope, the number of antennas can be enormous (e.g., around 20,000 dipole antennas in LOFAR). In that case, it becomes unrealistic to send all raw data collected by the antenna arrays to a centralized correlator and obtain visibility measurements for each antenna pair. One commonly used strategy for data compression involves grouping antennas as stations and applying beamforming to antenna signals within each station. The visibility measurements are then obtained by taking the cross-correlations at the station level. Specifically, the beamformed visibility measurement from two stations  $(i, j)$  are:

$$V_{ij} = \sum_{p=1}^Q \sum_{q=1}^Q w_p^{(i)} \bar{w}_q^{(j)} \iint_{\mathbb{R}^2} I(\mathbf{r}) e^{-j2\pi \langle \mathbf{r}, \Delta\mathbf{p}_{pq}^{(i,j)} \rangle} d^2\mathbf{r}, \quad \forall i, j = 1, \dots, L. \quad (4.2)$$

Here,  $L$  is the total number of stations,  $Q$  is the number of antennas per station;  $\Delta\mathbf{p}_{pq}^{(i,j)} = (\mathbf{p}_p^{(i)} - \mathbf{p}_q^{(j)})/\lambda$  is the normalized baseline for the pair formed by the  $p$ -th antenna in station  $i$  and the  $q$ -th antenna in station  $j$ ; and  $w_q^{(i)}$  are beamforming weights. A typical beamforming strategy is matched beamforming, which amounts to choosing  $w_p^{(i)} = \frac{1}{\sqrt{Q}} e^{j2\pi \langle \mathbf{r}_0, \mathbf{p}_p^{(i)} \rangle}$ , with  $\mathbf{r}_0$  the focus direction of the matched beamforming (e.g., the zenith).

When digital beamforming is performed, either the beamshapes at each station must be accounted for—rendering the conventional imaging pipeline computationally more intensive—or neglected with potentially strong consequences for image accuracy [124]. The latter is the strategy most commonly employed in off-the-shelf CLEAN implementations [92]. In contrast, as we will show in Section 4.4, the proposed approach readily accounts for the beamformed visibilities with minimal effort.

In summary, LEAP can transparently work with either measurement equation (4.1) (without beamforming), or (4.2) (factoring in beamforming). To this end, we estimate the Fourier transform of the sky image on a (resolution-independent) uniform grid (e.g.,  $55 \times 55$  in a typical setup for radioastronomy) from the visibility measurements (4.2). The point source locations and amplitudes are subsequently obtained from these uniform Fourier transform reconstructions with the FRI sampling technique, which we review now in the next section.

### 4.3 Continuous domain sparse recovery with FRI sampling

In this section, we briefly revisit the continuous-domain sparse recovery approach based on the finite rate of innovation sampling. This technique will be adapted to solve point source estimation in radio astronomy in the next section (Section 4.4).

We generalized the classic FRI framework to cases with *non-uniform* samples in the previous two chapters, making it applicable to point source reconstruction in radio astronomy. There, generic FRI reconstruction is recast as an approximation problem, where one would like to recover an FRI signal consistent with the given measurements. The re-synthesized measurements (based on the reconstructed FRI signal) should match the given (noisy) measurements up to the noise level. A valid solution to the approximation problem is obtained with the help of a constrained optimization, where the fitting error (e.g., the  $\ell_2$  norm of the discrepancies) is minimized subject to the annihilation constraint.

An efficient algorithm [96] was proposed to solve the constrained optimization iteratively, where an  $\ell \times \ell$  linear system of equations was solved at each iteration for a set of uniform sinusoidal samples  $\mathbf{b}$  of size  $\ell$ . The simplicity of the algorithm is beneficial for point source reconstructions in radio interferometric imaging. The recovery estimates point sources in the *continuous domain* directly from visibilities, and the complexity depends only on the dimension of  $\mathbf{b}$  (typically around a few thousand). In terms of computational complexity, solving a dimension  $\ell$  linear system of equations is at most  $\mathcal{O}(\ell^3)$  [see 61, Chapter 3]. In contrast, CLEAN or compressed sensing based approaches have to estimate an intermediate sky image defined on a grid first before applying local peak detections in order to identify point sources. Consequently, the complexity of these algorithms is related to the size of the *discrete* image (around one million pixels or more), which is significantly larger than the dimension of the uniform sinusoidal samples  $\mathbf{b}$  in a typical setup (see Section 4.5.5 for a concrete example).

Although the focus in this chapter is on point source reconstructions, the FRI-based approach can also deal with extended source recovery. Given a suitable set of bases in which the extended sources have a sparse representation, the same algorithm can be applied in the transform domain. However, substantially more work would be required to design a continuous domain “sparsifying” transformation for celestial sources [see 117, for examples in a discrete setup], and hence this is left for future work.

### 4.4 Algorithm

In the previous section, we reviewed the generic form of an FRI-based sparse reconstruction. In this section, we adapt this continuous-domain sparse recovery framework to point source reconstructions in radio astronomy. The FRI-based approach estimates source locations first (Section 4.4.1) before solving a least square minimization for the source intensities (Section 4.4.2). A multi-band formulation, which may potentially reduce the amount of data needed significantly, are proposed in Section 4.4.3. Finally, an iterative strategy to refine the source estimation based on the current reconstruction is discussed in Section 4.4.4.

### 4.4.1 Estimation of point source locations

For point source reconstruction, the sky image consists of a sum of Dirac deltas:

$$I(\mathbf{r}) = \sum_{k=1}^K \alpha_k \delta(\mathbf{r} - \mathbf{r}_k). \quad (4.3)$$

The goal is to reconstruct the source locations  $\mathbf{r}_k$  and intensities  $\alpha_k > 0$  from the beamformed visibility measurements<sup>2</sup>:

$$V_{ij} = \sum_{p=1}^Q \sum_{q=1}^Q w_p^{(i)} \bar{w}_q^{(j)} \sum_{k=1}^K \alpha_k e^{-j2\pi \langle \mathbf{r}_k, \Delta \mathbf{p}_{pq}^{(i,j)} \rangle}, \quad \forall i, j = 1, \dots, L. \quad (4.4)$$

The intensities of sources falling outside the telescope primary beam are significantly attenuated, hence it is reasonable to assume the sky image has finite spatial support, e.g.<sup>3</sup>  $\mathbf{r}_k \in [-\tau_1/2, \tau_1/2] \times [-\tau_2/2, \tau_2/2]$ . The Fourier transform of the sky image, then can be represented by sinc interpolation<sup>4</sup>:

$$\hat{I}(u, v) = \sum_{\xi_1 \in \mathbb{Z}} \sum_{\xi_2 \in \mathbb{Z}} \hat{I}\left(\frac{2\pi\xi_1}{\tau_1}, \frac{2\pi\xi_2}{\tau_2}\right) \text{sinc}\left(\frac{\tau_1}{2}\left(u - \frac{2\pi\xi_1}{\tau_1}\right), \frac{\tau_2}{2}\left(v - \frac{2\pi\xi_2}{\tau_2}\right)\right).$$

From the FRI reconstruction perspective, as long as we can estimate both the uniformly sampled sinusoids  $\hat{I}(2\pi\xi_1/\tau_1, 2\pi\xi_2/\tau_2)$  and the annihilating filter, then the source locations are given by finding roots of polynomials, whose coefficients are specified by the annihilating filters. In general, the zero-crossing of a 2D polynomial is a curve — any Dirac deltas that are located on the curve satisfy the annihilation constraints [95]. In order to uniquely determine the Dirac locations, it is necessary to find two annihilating filters: the Dirac locations are then obtained from the intersections of the two associated curves [99]. Once the source locations are reconstructed, it is a linear problem to estimate source intensities, which amounts to solving a simple least square minimization (see details in Section 4.4.2).

However, this would require the estimation of *infinitely* many sinusoidal samples from a *finite* number of visibility measurements. One way to address this challenge is to assume additionally that the Fourier transform  $\hat{I}(u, v)$  is periodic with period  $2\pi M \times 2\pi N$  for some  $M$  and  $N$  such that  $M\tau_1$  and  $N\tau_2$  are odd numbers<sup>5</sup>. From the Poisson sum formula, the Fourier transform of the sky image can be approximated as [see 96, for a similar treatment in 1D]:

$$\hat{I}(u, v) \approx \sum_{|\xi_1| \leq \lfloor \frac{M\tau_1}{2} \rfloor} \sum_{|\xi_2| \leq \lfloor \frac{N\tau_2}{2} \rfloor} \hat{I}\left(\frac{2\pi\xi_1}{\tau_1}, \frac{2\pi\xi_2}{\tau_2}\right) \varphi\left(u - \frac{2\pi\xi_1}{\tau_1}, v - \frac{2\pi\xi_2}{\tau_2}\right), \quad (4.5)$$

where

$$\varphi(u, v) = \frac{\sin(u) \sin(v)}{MN\tau_1\tau_2 \sin(u/(M\tau_1)) \sin(v/(N\tau_2))}.$$

<sup>2</sup> The proposed approach can also cope with non-beamformed measurements (see comments after (4.2) in Section 4.2).

<sup>3</sup> Without loss of generality, we can always shift the coordinates such that the telescope primary beam is centered at the origin.

<sup>4</sup> In cases where strong sources fall outside the assumed spatial support and still have significant contributions to the visibility measurements, the interpolation representation here will be less accurate.

<sup>5</sup> This is for the consideration of the convergence of Poisson sum equation [see 14, for details].

The beamformed visibility measurements (4.4) are linear combinations of irregularly sampled Fourier transform of the sky image at frequencies specified by the baselines of the antenna pairs  $\Delta \mathbf{p}_{pq}^{(i,j)}$ . From (4.5), we establish a linear connection between a finite set of uniform sinusoidal samples  $\hat{I}(2\pi\xi_1/\tau_1, 2\pi\xi_2/\tau_2)$  and the non-uniformly sampled Fourier transform. In terms of FRI sparse recovery, this amounts to solving a constrained minimization:

$$\begin{aligned} \min_{\substack{\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}, \\ \mathbf{b}}} \quad & \|\mathbf{a} - \mathbf{G}\mathbf{b}\|_2^2 \\ \text{subject to} \quad & \mathbf{b} * \mathbf{c}_1 = \mathbf{0} \text{ and } \mathbf{b} * \mathbf{c}_2 = \mathbf{0}, \end{aligned} \quad (4.6)$$

where

- $\mathbf{a}$  is the visibility measurements (4.4);
- $\mathbf{b}$  is the Fourier transform of the sky image on a uniform grid  $\hat{I}(2\pi\xi_1/\tau_1, 2\pi\xi_2/\tau_2)$ ;
- $\mathbf{G}$  is the linear mapping from the uniformly sampled Fourier transform  $\mathbf{b}$  to the visibilities based on (4.4) and (4.5);
- $\mathbf{c}_1$  and  $\mathbf{c}_2$  are two linearly independent annihilating filters, belonging to a certain feasible set  $\mathcal{C}$  (see Chapter 3 Section 3.2.1 for a typical choice).

Similar to the 1D case, each annihilating filter defines a mask function in the spatial domain, whose value vanishes on a certain curve. The source locations are then given by the intersections of the two curves. In spatial domain, the annihilation constraints can be considered as requiring the multiplication between the two mask functions with the sky image (that contains a few point sources) to be zero. Note that, instead of enforcing the reconstructed signal to follow the interpolation equation (4.5) exactly, we use it only as a metric to gauge the reconstruction quality in (4.6). This explains why a reasonably robust reconstruction is observed experimentally even when the periodicity assumption is violated [96]. However, see Section 4.4.4 for a strategy to refine the linear mapping based on the reconstructed source model.

#### 4.4.2 Estimation of point source intensities

The source intensity  $\alpha_k$  are estimated by solving a least-square fitting problem based on the measurement equation (4.4) once we have reconstructed the source locations  $\mathbf{r}_k$ :

$$\boldsymbol{\alpha}^{\text{opt}} = \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^K} \sum_{i,j=1}^L \left[ V_{ij} - \sum_{p,q=1}^Q w_p^{(i)} \bar{w}_q^{(j)} \sum_{k=1}^K \alpha_k e^{-j2\pi \langle \mathbf{r}_k, \Delta \mathbf{p}_{pq}^{(i,j)} \rangle} \right]^2. \quad (4.7)$$

Equation (4.7) can be re-written more compactly in matrix form. For this we need to introduce a few quantities:

- the visibility matrix  $\Sigma \in \mathbb{C}^{L \times L}$  whose terms are given by  $(\Sigma)_{ij} = V_{ij}$ , for all  $i, j = 1, \dots, L$ .

- the antenna steering matrix  $\mathbf{A} \in \mathbb{C}^{LQ \times K}$  defined by

$$\mathbf{A} = \begin{bmatrix} \boldsymbol{\rho}^{(1)}(\mathbf{r}_1) & \cdots & \boldsymbol{\rho}^{(1)}(\mathbf{r}_K) \\ \boldsymbol{\rho}^{(2)}(\mathbf{r}_1) & \cdots & \boldsymbol{\rho}^{(2)}(\mathbf{r}_K) \\ \vdots & \vdots & \vdots \\ \boldsymbol{\rho}^{(L)}(\mathbf{r}_1) & \cdots & \boldsymbol{\rho}^{(L)}(\mathbf{r}_K) \end{bmatrix},$$

where  $\boldsymbol{\rho}^{(i)}(\mathbf{r}_k) = [e^{-j2\pi\langle \mathbf{r}_k, \mathbf{p}_1^{(i)} \rangle}, \dots, e^{-j2\pi\langle \mathbf{r}_k, \mathbf{p}_Q^{(i)} \rangle}] \in \mathbb{C}^Q$  is the antenna steering vector for station  $i$  and  $\mathbf{r}_k$  are the reconstructed source locations.

- the beamforming matrix  $\mathcal{W} \in \mathbb{C}^{LQ \times L}$  is a block-diagonal matrix defined by

$$\mathcal{W} = \begin{bmatrix} \bar{\mathbf{w}}^{(1)} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \bar{\mathbf{w}}^{(L)} \end{bmatrix}, \quad (4.8)$$

where  $\mathbf{w}^{(i)} = [w_1^{(i)}, \dots, w_Q^{(i)}] \in \mathbb{C}^Q$  is the beamforming vector for station  $i$ .

With the notation introduced above, (4.7) reduces to:

$$\begin{aligned} \boldsymbol{\alpha}^{\text{opt}} &= \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^K} \left\| \Sigma - \mathcal{W}^H \mathbf{A} \text{diag}(\boldsymbol{\alpha}) \mathbf{A}^H \mathcal{W} \right\|_2^2 \\ &= \arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^K} \left\| \boldsymbol{\sigma} - \left[ (\mathcal{W}^T \bar{\mathbf{A}}) \circ (\mathcal{W}^H \mathbf{A}) \right] \boldsymbol{\alpha} \right\|_2^2, \end{aligned} \quad (4.9)$$

where  $\boldsymbol{\sigma} = \text{vec}(\Sigma)$  is the vectorization of the visibility matrix, and  $\circ$  denotes the Khatri-Rao product [see 130, for more details]. The closed-form solution of (4.9) is

$$\boldsymbol{\alpha}^{\text{opt}} = \left[ (\mathcal{W}^T \bar{\mathbf{A}}) \circ (\mathcal{W}^H \mathbf{A}) \right]^\dagger \boldsymbol{\sigma},$$

where  $\dagger$  denotes the Moore-Penrose pseudo-inverse. The optimization problem (4.9) could be further constrained by  $\boldsymbol{\alpha} > 0$  (since source intensities are positive) leading to a non-negative least-squares problem, which can be solved within a finite number of iterations [76]. Finally, when the number of sources is uncertain, a sparsity promoting penalty term  $\nu \|\boldsymbol{\alpha}\|_1$  could also be envisaged, with the parameter  $\nu$  acting as model selection parameter. Unfortunately, such a penalty term would bias the estimation of the source intensities. Instead, we propose an alternative model order selection procedure, based on the fitting error (see Section 4.5.4).

### 4.4.3 Coherent multiband reconstruction

Modern radio telescopes operate over a wide frequency range, e.g., 30MHz to 240MHz for LOFAR [131]. The emitted electromagnetic waves of celestial sources within the operation range are measured simultaneously, which are subsequently filtered into different sub-bands. If the consistency of the measurements across different sub-bands is exploited, it may potentially reduce significantly the integration times needed in order to have a reliable reconstruction.

Classic approaches, e.g., multi-frequency synthesis [35], and multi-frequency CLEAN [110], try to map multi-frequency visibility measurements into a single sub-band centered at a reference



frequency based on a frequency-dependent sky brightness distribution model.

With FRI-based sparse recovery, the mutual information shared across different sub-bands can be exploited in a coherent manner. It is usually reasonable to assume that the source locations remain the same across all subsequent sub-bands. Since the annihilating filter is uniquely specified by the point source locations alone, this implies that we should find *one* annihilating filter for *all* sub-bands such that the annihilation equations are satisfied<sup>6</sup>. In general, the source intensities  $\alpha_k$  differ from sub-band to sub-band. Hence, the uniformly sampled sinusoids, which are chosen as the interpolation knots  $\hat{I}(2\pi\xi_1/\tau_1, 2\pi\xi_2/\tau_2)$  in (4.5), are sub-band-dependent. Then, the multiband point source reconstruction amounts to solving

$$\begin{aligned} \min_{\substack{\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}, \\ \mathbf{b}_1, \dots, \mathbf{b}_J}} \sum_{i=1}^J \left\| \mathbf{a}^{(i)} - \mathbf{G}^{(i)} \mathbf{b}^{(i)} \right\|_2^2 \\ \text{subject to } \mathbf{b}^{(i)} * \mathbf{c}_1 = \mathbf{0} \text{ and } \mathbf{b}^{(i)} * \mathbf{c}_2 = \mathbf{0} \quad \text{for } i = 1, \dots, J, \end{aligned} \quad (4.10)$$

where  $\mathbf{a}^{(i)}$  and  $\mathbf{b}^{(i)}$  are the visibility measurements and the uniform sinusoidal samples in the  $i$ -th sub-band, respectively; and  $\mathbf{G}^{(i)}$  is the linear mapping based on (4.4) and (4.5) for each one of the  $J$  sub-bands. Note that (4.10) is in fact the same formulation<sup>7</sup> as (4.6) with a change of variables:

$$\mathbf{a} = \begin{bmatrix} \mathbf{a}^{(1)} \\ \vdots \\ \mathbf{a}^{(J)} \end{bmatrix}, \mathbf{b} = \begin{bmatrix} \mathbf{b}^{(1)} \\ \vdots \\ \mathbf{b}^{(J)} \end{bmatrix}, \text{ and } \mathbf{G} = \begin{bmatrix} \mathbf{G}^{(1)} & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{G}^{(J)} \end{bmatrix}.$$

Once we have estimated the common annihilating filter  $\mathbf{c}$  for all sub-bands, the source locations and intensities are determined in the same manner as in the single band case.

#### 4.4.4 Update strategy for the linear mapping $\mathbf{G}$ in (4.6)


Ideally,  $\mathbf{G}$  should be constructed based on the measurement equation (4.4) in the constrained optimization (4.6), and the discrepancies between the re-synthesized visibilities  $\mathbf{G}\mathbf{b}$  with the given measurements minimized. If the actual mapping, which links the Fourier transform of the sky image on a uniform grid to the visibility measurement, were available, then the FRI-based sparse recovery would give the exact reconstruction from the noiseless measurements. However, this is not feasible, as the exact mapping based on (4.4) contains the source locations and intensities, which are unknown a priori. One possible strategy is to use the initial reconstruction, where  $\mathbf{G}$  was approximated with the periodic-sinc interpolation (4.5), to update the objective function in (4.6).

To describe this update strategy, we need to introduce some notation (see Figure 4.2 for a summary):

- Denote by  $\mathbf{G}_0$  the linear operator that maps source intensities  $\boldsymbol{\alpha}$  to the Fourier transform

<sup>6</sup>The same approach can be applied to measurements of different polarizations within the same sub-band, where the source locations are common but intensities differ for each polarization.

<sup>7</sup>To be precise, the convolution in the annihilation constraints in the multiband case should be understood as a convolution for each sub-band, which amounts to vertically stacking the convolution matrices associated with all sub-bands.



**Figure 4.2:** Diagram of the various linear operators involved in the update strategy for the forward operator  $\mathbf{G}$  (see details in Section 4.4.1 and Section 4.4.4).

of the sky image on a uniform grid  $(u, v) = (2\pi\xi_1/\tau_1, 2\pi\xi_2/\tau_2)$  as in (4.5):

$$\mathbf{G}_0 : \boldsymbol{\alpha} \mapsto \hat{I}\left(\frac{2\pi\xi_1}{\tau_1}, \frac{2\pi\xi_2}{\tau_2}\right) = \sum_{k=1}^K \alpha_k e^{-j2\pi(l_k\xi_1/\tau_1 + m_k\xi_2/\tau_2)},$$

with some given source locations  $\mathbf{r}_k = (l_k, m_k)$ .

- Denote by  $\mathbf{G}_1$  the matrix mapping the source intensities  $\boldsymbol{\alpha}$  to the non-gridded Fourier samples  $\left(u_{pq}^{(i,j)}, v_{pq}^{(i,j)}\right)$ :

$$\mathbf{G}_1 : \boldsymbol{\alpha} \mapsto \hat{I}\left(u_{pq}^{(i,j)}, v_{pq}^{(i,j)}\right), \quad i, j = 1, \dots, L; p, q = 1, \dots, Q.$$

The matrix  $\mathbf{G}_1$  can also be expressed of the antenna steering matrix  $\mathbf{A}$  as  $\mathbf{G}_1 = \bar{\mathbf{A}} \otimes \mathbf{A}$ .

- Denote by  $\mathbf{W} \in \mathbb{C}^{L^2 \times L^2 Q^2}$  the cross-beamforming matrix that beamforms the off-grid Fourier samples:

$$\mathbf{W} : \hat{I}\left(u_{pq}^{(i,j)}, v_{pq}^{(i,j)}\right) \mapsto V_{ij} = \sum_{p,q=1}^Q w_p^{(i)} \bar{w}_q^{(j)} \hat{I}\left(u_{pq}^{(i,j)}, v_{pq}^{(i,j)}\right).$$

Here  $\mathbf{W}$  is related to the beamforming matrix  $\mathcal{W}$  in (4.8) as:  $\mathbf{W} = \mathcal{W}^T \otimes \mathcal{W}^H$ .

- Finally, denote by  $\Phi$  the periodic-sinc interpolation (4.5) evaluated at non-gridded Fourier

**Algorithm 3:** Point source estimation with FRI

---

**Input** : Visibility measurements  $\mathbf{V}$ , radio-telescope antenna layout  $\mathbf{p}$ , number of sources  $K$  to reconstruct

**Output**: Source locations  $\mathbf{r}_k^{\text{opt}}$ , source intensities  $\boldsymbol{\alpha}^{\text{opt}} = [\alpha_1, \dots, \alpha_K]$

- 1 Initialize  $\mathbf{G} = \mathbf{W}\boldsymbol{\Phi}$ ,  $\text{MinError} = \infty$ ;
- for**  $\text{loop} \leftarrow 1$  **to**  $\text{MaxIterations}$  **do**
- 2 Reconstruct annihilating filter coefficients  $\mathbf{c}_1$  and  $\mathbf{c}_2$  from (4.6);
- 3  $\mathbf{r}_k \leftarrow$  common roots of two polynomials with coefficients  $\mathbf{c}_1$  and  $\mathbf{c}_2$ , respectively;
- 4 Update  $\mathbf{G}_1$  with the reconstructed  $\mathbf{r}_k$ ;
- 5  $(\boldsymbol{\alpha}, \text{FittingError}) \leftarrow \min_{\boldsymbol{\alpha}} \|\mathbf{V} - \mathbf{W}\mathbf{G}_1\boldsymbol{\alpha}\|_2^2$ ;
- if**  $\text{FittingError} < \text{MinError}$  **then**
- 6    $\mathbf{r}_k^{\text{opt}} \leftarrow \mathbf{r}_k$ ,  $\boldsymbol{\alpha}^{\text{opt}} \leftarrow \boldsymbol{\alpha}$ ,  $\text{MinError} \leftarrow \text{FittingError}$ ;
- 7   Update  $\mathbf{G}_0$  and  $\mathbf{G}_1$  with  $\mathbf{r}_k^{\text{opt}}$  and  $\mathbf{G} = \mathbf{W}(\mathbf{G}_1\mathbf{G}_0^\dagger + \boldsymbol{\Phi}\mathbf{P}_{\mathcal{N}(\mathbf{G}_0^\dagger)})$ ;
- end**
- end**

---

samples  $(u_{pq}^{(i,j)}, v_{pq}^{(i,j)})$ :

$$\boldsymbol{\Phi} : \hat{I}\left(\frac{2\pi\xi_1}{\tau_1}, \frac{2\pi\xi_2}{\tau_2}\right) \mapsto \hat{I}(u_{pq}^{(i,j)}, v_{pq}^{(i,j)}).$$

Then the linear mapping is chosen as  $\mathbf{G} = \mathbf{W}\boldsymbol{\Phi}$  in the initial estimate. In comparison, if we knew the ground truth source locations and intensities, the optimal mapping would be  $\mathbf{G} = \mathbf{W}\mathbf{G}_1\mathbf{G}_0^\dagger$ , where  $\mathbf{G}_0^\dagger$  is the Moore-Penrose pseudo-inverse  $(\mathbf{G}_0^H\mathbf{G}_0)^{-1}\mathbf{G}_0^H$ .

Note that  $\mathbf{G}_1\mathbf{G}_0^\dagger$ , which transforms the uniformly sampled Fourier data to the irregularly sampled ones<sup>8</sup>, has rank at most  $K$ . At any intermediate step, we may choose the linear mapping  $\mathbf{G}$  as:  $\mathbf{W}(\mathbf{G}_1\mathbf{G}_0^\dagger + \boldsymbol{\Phi}\mathbf{P}_{\mathcal{N}(\mathbf{G}_0^\dagger)})$ , where  $\mathbf{G}_0$  and  $\mathbf{G}_1$  are built with the reconstructed point source locations and  $\mathbf{P}_{\mathcal{N}(\mathbf{G}_0^\dagger)}$  is the orthogonal projection onto the null space of  $\mathbf{G}_0^\dagger$ :  $\mathbf{I} - \mathbf{G}_0(\mathbf{G}_0^H\mathbf{G}_0)^{-1}\mathbf{G}_0^H$ . Experimentally, such an iterative strategy manages to refine the linear mapping and results in a reliable reconstruction (see an example in Section 4.5.2).

We emphasize that the reconstruction quality should always be measured based on (4.4), with  $\mathbf{r}_k$  and  $\alpha_k$  the *reconstructed* source locations and intensities, respectively, regardless of the update strategy for the linear mapping  $\mathbf{G}$ . We summarize the FRI-based point source reconstruction in Algorithm 3.

---

<sup>8</sup>Superficially, this looks similar to the "gridding" in a conventional approach, e.g., CLEAN. However, unlike CLEAN, the final resolution that can be achieved by the FRI-based algorithm is not related to the grid step size but only the noise level in the given measurements [see 96].

## 4.5 Results

### 4.5.1 Data and experiment setup

The proposed FRI-based sparse recovery approach for point source estimation (LEAP) was validated with both simulated visibilities and real observations from LOFAR. In simulation, visibilities were generated from ground truth point source parameters (locations and intensities) with the LOFAR core station antenna layout. In experiments with real LOFAR observations, we sub-sampled the visibility measurements over different integration times such that only 2% or 0.25% of the total integration times in the measurement set were available to the reconstruction algorithms in single band and multi-band scenarios, respectively.

We should point out that it is not only the number of integration times that matters: with the same number of integration times taken *consecutively*, a much worse reconstruction is obtained by both CLEAN and LEAP. Experimentally, we observe that it is better to take measurements that are well-spread over the whole acquisition time. One explanation might be that with a larger time separation between adjacent measurements, the earth has more significant displacement in space. Thus, it allows the radio interferometer to sample the  $(u, v)$ -plane sparsely but over a large area (instead of densely sampling a local area as in the case with consecutive integration times). Spatial diversity in the  $(u, v)$  domain sampling makes the reconstruction algorithms more resilient to noise.

We summarize the experimental setups in terms of antenna layouts, integration time and sub-band selections in Table 4.1. The reconstruction quality of the FRI-based approach was measured by the average distance<sup>9</sup> between the recovered and the ground truth source locations (in simulation) or the catalog data (in real-data experiments).

The LEAP results were compared to those obtained from WSClean<sup>10</sup>, which implements the state-of-the-art w-stacking CLEAN algorithm [92]. The maximum number of iterations for the WSClean algorithm was set to  $4 \times 10^4$ . The threshold level was determined automatically by the WSClean algorithm based on the estimated background noise level. Unlike LEAP, which reconstructs point source locations and intensities directly, the CLEAN algorithm typically produces discrete sky images, which will be compared to the FRI reconstruction by visual inspection.

In the following part, we first conduct simulations to verify the effectiveness of the updating strategy of the linear mapping in Section 4.5.2. Next, we investigate the resolvability of the proposed algorithm by simulating visibilities from two point sources that are separated by various distances in Section 4.5.3. Further, a strategy to avoid false detections by selecting an adequate model order is validated through simulations in Section 4.5.4. Finally, we apply LEAP to actual LOFAR observations from the Boötes field (Section 4.5.5), which consists of mostly point sources; and the “Toothbrush” cluster (Section 4.5.6), which has an extended structure in addition to many point sources within the field of view.

### 4.5.2 Iterative refinement of linear mapping

One challenge in applying the FRI-based sparse recovery technique to radio astronomy is identifying a suitable surrogate function to gauge reconstruction quality—the ideal MSE criteria

<sup>9</sup>The correspondence between the reconstructed and ground truth source locations were obtained by permuting the source locations such that the distance between the permuted reconstruction and the ground truth source locations was minimized.

<sup>10</sup>Available at <https://sourceforge.net/projects/wsclean/>

**Table 4.1:** Summary of various datasets used in the experiments. The LOFAR antenna layout is used in simulation (Dataset I). We have sub-sampled the real observations over different integration times such that only 2% (for the single band cases: Dataset II and IV) or 0.25% (for the multi-band case: Dataset III) of the total integration times is available to the FRI reconstruction algorithm.

\*NOTE: 4 of the 24 LOFAR core stations (HBA) were not working during the observation. We have used 4 additional remote stations that are closest to the telescope center.

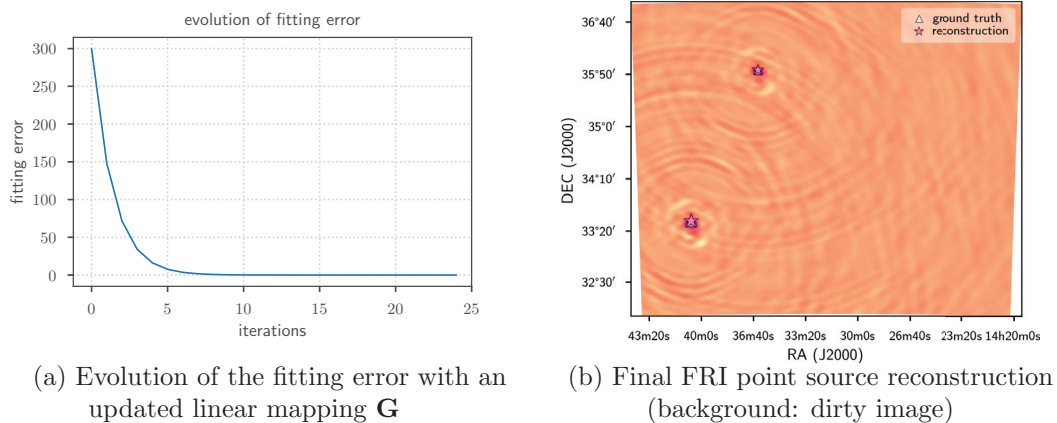
Dataset	I	II	III	IV
Sources	Point sources	Boötes field		“Toothbrush” cluster RX J0603.3+4214
Data type	Simulated	Actual observation		Actual observation
Antenna layout	LOFAR core stations (HBA)	LOFAR core and remote stations (HBA)		LOFAR core and remote stations (HBA)
Number of stations	24	28*		36
Telescope focus (RA, DEC)	(14h32m00.0s, +34°30′00.1″)	(14h32m00.0s, +34°30′00.1″)		(06h03m20.0s, +42°14′00.0″)
Longest baseline (× wavelength)	713.3	1015.4	1019.9	1150.3
Instrument angular resolution	4′49.2″	3′23.1″	3′22.2″	2′59.3″
Polarization	I			
Observation start time (UTC)	2014-08-10-13:00:04.01	2014-08-10-13:00:04.01		2013-02-24-15:32:01.42
Observation end time (UTC)	2014-08-10-19:53:58.48	2014-08-10-20:07:19.59		2013-02-25-00:09:24.51
Time resolution (s)	400.56	400.56	3204.45	500.70
Number of integration times	63	72	9	63
Number of sub-bands	1	1	8	1
Sub-band frequency (MHz)	145.8	145.8	145.8 ~ 146.5	132.1
Section	4.5.2, 4.5.3 and 4.5.4	4.5.5	4.5.5	4.5.6

based on the measurement equation (4.4) requires the knowledge of the (unknown) ground truth source locations and intensities. We proposed one possible strategy that allows us to refine the objective function based on the current reconstruction in Section 4.4.4. In order to verify the effectiveness of such a strategy, we generated an empty measurement set (MS) with the LOFAR antenna layout as specified in Table 4.1 Dataset I. The MS file is then filled with noiseless visibilities that are simulated based on (4.4) from two point sources with randomly generated intensities and locations within the field-of-view ( $5^\circ \times 5^\circ$ ).

The evolution of the fitting error between the re-synthesized visibilities (4.4) (based on the reconstructed point source parameters) and the given visibility measurements is shown in Figure 4.3. The reconstructed sources are included for visual comparison, where the dirty image is overlaid with the reconstructed and ground truth point sources. With this simple updating strategy, we indeed obtain the exact reconstruction after a few iterations.

### 4.5.3 Source resolution

In this section, we investigate the resolving power of the proposed FRI-based sparse recovery, by comparing the performance to that of CLEAN. The antenna layout of the 24 LOFAR core stations



**Figure 4.3:** Iterative refinement of the linear mapping  $\mathbf{G}$  in (4.6) for exact point source reconstruction (see details in Section 4.4.4). (a) The evolution of the fitting error between the re-synthesized and given noiseless visibility measurements. (b) Exact reconstruction of the point sources with LEAP (background image: dirty image from the same visibility measurements).

was used to simulate a 7 hour single sub-band observation with center frequency 145.8MHz (HBA band). The visibility measurements were taken every 400.56 seconds, leading to a total 63 sets of visibilities at different time instances. The maximum baseline was 713.3 wavelengths, which corresponds to an instrument angular resolution of  $4'49.2''$ . For simplicity, we did not account for polarization effects in this chapter. We note, however, that the technique described in Section 4.4.3 to reconstruct point sources from multi-bands may be adapted to treat together different polarizations in a coherent manner (see remarks in footnote 6).

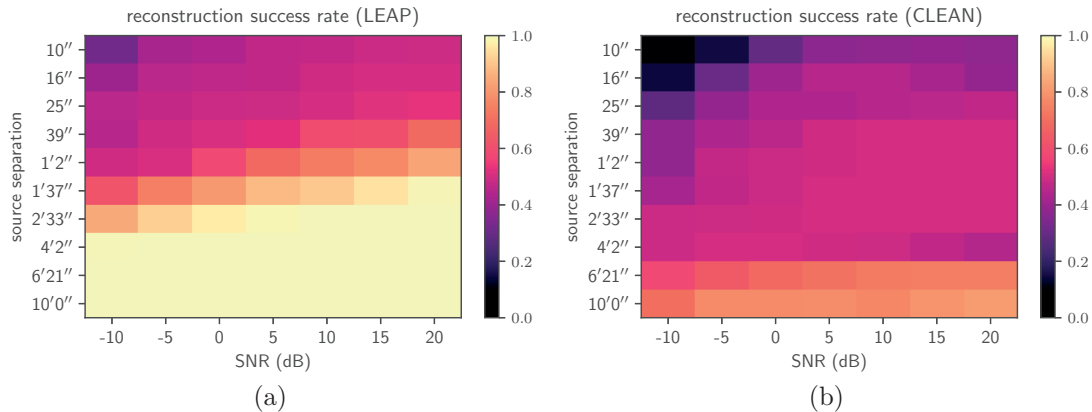
We simulated visibilities of two point sources with unitary intensities. The two sources were separated by distances varying from  $10''$  to  $10'$  on a log scale. The particular antenna layout means there is not always the same sensitivity along all directions. To alleviate this potential direction-dependent bias, we averaged the results over 100 signal realizations with different relative orientations between the two sources for each separation distance. Circularly symmetric complex Gaussian white noise is added to the noiseless visibilities such that the signal to noise ratio (SNR) in the visibility measurements ranges from  $-10$  dB to  $20$  dB with a step size of  $5$  dB.

A point source was considered to be successfully reconstructed if the estimated source location is within half the separation distance between the two sources, from the ground truth source location. Hence, the average reconstruction success rate for  $K$  point sources is:

$$\text{success rate} = \frac{\# \text{ of } k \text{ such that } \text{dist}(\mathbf{r}'_k, \mathbf{r}_k) < \Delta \mathbf{r} / 2}{K}, \text{ for } k = 1, 2. \quad (4.11)$$

Here  $\text{dist}(\cdot, \cdot)$  computes the distance between the reconstructed  $\mathbf{r}'_k$  and ground truth source locations  $\mathbf{r}_k$ ; and  $\Delta \mathbf{r}$  is the separation between the two sources. We extracted the reconstructed source locations from the CLEAN model image with a pixel size  $3.5''$ . The average reconstruction success rate is shown in Figure 4.4 for both CLEAN and LEAP.

Note that while the instrument angular resolution was close to  $5'$  here, the FRI-based sparse recovery still manages to resolve two sources beyond the instrument limitation in many cases.



**Figure 4.4:** Average reconstruction success rate (4.11) of two point sources that are separated by various distances from noisy visibility measurements with (a) LEAP and (b) CLEAN. The reconstructed point sources with CLEAN are extracted from the corresponding CLEAN model image. A point source is considered to be successfully reconstructed if the estimated source location is within half the separation distance with the ground truth source location. Results for each noise level and source separation are averaged over 100 different signal and noise realizations. Instrument angular resolution is  $4'49.2''$ .

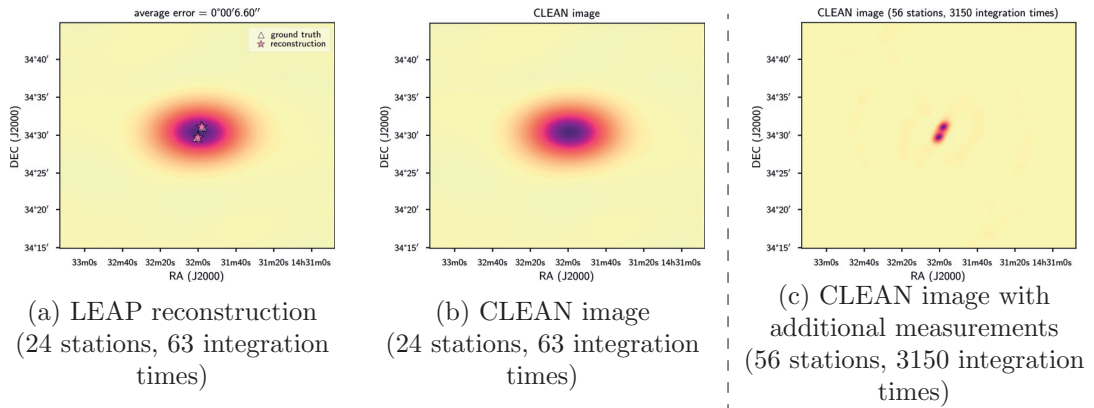
This is in stark contrast to image-based approaches such as CLEAN or compressed sensing (CS), where the reconstructions are typically spatial domain images—In order to determine point source locations (and intensities), an additional blob detection algorithm needs to be employed. In contrast, the FRI-based approach starts from a point source assumption, and reconstructs the source parameters directly without going through an intermediate spatial domain image.

To better illustrate the difference between the proposed method and other image-based reconstruction methods, we include two examples where it would otherwise not be possible to recover all point sources based on the estimated sky images (Figure 4.5 and Figure 4.6):

- Figure 4.5 shows how two closely located sources,  $1'30''$  apart, are accurately estimated from mildly noisy visibility measurements (SNR = 20 dB) with LEAP, while the CLEAN image contains one big blob encompassing both sources. In fact, neither one of the two source locations corresponds to the peak of the blob.

Further, we considered another case for CLEAN, where additional measurements from 32 LOFAR remote stations were added. With this configuration, the telescope has a much smaller angular resolution  $6.10''$  (compared with  $4'49.2''$  with 24 LOFAR core stations only). The time resolution of the visibility measurements from all stations is also increased: adjacent integration times are separated by 8.01 seconds. In total, visibilities from all 56 stations at 3150 integration times are given to the CLEAN algorithm. The blob size is significantly reduced and both sources are resolved by CLEAN (Figure 4.5 (c)).

- Figure 4.6 shows two well-separated sources,  $1^\circ30'$  apart, are reliably reconstructed from highly noisy visibility measurements (SNR =  $-10$  dB) with LEAP, while the weaker source, whose intensity is  $1/5$  of that of the strong source, is completely buried in the noisy background, and cannot be detected from the estimated sky image using CLEAN. With additional visibilities from a higher time resolution (8.01 seconds between adjacent integration



**Figure 4.5:** Resolve closely located sources accurately with LEAP (SNR = 20 dB, source separation:  $1'30''$ , antenna layout: 24 LOFAR core stations, background: CLEAN image). (a) LEAP reconstruction. (b) CLEAN image from the same set of visibility measurements. (c) CLEAN image from *additional* visibilities with higher time resolution (8.01 seconds between adjacent integration times) and longer baselines (56 LOFAR core and remote stations). The setup in (c) with longer baselines reduces the telescope angular resolution to  $6.10''$ .

times), both the strong source in the middle of the field of view and the weak source are correctly reconstructed by CLEAN (Figure 4.6 (c)).

#### 4.5.4 Model order selection

The point source model requires a choice of a certain model order  $K$ . In simulation, this parameter is assumed to be given a priori, while with actual observations, we have to specify / estimate it. A strategy is then needed to determine if a particular choice of  $K$  over-estimates<sup>11</sup> the model order and leads to false detections. One method is based on the fitting error between the reconstructed source model and the given measurements [59]. From our experience, the FRI reconstruction algorithm can reliably estimate a sparse signal that fits the given measurements up to the noise level [96]. The strategy is constructive since we can usually estimate the noise level from a source finding algorithm, e.g., Duchamp [139], SoFiA [113], and PyBDSF<sup>12</sup>, or manually specify the approximation error allowed in the reconstruction.

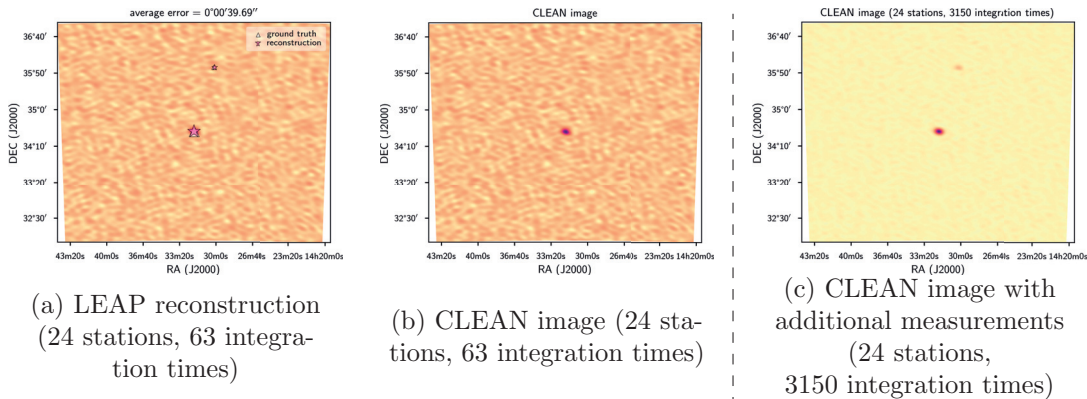
In particular, if the fitting error achieved by FRI-based reconstruction is still within the noise level with a lower model order, then we can further reduce the number of sources by one. The final model order is obtained by repeating the process until any further reduction in  $K$  leads to a fitting error above the noise level.

We validated this approach by simulating visibility measurements for  $K = 3$  point sources from 24 LOFAR core stations (Dataset I in Table 4.1). The noiseless visibilities are contaminated by circularly symmetric complex Gaussian white noise such that the SNR in the visibility measurements is 0 dB. We applied LEAP with different model orders, and compared the noise level with the fitting errors between the given and the re-synthesized visibility measurements

<sup>11</sup>In the case of under-estimation, the most dominating point sources will be reconstructed by the algorithm with a fitting error above the noise level (see an example in Figure 4.7 (c)).

<sup>12</sup>Available at <https://github.com/lofar-astron/PyBDSF>.





**Figure 4.6:** Reliable estimation of sources from highly noisy measurements with LEAP (SNR =  $-10$  dB, source separation:  $1^{\circ}30'$ , antenna layout: 24 LOFAR core stations, background: CLEAN image). (a) LEAP reconstruction. (b) CLEAN image from the same set of visibility measurements. (c) CLEAN image from *additional* visibilities with higher time resolution (8.01 seconds between adjacent integration times) from 24 core stations.

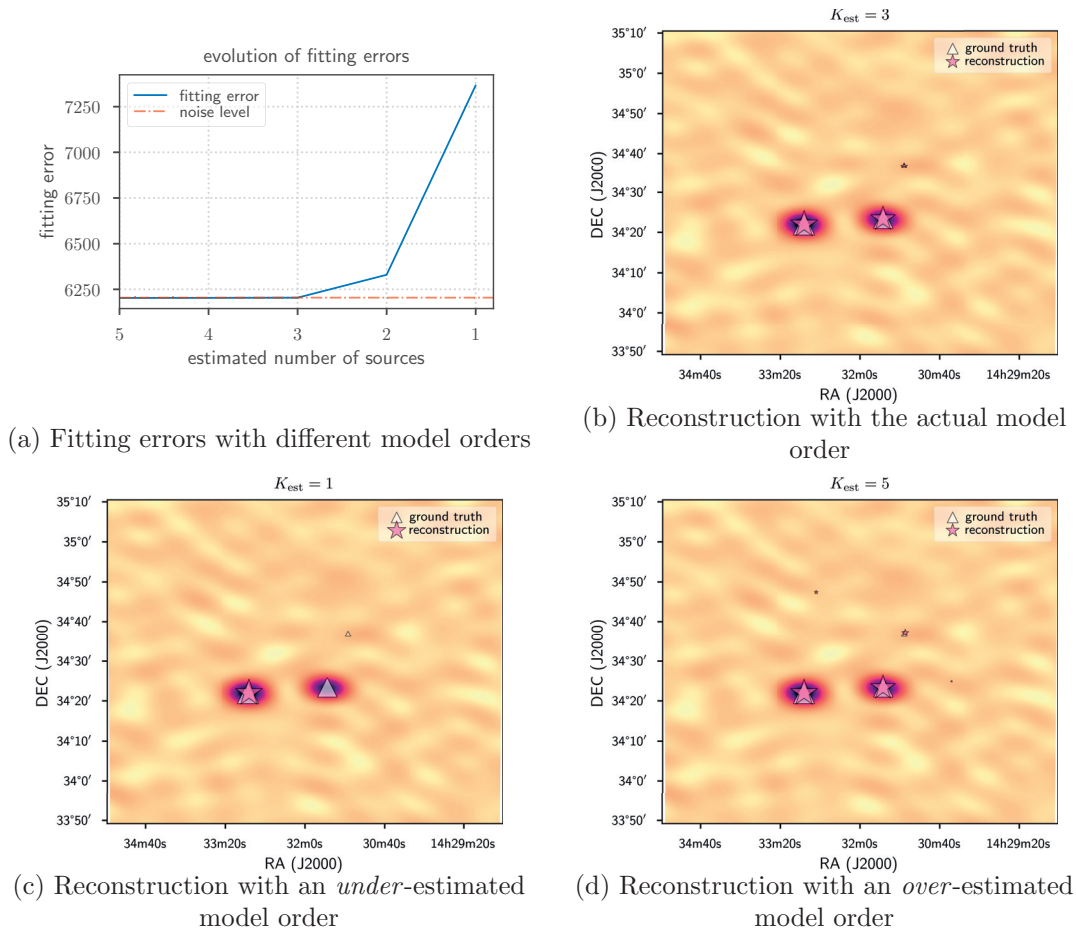
based on (4.4) (Figure 4.7(a)). As soon as the model order is reduced below the actual number of sources, the fitting error jumps above the noise level, while this error remains stable for over-estimation cases. The reconstructed point source model that corresponds to under-estimation and over-estimation cases are also included in Figure 4.7.

#### 4.5.5 Actual LOFAR observation: Boötes Field

In the previous section, we validated the robustness of the proposed FRI-based sparse recovery method with simulated point sources. In this section, we apply LEAP to an actual LOFAR observation from the Boötes field. Source estimation from real observations of a radio telescope presents a significant challenge over and above ideal simulation conditions. Visibility measurements in a typical setting suffer from severe noise contamination, which may arise from thermal noise at antennas, the planar approximation (4.1), as well as directional dependent artifacts due to ionosphere variations [142]. We used<sup>13</sup> the visibility measurements from the 24 LOFAR core stations and 4 remote stations closest to the telescope center, and considered two different scenarios for the point source reconstruction in the Boötes field:

1. Single-band visibility measurements at 145.8MHz were extracted from the measurement set (MS). We sub-sampled the MS file uniformly every 50 integration times ( $\sim 7$  minutes), leading to a total 72 sets of visibilities from 2% of the all integration times. See Section 4.4.1 and Section 4.4.2 for the algorithmic details.
2. Visibility measurements within 8 sub-bands centered at frequencies from 145.8MHz to 146.5MHz were extracted from the MS file. In total, 9 sets of visibilities at different integration times were given to the reconstruction algorithms by sub-sampling the MS file every 400 integration times ( $\sim 53$  minutes). Sources were reconstructed with the multi-band formulation presented in Section 4.4.3.

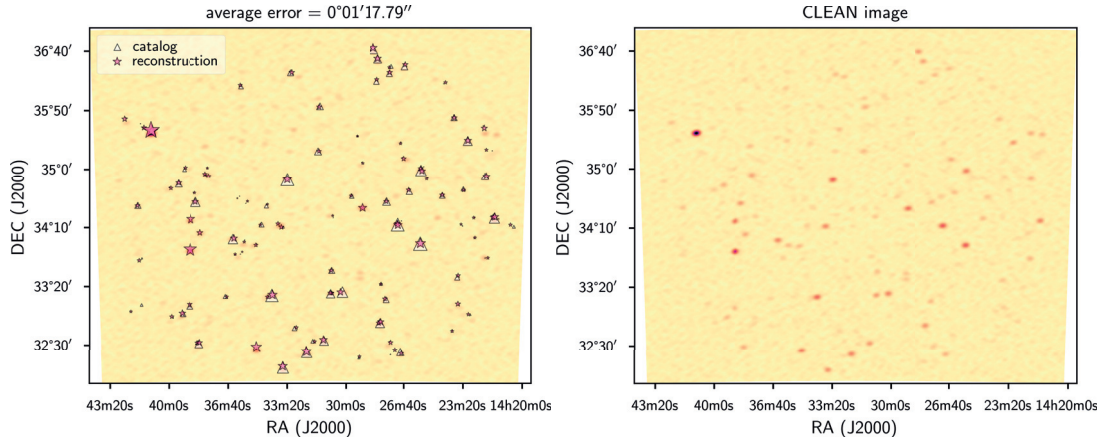

<sup>13</sup>During the observation, 4 out of the 24 core stations were not working.



**Figure 4.7:** Model order selection based on the fitting error of the reconstructed point source model and the given visibility measurements (Dataset I). The model order is the minimum number of sources such that the fitting error reaches the noise level of the measurements (Background images of (b) – (d): CLEAN images; signal to noise ratio (SNR) in the visibilities: 0 dB; Actual model order: 3). (a) Evolution of the fitting errors against different point source model orders  $K$  in (4.3). (b) Reconstructed source models with the correct model order  $K = 3$ . (c) Reconstruction with under-estimated model order  $K = 1$ . (d) Reconstruction with over-estimated model order  $K = 5$ .

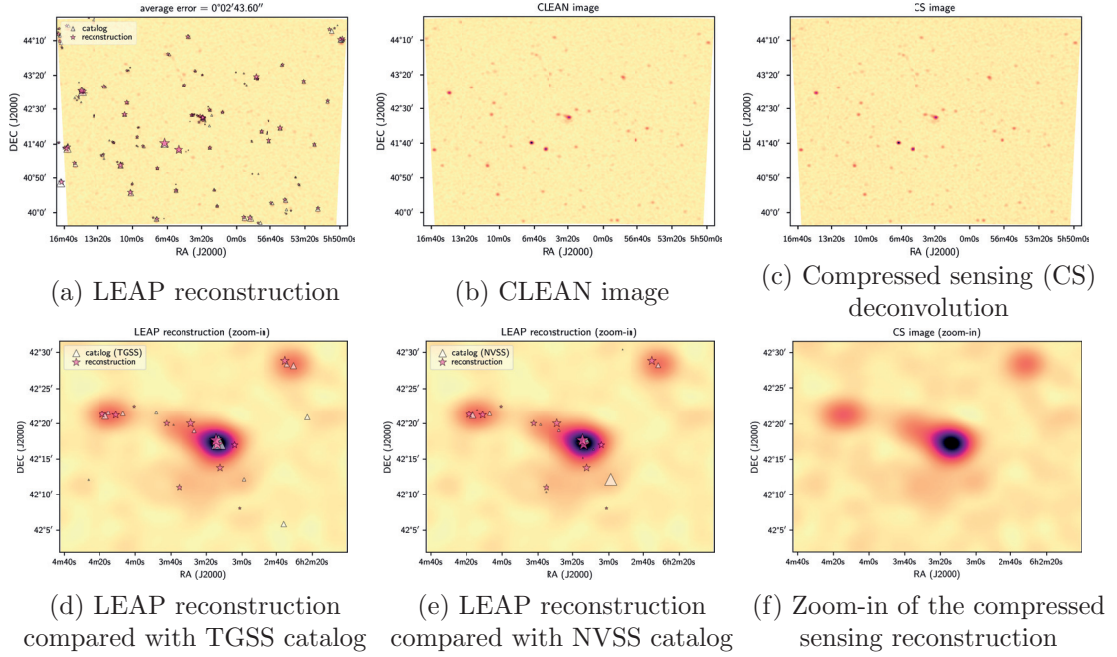
In both cases, LEAP reconstructed Fourier transform of the sky image within each frequency sub-band, on a uniform grid of size  $57 \times 63$  that spanned the telescope  $uv$ -coverage. The source locations were subsequently estimated from these uniform Fourier data. We fixed a priori the number of sources to be  $K = 100$  in (4.3) for the LEAP reconstruction. Of course, in practice a non-arbitrary choice of an adequate model order  $K$  is needed, and this is discussed in Section 4.5.4.

The reconstructed point sources for the single-band and multi-band cases are plotted in Figure 4.8, where the background image is the corresponding CLEAN image reconstructed from

(a) Source detection from *single-band* LOFAR measurements (Dataset II)(b) Source detection from *multi-band* LOFAR measurements (Dataset III)

**Figure 4.8:** Source detection from actual LOFAR observations from the Boötes field (Dataset II and III) with LEAP (the left column) and CLEAN (the right column). (a) LEAP and CLEAN reconstruction from single band visibility measurements at frequency 145.8MHz with 72 integration times. (b) LEAP and CLEAN reconstruction from visibility measurements within 8 sub-bands at frequencies from 145.8MHz to 146.5MHz with 9 integration times.

the same sets of visibility measurements in each case. We compare the estimated source parameters with a catalog of the Boötes field at 130 ~ 169MHz [142]. The errors of the estimated source locations are  $1'17.79''$  and  $1'38.01''$  in the single-band and multi-band cases, respectively. Similar to the observations in simulated cases, we find that LEAP can reliably resolve closely located sources from the actual LOFAR measurements. The advantage of the FRI-based sparse recovery, which gives a direct estimate of the point source parameters, is evident by comparing the CLEAN image in both single-band and multi-band cases: several sources would otherwise be too weak to be reliably detected with a blob detection algorithm applied to the CLEAN image without introducing many false detections.



**Figure 4.9:** Point source reconstruction from actual LOFAR observations from the Toothbrush cluster RX J0603.3+4214 (Dataset IV). Point sources are reliably estimated even in the presence of extended sources (the Toothbrush cluster). Existence of sources are validated by referencing both the TGSS and the NVSS catalog. (a) The reconstructed point sources with LEAP. (b) The CLEAN image from the same set of visibility measurements. (c) The deconvolved image with a compressed sensing (CS) based approach [39]. Zoom-in comparisons of the LEAP reconstruction against (d) the TGSS catalog and (e) the NVSS catalog around the Toothbrush cluster. (f) Zoom-in plot of the CS reconstruction around the toothbrush cluster.

#### 4.5.6 Actual LOFAR observation: Toothbrush Cluster

A natural question to ask is how well the algorithm recovers point sources in the presence of extended sources within the field of view. In order to test their influence, we applied LEAP to a LOFAR observation from the Toothbrush cluster RX J0603.3+4214, which contains one of the brightest radio relics. The toothbrush shape is suggested to be a consequence of a triple merger event based on simulation [20]. We used the LOFAR observation with both the 24 core stations and 12 remote stations within a single-band at frequency 132.1MHz (Dataset IV in Table 4.1).

The LEAP reconstruction results are compared to the CLEAN image. We also included the deconvolved image obtained with a compressed sensing (CS) based approach [39] for reference. We overlaid the CLEAN image with the reconstructed point sources from LEAP in Figure 4.9, and validated the existence of sources by comparing the reconstructions with the 150MHz TIFR GMRT sky survey (TGSS) [69]. Even in the presence of the Toothbrush cluster at the center of the field of view, LEAP is robust enough to estimate the point sources reliably. The average reconstruction error of the source locations compared with the TGSS catalog is  $2'43.60''$ . A zoom-in plot, Figure 4.9(d), of the area around the toothbrush cluster, reveals that LEAP reconstructed a few sources not in the TGSS catalog (and hence are mismatched to other sources

in the catalog). In order to determine whether these sources were false detections, we cross-referenced with the NRAO VLA sky survey (NVSS) [34], which observes the sky at a much higher frequency (1.4GHz). The extra sources reconstructed by LEAP were indeed confirmed to be actual radio sources (Figure 4.9 (e)). The average LEAP reconstruction with respect to the NVSS catalog was  $2'1.36''$ .

## 4.6 Discussion

### 4.6.1 Resolution

In CLEAN-based source estimation algorithms, the reconstructed sky model (which consists of a few non-zero pixels around sources) are convolved with a point spread function (a.k.a. the “CLEAN beam”). The motivation for such an additional smoothing step is to reflect the angular resolution of a given instrument — the size of the CLEAN beam is determined by the diffraction limit imposed by the instrument with a given maximum baseline. Consequently, it is not possible to resolve sources beyond the instrument angular resolution. In practice, the minimum angular resolution that can be achieved by CLEAN is much larger than the instrument diffraction limit as noticed in [57]. This is also observed in the two-source simulations in Section 4.5.3: within the simulation setup (which has a maximum source separation of  $10'$  and an instrument angular resolution  $4'49.2''$ ), CLEAN cannot always resolve both sources consistently even for cases with relative large source separation and low noise levels.

In comparison, LEAP directly reconstructs the source locations in continuous space. The final resolution achievable is only related to the noise level in the visibility measurements. This is a direct consequence of enforcing continuous domain sparsity (i.e., the point source model) in the reconstruction process. LEAP tries to fit the given visibility measurements optimally (in the least square sense) with a point source model. From the experimental results in Figure 4.4 and Figure 4.5 (a), LEAP indeed manages to resolve sources separated by a distance well below the instrument angular resolution. Further, experiments with actual LOFAR observations from the Boötes field and the Toothbrush cluster (Section 4.5.5 and Section 4.5.6, respectively) also confirm the ability of LEAP to identify closely located sources reliably.

### 4.6.2 Efficient use of visibility measurements

In the experiments with actual LOFAR observations, we sub-sampled<sup>14</sup> the measurement set along different integration times: 2% and 0.25% of the total integration times is used in the single-band and the multi-band reconstruction respectively. Even with such little data, the point sources were estimated accurately by LEAP. This is potentially useful for a modern radio-interferometer like LOFAR or the next generation Square Kilometer Array (SKA), which consists of an array of numerous omnidirectional antennas. With an efficient algorithm, like LEAP, using only a fraction of the total observations can still yield reconstruction accuracy comparable to a conventional method (e.g., CLEAN) requiring significantly more integration times and longer baselines (e.g., Figure 4.6). Another potential application of LEAP could be for the Very Long Baseline Interferometry (VLBI), which has sparse coverage in the  $uv$  domain.

<sup>14</sup>This is different than taking the measurements with same number of *consecutive* integration times. See our comments at the beginning of Section 4.5.1.

## 4.7 Conclusion

We investigated the FRI-based continuous-domain sparse recovery framework in the challenging radio astronomy setting with both simulations and actual LOFAR observations, and demonstrated that the proposed approach (LEAP) can accurately resolve closely located sources under various noise conditions. In particular, we showed that it is possible to go beyond the instrument angular resolution. A reconstruction accuracy comparable to CLEAN can be achieved with significantly fewer measurements. Further, we developed a multi-band reconstruction scheme that estimates the point sources consistently among different sub-bands, and demonstrated the effectiveness of the multi-band reconstruction strategy given actual LOFAR observations. In order to facilitate the application of this new approach for source estimation in radio astronomy, we have made the Python implementation available online<sup>15</sup>.

For future work, we will consider the potential application in calibration, where an accurate source estimation is essential for subsequent imaging. Another interesting application could be for pulsar detection (from the antenna signals), given that LEAP can have reliable reconstructions even from very limited data. We will also extend the current framework to the sphere, to improve performance for instruments with a wide field-of-view, such as the *Murchison Wide-field Array* (MWA) [80].

---

<sup>15</sup>The Python package is available at <https://github.com/hanjiapan/LEAP>.

## Chapter 5

# Direction of Arrival Estimation with Microphone Arrays\*

风声雨声读书声，声声入耳。家事国事天下事，事事关心。

The sounds of the wind, the rain and words read aloud — all these sounds enter my ears. So too the affairs of the family, the country and the world — let them all be of concern to my heart.

---

顾宪成 GU, XIANCHENG

### 5.1 Introduction

In audio processing, an important problem is to estimate sound source locations from measurements with microphone arrays in audio processing. Even when the locations themselves are not of direct interest, the source location estimation often serves as a fundamental pre-processing step to more sophisticated processing, e.g., speech source separation. In this chapter, we consider a compact microphone array with sources in far-field, that is the distances between sources and microphones are much larger than the inter-microphones distance. In this regime, the wavefronts at microphone locations are very close to parallel and the source signals can be approximated as plane waves. Consequently only the directions of arrival (DOA) of these plane waves can be recovered from the measurements.

---

\*The material in this chapter is the result of joint work with Robin Scheibler, Eric Bezzam, Ivan Dokmanić and Martin Vetterli [97].

An ideal DOA estimation algorithm should be able to resolve multiple sources even when some of them are closely located and with possibly very noisy measurements. Many audio signals, like human voices, have a relatively wide spread across different frequencies. Therefore, it is imperative that a DOA estimation algorithm can exploit wide-band measurements coherently. In the classic direction of arrival estimation problem, most of the work is restricted to cases where measurements are obtained from arrays with regular geometries, e.g., uniform linear arrays, circular or spherical arrays. However, not all arrays have regular microphone layout patterns as is the case with the “Pyramic Array” [7], which is used in our experiments. It is therefore important for practical purposes that the algorithm is applicable to cases with arbitrary array layouts.

Various algorithms have been proposed to reconstruct the direction of arrival estimation problem. But it is uncommon for a DOA algorithm to meet all the desirable requirements. For example, the steered response power (SRP) methods [24], do not require a specific array geometry and are immune to coherence in signals. However, since these methods are based on beamforming, they cannot resolve closely located sources [44]. While subspace-based methods, such as MUSIC [111] and root-MUSIC [55], are able to provide high-resolution DOA estimation, these methods are fundamentally narrowband. This is because subspace methods exploit the property that source steering vectors span a subspace with bases given by the eigenvectors of the covariance matrix, which have the largest eigenvalues [111]. But the signal subspaces are sub-band frequency dependent. It is possible to extend subspace-based methods to wide-band scenarios by combining individual narrowband estimates, or transforming the array manifold at each frequency to a reference frequency with so-called “focusing matrices” (CSSM [135], WAVES [43] and TOPS [146]). In general, these methods require a search over space except when the array is a uniform linear array (ULA) [9]. It is a non-trivial task to define proper focusing matrices, which essentially are the initial guesses of the source locations. Grid-search free methods usually rely on polynomial rooting and are applicable to uniform linear arrays [9, 63] or circular arrays [65]. These methods can be further extended to arbitrary array layouts by using array interpolation [10, 19, 55, 107]. Finally, in recent years, techniques based on sparsity inducing norm optimization have gained attention [68, 83]. A notable example is [82], which allows the localization of more sources than the number of microphones available based on the Khatri-Rao subspace techniques.

In this chapter, we adapt the generic FRI reconstruction framework for DOA finding—FRIDA, which works for arbitrary array geometries. Moreover, FRIDA is capable to cope with wide-band signals coherently. The proposed method does not require grid search and is not sensitive to preprocessing akin to focusing matrices as in previous wide-band methods. We show that FRIDA can reconstruct more sources than the number of microphones as it is the cross-correlations that are used in the reconstruction algorithm. The tradeoff is that the cross-correlation-based approach cannot handle correlated signals, e.g., in an indoor environment with strong reflections from surrounding walls. Further, FRIDA can be adapted to deal with correlated sources by working with the raw microphones signals directly.

Because of the similarities with the radio astronomy point source reconstruction problem in the previous chapter, our focus will primarily be on the formulation of the DOA problem in terms of the generic FRI reconstruction. The rest of the chapter is organized as follows. In Section 5.2, we define the sources model and measurements used in the direction of arrival estimation. Then, we adapt the generic FRI-based sparse recovery algorithm by identifying a set of unknown sinusoidal samples and the linear mapping between from the sinusoids to the measurements (Section 5.3). Results with both synthetic and recorded signals are presented in



Section 5.4 before we conclude in Section 5.5.

## 5.2 Signal and measurement models

In this section, we introduce notations and setups used throughout the chapter. We describe the acoustic point source model and derive the analytical form of the cross-correlations of the received microphone signals.

For notational brevity, we use  $\mathbb{S}$  for both the unit circle  $\mathbb{S}^1$  (in 2D) and the unit sphere  $\mathbb{S}^2$  (in 3D) wherever there is no confusion. Unit propagation vectors are denoted by  $\mathbf{p}$ .

- In 2D, only the azimuth direction  $\varphi \in [0, 2\pi)$  is required so that  $\mathbf{p} \in \mathbb{S}^1$  is  $\mathbf{p} = (\cos \varphi, \sin \varphi)$ .
- In 3D, we use azimuth  $\varphi \in [0, 2\pi)$  and colatitude  $\theta \in [0, \pi)$ , thus a unit vector on the unit sphere  $\mathbf{p} \in \mathbb{S}^2$  is  $\mathbf{p} = (\cos \varphi \sin \theta, \sin \varphi \sin \theta, \cos \theta)$ .

Denote with  $\{\mathbf{r}_q \in \mathbb{R}^2\}_{q=1}^Q$  the  $Q$  microphones locations, and  $\{\mathbf{p}_k\}_{k=1}^K$  the propagation vector of  $K$  monochromatic sources in the far-field. Within a narrow band centered at frequency  $\omega$ , the baseband representation of the signal coming from direction  $\mathbf{p} \in \mathbb{S}$  is  $x(\mathbf{p}, \omega, t) = \tilde{x}(\mathbf{p}, \omega)e^{j\omega t}$ , where  $\tilde{x}(\mathbf{p}, \omega)$  is the emitted sound signal from a source located at  $\mathbf{p}$  and frequency  $\omega$ . The intensity of the sound field is

$$I(\mathbf{p}, \omega) \stackrel{\text{def}}{=} \mathbb{E} [|x(\mathbf{p}, \omega, t)|^2] = \sum_{k=1}^K \sigma_k^2(\omega) \delta(\mathbf{p} - \mathbf{p}_k), \quad (5.1)$$

where  $\sigma(\omega)_k^2$  is the power of the  $k$ -th source and  $\delta(\cdot)$  is a Dirac. We assume frame-based processing, and the expectation is over the randomness of  $\tilde{x}$  from frame to frame. As  $\tilde{x}$  carries the phase, the assumption  $\mathbb{E}[\tilde{x}] = 0$  holds.

The received signal at the  $q$ -th microphone located at  $\mathbf{r}_q$  is the integration of plane waves from all directions. Under the point source assumption, the integral reduces to a sum of  $K$  plane waves:

$$y_q(\omega, t) = \int_{\mathbb{S}} x(\mathbf{p}, \omega, t) e^{-j\omega \langle \mathbf{p}, \frac{\mathbf{r}_q}{c} \rangle} d\mathbf{p} = \sum_{k=1}^K x(\mathbf{p}_k, \omega, t) e^{-j\omega \langle \mathbf{p}_k, \frac{\mathbf{r}_q}{c} \rangle}, \quad \text{for } q = 1, \dots, Q, \quad (5.2)$$

where  $c$  is the speed of sound.

The cross-correlation between the received signals for a microphone pair  $(q, q')$  is:

$$V_{q,q'}(\omega) \stackrel{\text{def}}{=} \mathbb{E} [y_q(\omega, t) y_{q'}^*(\omega, t)] = \int_{\mathbb{S}} \int_{\mathbb{S}} \mathbb{E} [x(\mathbf{p}, t) x^*(\mathbf{p}', t)] e^{-j\omega \langle \mathbf{p}, \frac{\mathbf{r}_q - \mathbf{r}_{q'}}{c} \rangle} d\mathbf{p} d\mathbf{p}' \quad (5.3)$$

for  $q, q' \in [1, Q]$  and  $q \neq q'$ . In practice,  $V_{q,q'}$  is estimated by averaging over a finite number of frames (sometimes called snapshots); for simplicity we use the same symbol for the empirical estimate of the cross-correlation. When sources are uncorrelated, the cross-correlations take a form similar to that of (5.2):

$$V_{q,q'}(\omega) = \int_{\mathbb{S}} I(\mathbf{p}, \omega) e^{-j\omega \langle \mathbf{p}, \Delta \mathbf{r}_{q,q'} \rangle} d\mathbf{p} = \sum_{k=1}^K \sigma_k^2(\omega) e^{-j\omega \langle \mathbf{p}_k, \Delta \mathbf{r}_{q,q'} \rangle}, \quad (5.4)$$

**Table 5.1:** Summary of the linear mappings for various DOA estimation cases. See Section 5.3 for details.

Measurements <b>a</b>	Mapping <b>G</b> (2D: $g_{q,m}$ , 3D: $g_{(q,q'),(m,n)}$ )	Sinusoids <b>b</b>
<b>cross-correlations</b>		
$\sum_{k=1}^K \sigma_k^2(\omega) e^{-j\omega \langle \mathbf{p}, \Delta \mathbf{r}_{q,q'} \rangle}$	2D $(-j)^m J_m(\ \omega \Delta \mathbf{r}_{q,q'}\ _2) Y_m\left(\frac{\Delta \mathbf{r}_{q,q'}}{\ \Delta \mathbf{r}_{q,q'}\ _2}\right)$	$\sum_{k=1}^K \sigma_k^2(\omega) e^{-jm\varphi_k}$
	3D $(2\pi)^{3/2} \sum_{\ell=0}^L (-j)^\ell \frac{j^{\ell+1/2} (\ \omega \Delta \mathbf{r}_{q,q'}\ _2)^\ell}{(\ \omega \Delta \mathbf{r}_{q,q'}\ _2)^{1/2}} Y_{\ell,m}\left(\frac{\omega \Delta \mathbf{r}_{q,q'}}{\ \omega \Delta \mathbf{r}_{q,q'}\ _2}\right) N_{\ell,m} p_{\ell,n, m }$	$\sum_{k=1}^K \sigma_k^2(\omega) (\cos \theta_k)^n (\sin \theta_k)^{ m } e^{-jm\varphi_k}$
<b>microphone signals</b>		
$\sum_{k=1}^K \alpha_k(\omega) e^{-j\omega \langle \mathbf{p}, \frac{\mathbf{r}_q}{c} \rangle}$	2D $(-j)^m J_m(\ \omega \mathbf{r}_q/c\ _2) Y_m\left(\frac{\mathbf{r}_q}{\ \mathbf{r}_q\ _2}\right)$	$\sum_{k=1}^K \alpha_k(\omega) e^{-jm\varphi_k}$
	3D $(2\pi)^{3/2} \sum_{\ell=0}^L (-j)^\ell \frac{j^{\ell+1/2} (\ \omega \mathbf{r}_q/c\ _2)^\ell}{(\ \omega \mathbf{r}_q/c\ _2)^{1/2}} Y_{\ell,m}\left(\frac{\omega \mathbf{r}_q}{\ \omega \mathbf{r}_q\ _2}\right) N_{\ell,m} p_{\ell,n, m }$	$\sum_{k=1}^K \alpha_k(\omega) (\cos \theta_k)^n (\sin \theta_k)^{ m } e^{-jm\varphi_k}$

where  $\Delta \mathbf{r}_{q,q'} \stackrel{\text{def}}{=} \frac{\mathbf{r}_q - \mathbf{r}_{q'}}{c}$  is the normalized baseline.

In the next section, we first consider the DOA estimation problem where the cross-correlations of the microphone signals are taken as the measurements in the FRI-based reconstruction algorithm. Then, in Section 5.3.3 we describe the adaptations needed in order to work with microphone signals directly.

## 5.3 The acoustic direction of arrival problem

From our analysis in Chapter 2 and Chapter 3, it is essential to identify two key elements in a particular sparse recovery problem:

- a set of (unknown) uniform sinusoidal samples, which contains the sparse signal parameters to be recovered;
- a linear mapping between the given (possibly non-uniformly sampled) measurements and these uniform sinusoidal samples.

The DOA estimation problem is then formulated as a constrained optimization and solved within the FRI-based sparse recovery framework, which we developed in Chapter 2 and Chapter 3.

We will define the set of unknown sinusoidal samples and its relation to the given measurements. The mapping is described for four cases, where the measurements are either the raw microphone signal (Section 5.3.3) or the cross-correlations, in both 2D (Section 5.3.1) and 3D (Section 5.3.2), which are summarized in Table 5.1.

### 5.3.1 DOA estimation from cross-correlations in 2D

**Uniform samples of sinusoids** In 2D, the propagation vector  $\mathbf{p} = (\cos \varphi, \sin \varphi)$  is on the unit circle. Therefore, the sound field intensity (5.1) is  $2\pi$ -periodic. The Fourier series representation of  $I(\mathbf{p}, \omega)$  is:

$$I(\mathbf{p}, \omega) = \sum_{m \in \mathbb{Z}} \hat{I}_m(\omega) Y_m(\mathbf{p}),$$

where  $Y_m(\cdot)$  is the Fourier series basis:  $Y_m(\mathbf{p}) = Y_m(\varphi) = e^{jm\varphi}$ , and  $\hat{I}_m(\omega)$  is the Fourier series coefficient within a sub-band centered at frequency  $\omega$ :

$$\hat{I}_m(\omega) = \frac{1}{2\pi} \sum_{k=1}^K \sigma_k^2(\omega) e^{-jm\varphi_k}.$$

Note that  $\hat{I}_m(\omega)$  for  $m \in \mathbb{Z}$  are uniform samples of sinusoids.

**Relation with the given measurements** The cross-correlation (5.4) of the microphone signals are related with the Fourier series coefficients  $\hat{I}_m(\omega)$  as:

$$\begin{aligned} V_{q,q'}(\omega) &= \int_{\mathbb{S}^1} \sum_{m \in \mathbb{Z}} \hat{I}_m(\omega) Y_m(\mathbf{p}) e^{-j\omega \langle \mathbf{p}, \Delta \mathbf{r}_{q,q'} \rangle} d\mathbf{p} \\ &\stackrel{(a)}{=} 2\pi \sum_{m \in \mathbb{Z}} (-j)^m J_m(\|\omega \Delta \mathbf{r}_{q,q'}\|_2) Y_m\left(\frac{\Delta \mathbf{r}_{q,q'}}{\|\Delta \mathbf{r}_{q,q'}\|_2}\right) \hat{I}_m. \end{aligned} \quad (5.5)$$

Here (a) results from the Jacobi-Anger expansion [31] of complex exponentials and  $J_m(\cdot)$  is the order  $m$  Bessel function of the first kind.

Therefore, we establish a linear mapping from the uniformly sampled sinusoids  $\hat{I}_m$  to the given measurements  $V_{q,q'}$ . Concretely, denote the vectorized cross-correlation (5.3) by  $\mathbf{a}(\omega) \in \mathbb{C}^Q$ , and let the vector  $\mathbf{b}(\omega)$  be  $2\pi \hat{I}_m(\omega)$  for  $m = -M, \dots, M$ . Then the linear mapping  $\mathbf{G}(\omega)$  is:

$$[\mathbf{G}]_{(q,q'),m}(\omega) = (-j)^m J_m(\|\omega \Delta \mathbf{r}_{q,q'}\|_2) Y_m\left(\frac{\Delta \mathbf{r}_{q,q'}}{\|\Delta \mathbf{r}_{q,q'}\|_2}\right),$$

where rows of  $\mathbf{G}$  are indexed by the cross-correlation between microphone  $q$  and  $q'$ , and columns of  $\mathbf{G}$  are indexed by Fourier bins  $m$ . We can then concisely write (5.5) as  $\mathbf{a}(\omega) = \mathbf{G}(\omega)\mathbf{b}(\omega)$ .

### 5.3.2 DOA estimation from cross-correlations in 3D

**Uniform samples of sinusoids** In 3D, the propagation vector  $\mathbf{p} = (\cos \varphi \sin \theta, \sin \varphi \sin \theta, \cos \theta)$  is on the unit sphere. Similar to the 2D cases in the previous section, the sound field intensity also has a Fourier domain expansion, which involves the spherical harmonics:

$$I(\mathbf{p}, \omega) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \hat{I}_{\ell,m} Y_{\ell,m}(\mathbf{p}). \quad (5.6)$$

Here  $Y_{\ell,m}(\cdot)$  is the spherical harmonic basis of order  $\ell$  and degree  $m$ :  $Y_{\ell,m}(\mathbf{p}) = Y_{\ell,m}(\theta, \varphi) = N_{\ell,m} e^{jm\varphi} P_{\ell}^m(\cos \theta)$ , where  $P_{\ell}^m$  is an associated Legendre polynomial, and  $N_{\ell,m} = (-1)^{(m+|m|)/2} \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-|m|)!}{(\ell+|m|)!}}$  is a normalization factor (see e.g., [104] for details).

From our analysis in Chapter 2 (Section 2.4.4 and Appendix 2.D), we know that the spherical harmonic coefficients  $\hat{I}_{\ell,m}$  are linear combinations of uniform sinusoidal samples:

$$b_{n,m} = \sum_{k=1}^K \sigma_k^2(\omega) (\cos \theta_k)^n (\sin \theta_k)^{|m|} e^{-jm\varphi_k} \quad (5.7)$$

and

$$\hat{I}_{\ell,m} = N_{\ell,m} \sum_{n=0}^{\ell-|m|} p_{\ell,n,|m|} b_{n,m}, \quad (5.8)$$

where  $p_{\ell,n,|m|}$  are some fixed coefficients that can be pre-computed (see derivations in Appendix 2.E).

**Relation with the given measurements** We can decompose the connection from the uniform sinusoidal samples to the cross-correlation measurements into two stages:

- (i) the linear mapping from  $b_{n,m}$  to the spherical harmonic coefficients  $\hat{I}_{\ell,m}$  (5.8);
- (ii) the linear mapping from the spherical harmonic coefficients to the cross-correlations (5.6):

$$\begin{aligned} V_{q,q'}(\omega) &= \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \hat{I}_{\ell,m} \iint_{\mathbb{S}^2} e^{-j\omega \langle \mathbf{p}, \Delta \mathbf{r}_{q,q'} \rangle} Y_{\ell,m}(\mathbf{p}) d\mathbf{p} \\ &\stackrel{(a)}{=} (2\pi)^{3/2} \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} (-j)^{\ell} \frac{J_{\ell+1/2}(\|\omega \Delta \mathbf{r}_{q,q'}\|_2)}{(\|\omega \Delta \mathbf{r}_{q,q'}\|_2)^{1/2}} Y_{\ell,m} \left( \frac{\omega \Delta \mathbf{r}_{q,q'}}{\|\omega \Delta \mathbf{r}_{q,q'}\|_2} \right) \hat{I}_{\ell,m}, \end{aligned}$$

where  $J_{\ell}(\cdot)$  is Bessel function of the first kind and (a) is from the expansion of complex exponentials with spherical harmonic basis (see eq. (2.37) in [104]).

In the actual implementation, the infinite summation in the spherical harmonic expansion is truncated up to a certain maximum order  $L$  and degree  $M \leq L$ . Denote the vectorized cross-correlations (5.3) by  $\mathbf{a}(\omega) \in \mathbb{C}^Q$ , and the uniform sinusoidal samples (5.7) as a vector  $\mathbf{b}(\omega)$  for  $m = -M, \dots, M$ , then the linear mapping in the constrained FRI reconstruction algorithm  $\mathbf{G}(\omega)$  is:

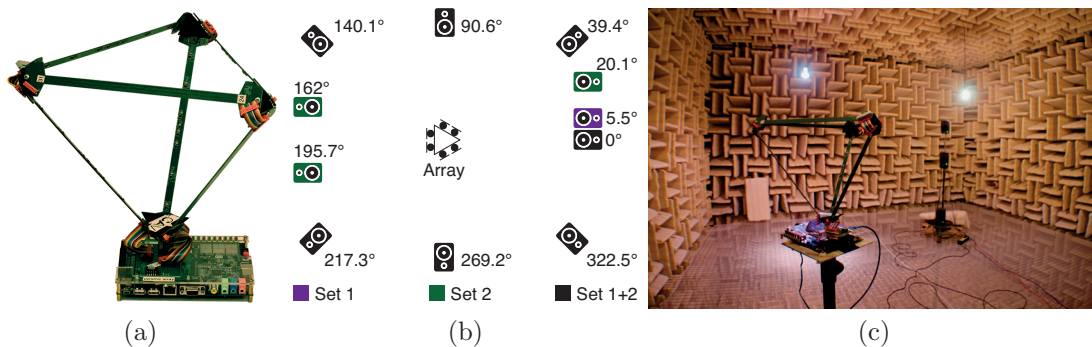
$$[\mathbf{G}]_{(q,q'),(m,n)}(\omega) = (2\pi)^{3/2} \sum_{\ell=0}^L (-j)^{\ell} \frac{J_{\ell+1/2}(\|\omega \Delta \mathbf{r}_{q,q'}\|_2)}{(\|\omega \Delta \mathbf{r}_{q,q'}\|_2)^{1/2}} Y_{\ell,m} \left( \frac{\omega \Delta \mathbf{r}_{q,q'}}{\|\omega \Delta \mathbf{r}_{q,q'}\|_2} \right) N_{\ell,m} p_{\ell,n,|m|}.$$

### 5.3.3 DOA estimation from raw microphone signals

In the previous two cases, we take the cross-correlations of the received microphone signals as the measurement in the FRI-base reconstruction. In the case of correlated sources, e.g, in an indoor environment, the cross-correlation based approach is no longer applicable. However, it turns out that very little modifications are needed in order to estimate the source direction of arrival directly from the received microphone signals—hence overcoming the limitations of working only with uncorrelated sources in the previous cases. More specifically, the microphone measurements, the normalized microphone locations (by the speed of sound), and the source signals replace the cross-correlation measurements, the difference of microphone locations, and the source signals variances, respectively:

$$V_{q,q'}(\omega) \rightarrow y_q(\omega), \quad \Delta \mathbf{r}_{q,q'} \rightarrow \mathbf{r}_q/c, \quad \sigma_k^2(\omega) \rightarrow \alpha_k(\omega) \stackrel{\text{def}}{=} \tilde{x}(\mathbf{p}_k, \omega).$$

We summarize the linear mappings of both 2D and 3D cases in Table 5.1.



**Figure 5.1:** (a) Pyramic array, a compact microphone array with 48 MEMS microphones distributed on the edges of a tetrahedron. For the 2D DOA experiments, only the top triangle is used; while in 3D DOA experiments, measurements from all 48 microphones are used. (b) Locations of the loudspeakers and microphone array in experiments. (c) Anechoic chamber used for the 3D DOA estimation experiments.

## 5.4 Results

In this section, we demonstrate the effectiveness of the proposed algorithm in simulations and experiments with actual recordings. We compare FRIDA with other wide-band algorithms: MUSIC [111], SRP-PHAT [44], CSSM [135], WAVES [43], and TOPS [146].

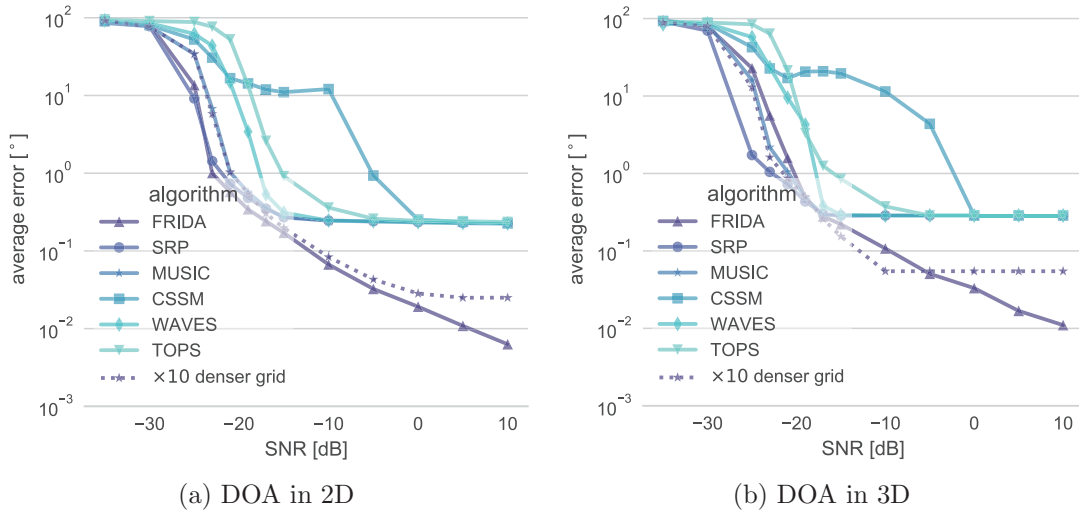
In experiments, sampling frequency of the recordings is fixed at 16 kHz. Narrow-band sub-carriers are extracted from a 256-tap short-time Fourier transform (STFT) with a Hanning window and no overlap. We use a tetrahedron-shaped array of 24 microphones (the “Pyramic array” [7]). Each side is 30 cm long and contains 8 microphones. The spacing between microphone pairs ranges from 8 mm to 25 cm. We use the Pyramic arrays to collect recordings in experiments (see Figure 1.2 for both 2D and 3D setups). We have taken 20 sub-bands uniformly spaced between 2.5 kHz and 4.5 kHz for all algorithms except TOPS, which requires low frequency sub-bands in order to have a reasonable reconstruction. For TOPS, we have chosen 60 sub-bands from 100 Hz to 5 kHz. For algorithms based on grid search, we have chosen a grid with 720 uniform grid points within  $[0^\circ, 360^\circ)$  in 2D and 20 000 approximately equidistant grid points on the sphere [71] with median inter-point distance of  $1.53^\circ$ .

The reconstruction errors are measured with the great-circle distance<sup>1</sup>. The correspondence between the reconstructed source locations and the ground truth ones are obtained by finding the permutation that minimizes the distance between the permuted source locations.

### 5.4.1 Influence of noise

We study the performance of various algorithms with measurements at different noise levels in simulations. Sound propagation is simulated by applying fractional delay filters to the source signal. The microphone signals are generated based on the array geometry of the Pyramic arrays, which are used in experiments with actual speech recordings (Section 5.4.3). Gaussian white noise was added to the received microphone signals based on the desired SNR levels. We estimated

<sup>1</sup>In a planar setup (i.e., DOA in 2D), this corresponds to choosing  $0^\circ$  latitudes for all sources.



**Figure 5.2:** Average DOA reconstruction error in (a) 2D and (b) 3D, from simulated microphone signals with different noise levels. The DOA estimation errors are averaged over 500 different noise realizations at each noise level. With denser grid, grid-search based algorithms achieves lower DOA estimation errors but at the cost of higher computational burden.

the covariance matrix by averaging over 256 snapshots of 256 samples each. The reconstruction errors are averaged over 500 different noise realizations at each noise level.


From the average DOA estimation errors in Figure 5.2, FRIDA along with SRP-PHAT and MUSIC are the most robust methods. For cases with little noise, the DOA estimation of grid search based methods all stagnate to an error that is related to the grid search step-size; while FRIDA, which is based on the continuous-domain sparse recovery, does not have such a limitation and continues decreasing the DOA estimation errors with increased SNRs.

With a denser grid, the minimum DOA estimation error of grid search based algorithms might decrease but at the expense of increased computational cost, which is proportional to the total number of grid points. We repeat the DOA estimation with a grid that consists of 10 times more grid points. The DOA estimation errors is show as dotted line in Figure 5.2. The saturation level indeed decreases. Even though a denser grid improves the estimation accuracy of one source, it is not clear whether the dense grid will benefit the resolution of close sources, which is what we will investigate in the next section.

## 5.4.2 Resolving close sources

Next, we study the minimum separation that an algorithm can resolve distinct sources in a 2D setup. Two sources are separated by an angle that varies linearly on a log scale from  $1^\circ$  to  $45^\circ$ . Each source emits Gaussian white noise with equal power. The average error is computed over 120 uniformly space source locations<sup>2</sup> from  $0^\circ$  to  $120^\circ$ . The source directions of arrival were estimated from microphone signals with 0 dB SNR. A source is considered as successfully recovered if the reconstructed DOA is less than half the separation between the two sources from

<sup>2</sup>Due to the symmetry of the Pyramic array, we do not need to use the full  $360^\circ$ .



**Figure 5.3:** Average number of resolved sources against different separation distances between two sound sources (noise level in microphone measurements SNR = 0 dB).

DOA	FRIDA	MUSIC	SRP-PHAT
$0^\circ$	$-0.5^\circ$	$1.6^\circ$	$1.4^\circ$
$5.5^\circ$	$4.6^\circ$	$-93.9^\circ$	$-38.1^\circ$

**Table 5.2:** Resolving two closely located sources (DOAs:  $0^\circ$  and  $5.5^\circ$ ) from actual speech recordings in an audio room (see experimental setup in Figure 5.1(b)).

the ground truth source location. Unfortunately, we ran into numerical difficulties with CSSM, WAVES and TOPS when two sources are too close to each other. Additionally, these methods do not appear to be competitive based on the simulations in the previous section. Hence, we have omitted them from this comparison.

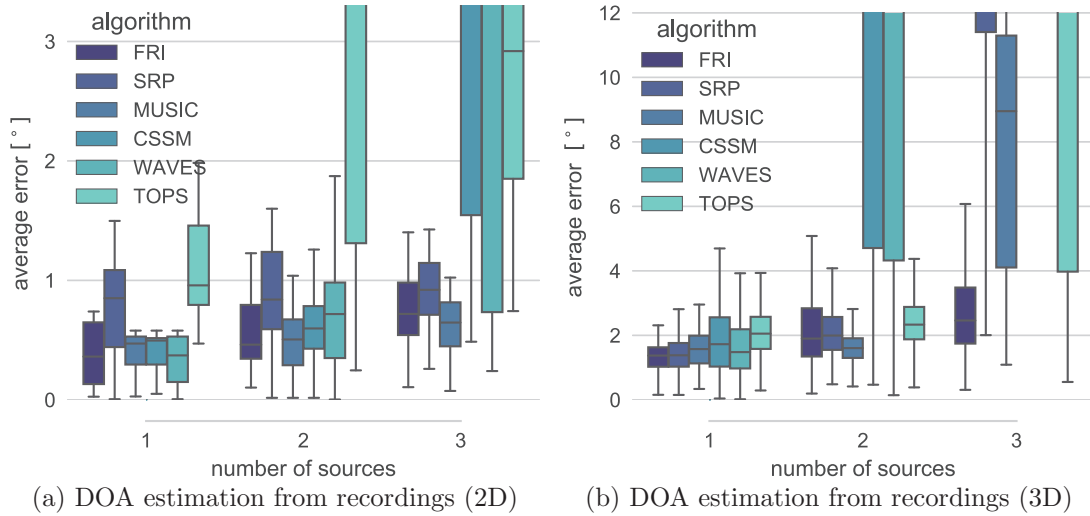
FRIDA significantly outperforms MUSIC and SRP-PHAT (Figure 5.3): Sources that are separated by  $3.56^\circ$  or more are always resolved; while the closest contenders, MUSIC can only distinguish two sources that are at least  $10.80^\circ$  apart and SRP-PHAT needs to have a separation of  $20.36^\circ$ . The distance between the adjacent grid points is  $0.5^\circ$ , which is much smaller than the minimum separation that either MUSIC or SRP-PHAT can resolve two distinctive sources.

### 5.4.3 Results with actual speech recordings

Finally, we perform experiments with recorded data to validate the algorithm in real-world conditions for 2D DOA and 3D DOA estimations.

For the 2D DOA experiment, the Pyramic array was placed at the center of eight loudspeakers (Figure 5.1(b), Set 1). All loudspeakers were between 1.45 m and 2.45 m away from the array. Recordings were made with all possible combinations of one, two, and three speakers playing simultaneously distinct speech segments. Two of the speakers were separated by  $5.5^\circ$  from each other. We will use these two loudspeakers to test the resolving power of different algorithms.

For the 3D DOA experiment, the Pyramic array was fixed on top of a turn table inside an




**Figure 5.4:** Average DOA estimation errors with different algorithms applied to the recorded speech signals in (a) 2D and (b) 3D setups, respectively.

anechoic chamber. Three loudspeakers located at different heights are between 3.94 m and 4.08 m away from the array. Instead of moving loudspeakers around to different locations, the array was rotated to different angles (using the turn table). Each time, only one loudspeaker played speech segments. In order to test the DOA estimations with multiple sources, recordings of individual sources (at certain orientation angles) were added up. This inevitably increased noise levels in the re-synthesized recordings compared with the direct recording of multiple simultaneously playing loudspeakers. However, this does not seem to be an issue based on our experience: For instance, a synthesized recording of three sources has an estimated SNR level around 25 dB. We select recordings that correspond to 61 uniformly spaced azimuth values from  $0^\circ$  to  $120^\circ$  when DOA estimations are performed for 1 source; and 16 uniformly spaced azimuth values from  $0^\circ$  to  $30^\circ$  when DOA estimations are performed for multiple sources. Combinations of up to three sources at different locations are tested in the experiments.

The statistics of the reconstruction errors for the different algorithms are shown in Figure 5.4 in both 2D and 3D setups. In 2D experiments, the coherent methods WAVES and CSSM perform well for one and two sources, but break down for three sources. TOPS has difficulty in estimating multiple sources. FRIDA, MUSIC and SRP-PHAT perform the best with a median error within one degree from the ground truth. In 3D experiments, WAVES and CSSM fails to perform well when there is more than one source. Similar to the 2D experiments, FRIDA and MUSIC give most reliable DOA estimations. However, FRIDA is the only algorithm that gives reliable reconstructions when 3 sources are present. The median DOA estimation error is less than  $3^\circ$ . FRIDA distinguishes itself from the conventional methods in resolving closely spaced sources. This is highlighted in Table 5.2 for the 2D setup, where the reconstructed DOA for the sources located at  $0^\circ$  and  $5.5^\circ$  is shown. While all three methods correctly identify one source, only FRIDA is able to resolve the second one. We have observed similar performance with FRIDA in resolving close sources in 3D cases. In Figure 5.5, we show one example in the 3D setup. Two of the three sources are closely located with a great-circle distance of  $12.76^\circ$ . In the least square

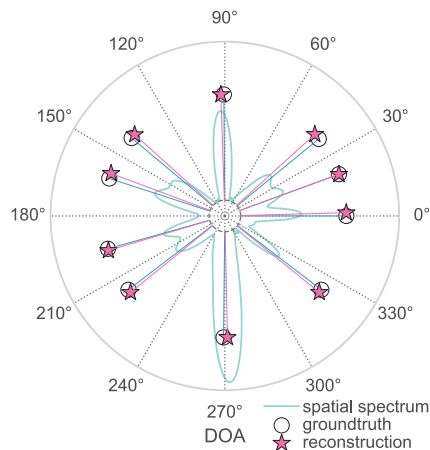




**Figure 5.5:** Accurate direction of arrival estimation with FRIDA of closely located sources from recordings of speech signals in an anechoic chamber (The estimates SNR in the recording: 27.26 dB. The noise level is estimated from the silence recording). FRIDA accurately reconstructs all source locations (average estimation error: 1.64°); while a grid-search based algorithm, e.g., MUSIC, fails to resolve two closely located sources even with a dense grid (see Figure 1.2(b) in the Introduction Chapter).

estimate, both sources are contained within one bright blob. With FRIDA, all three DOAs are accurately estimated. The average reconstruction error is 1.64°.

Finally, we tested the ability of FRIDA to accurately resolve a number of sources larger than the number of microphones used. We placed ten loudspeakers (Figure 5.1(b), Set 2) around the Pyramic array and recorded them simultaneously playing white noise. Then, we discarded the signals of all but 9 microphones (out of 24) and applied FRIDA to the recordings. The algorithm successfully reconstructed all DOAs within 2° of the ground truth, as shown in Figure 5.6. Note that none of the subspace methods can achieve this result. While SRP-PHAT is not limited in this way, its resolution is lower (its error is  $\sim 4^\circ$  on this recording).



**Figure 5.6:** Reconstruction of 10 sources from only 9 microphones. The average DOA estimation error is within 2°.

## 5.5 Conclusion

In this chapter, we introduced FRIDA, a new algorithm for DOA estimation of sound sources. FRIDA, which is based on the finite rate of innovation sampling technique that we developed in Chapter 2 and Chapter 3, estimates DOAs in continuous domain without grid search. The approach is applicable to arrays with arbitrary geometries as is demonstrated with both simulations and real experiments with the Pyramic arrays. Wide-band measurements can be efficiently exploited to obtain more robust reconstruction. We demonstrated that FRIDA compares favorably to the state-of-the-art, and outperforms all other algorithms in resolving close sources. Finally, FRIDA is notable for resolving more sources than microphones, as demonstrated experimentally with recorded signals.

## Chapter 6

# Image up-sampling via Edge Etching

不闻不若闻之，闻之不若见之。见之不若知之，知之不若行之。学至于行而止矣。

Tell me and I will forget. Show me and I will remember. Involve me and I will understand. Step back and I will act.

---


荀子·儒孝 *Xunzi · Ruxiao*

### 6.1 Introduction

The goal of image up-sampling is to reconstruct a high resolution image from a single low resolution one, which is obtained from the high resolution image via lowpass filtering followed by down sampling. A closely related classic problem is super-resolution, where the objective is to fuse a sequence of low resolution images into a higher resolution image [51, 127]. Image up-sampling offers a promising way to overcome resolution limitations of imaging sensors due to cost or physical constraints. Various approaches have been proposed to address the inherently ill-conditioned restoration problem and they may be broadly categorized into three classes: interpolation-based methods, reconstruction-based methods and example-based learning approaches.

Interpolation-based methods address the missing pixel values in the up-sampled image with a predefined interpolating kernel, e.g., bicubic spline [72], from the given pixel values in the neighborhood. Generic interpolation approaches are simple and intuitive but the up-sampled image tends to have annoying ringing artifacts and blurred edges. Improved interpolation methods with adaptive interpolating kernels (based on local geometric regularities [78, 121]) have been proposed so as to have more faithful up-sampling results around edges.

Reconstruction-based methods address the up-sampling problem by exploiting prior knowledge of the up-sampled image, e.g., piecewise smoothness [4, 87], or priors on edge smooth-



**Figure 6.1:** The proposed image up-sampling algorithm preserves sharp edges and reduces ringing artifacts. (a) Ground truth high resolution image; Up-sampled images (factor: 3) with (b) bicubic interpolation (PSNR = 29.36 dB), (c) Kim & Kwon [73] (PSNR = 32.98 dB), and (d) the proposed approach (PSNR = 33.70 dB), respectively.

ness [40, 53, 119]. The up-sampled image is then reconstructed by minimizing a regularization term subject to a data fidelity constraint, which requires that the lowpass filtered and down-sampled version of the up-sampled image should match the given low resolution image.

A different line of work based on machine learning tries to estimate the mapping functions between the high and low resolution image pairs (typically in the image patch space) from either internal or external examples [28, 54, 60, 64, 73, 81, 100, 136, 144, 147]. We may view the learning process as an implicit way to find the priors that one should impose on the up-sampled image. The reconstruction quality of the learning-based approaches is dependent on the training dataset: when the training and testing images are of different nature, the up-sampling performance usually deteriorates.

In this chapter, we propose a reconstruction-based image up-sampling approach where we incorporate a global edge model estimated from the given low resolution image in the up-sampling process. Motivated by our previous work on sampling curves with finite rate of innovation [95], we recast the edge modeling problem as finding a mask function such that its product with the image gradients is minimized. In a sense, the mask function, which vanishes around edges, serves as a complement to the gradients. Subsequently, we enforce the reconstructed model as an edge preserving constraint for image up-sampling. Different from [119], our goal is not to

match a gradient profile for the up-sampled image directly but rather using the mask function to “kill” (i.e., *annihilate*) large gradients. We show through numerical experiments that the estimated edge model is indeed beneficial in preserving sharp edges and achieving state of the art up-sampling performances (see an example in Figure 6.1). Our contributions are twofold:

- 1) We propose an edge modeling framework, which gives a continuous domain representation of image edges, both in a local area and for the whole image. We give an example of one possible application of the reconstructed local edge model for light-field images.
- 2) We formulate an edge-friendly image up-sampling problem by incorporating the reconstructed global edge model in the up-sampling process. An efficient algorithm is proposed to solve the constrained optimization problem based on alternating direction method of multipliers (ADMM).

The rest of the chapter is organized as follows. We briefly review our work on sampling curves with *finite rate of innovation* (FRI) and highlight an important spatial domain interpretation of the theoretical results in Section 6.2. Based on this new interpretation, we propose an edge modeling framework in Section 6.3, which serves as an edge-preserving regularization in image up-sampling. We devise a fast algorithm to solve the resultant optimization problem in Section 6.4. Experimental results on both edge modeling and image up-samplings are presented in Section 6.5 before we make the concluding remarks in Section 6.6.

## 6.2 Annihilating curves with a mask function

In this section, we will introduce a continuous domain mask function that allows us to estimate image edges very accurately. Our approach is based on recent work [95] that extends the classic FRI sampling and reconstruction framework [11, 14, 84, 115, 133] to two dimensional curves.

To introduce this approach, consider a continuous domain edge image  $I_C(x, y)$ , which is analytic almost everywhere except on some curve  $C$ , where the image becomes discontinuous (Figure 6.2(a)). Because of analyticity, the (complex) derivative of the edge image  $\partial I_C(x, y) = \left(\frac{\partial}{\partial x} + j\frac{\partial}{\partial y}\right)I_C(x, y)$  is different from zero only on the curve (Figure 6.2(b)). We would like to retrieve the edge positions in the continuous domain from a sampled version of  $I_C$  (i.e., find the curve  $C$ ).

Instead of directly looking for non-zero regions of the derivative image, we introduce a mask function  $\mu(x, y)$  which *annihilates* the curve; i.e., its product with the derivative image is zero (Figure 6.2):


$$\mu(x, y) \cdot \partial I_C(x, y) = 0. \quad (6.1)$$

The curve is then given by the zero-crossings of the reconstructed mask function  $\mu(x, y) = 0$ .

In [95], we showed that if the mask function has a particular form, i.e., a linear combination of sinusoids (with some unknown weights  $c_{k,l}$ )

$$\mu(x, y) = \sum_{k=-K_0}^{K_0} \sum_{l=-L_0}^{L_0} c_{k,l} e^{j\frac{2\pi k}{\tau_x}x + j\frac{2\pi l}{\tau_y}y}, \quad (6.2)$$

we can exactly reconstruct the curve from a few samples of the edge image. This is because the



**Figure 6.2:** Reconstruction of the positions of discontinuities in an edge image by finding a mask function that annihilates the derivative image (6.1). (a) Edge image, which is discontinuous across the curve  $C$ ; (b) The derivative image; (c) The mask function, which vanishes on the curve.

coefficients  $c_{k,l}$  satisfy the following *annihilation equations*<sup>1</sup>:

$$\sum_{k,l} c_{k,l} \widehat{\partial I}_C \left( \omega_x - \frac{2\pi k}{\tau_x}, \omega_y - \frac{2\pi l}{\tau_y} \right) = 0 \quad \forall \omega_x, \omega_y, \quad (6.3)$$

where  $\widehat{\partial I}_C = (\omega_x + j\omega_y) \hat{I}_C(\omega_x, \omega_y)$ . On the other hand, the Fourier transforms  $\hat{I}_C$  at particular frequencies have a one to one correspondence with the ideally low pass filtered samples of the edge image. Hence, with sufficiently many samples, we can solve for the coefficients  $c_{k,l}$  and uniquely define the mask function and the curve (see [95] for more details).

In practice, the edge model defined in (6.2) is not always suitable for natural images. Since the Fourier basis in (6.2) is *globally varying*, we have limited control over the local topology of the curve. In the following sections we will introduce more local, hence more practical, mask functions. However, it will not be possible to exactly satisfy (6.1). Instead, we will look to minimize


$$\|\mu \cdot \partial I\| \quad (6.4)$$

with respect to a certain norm  $\|\cdot\|$  to be specified (e.g.,  $\ell_2$  norm) and hence find the mask function that best annihilates the derivative image.

### 6.3 Annihilation-driven edge modeling

Although the mask function given in (6.2) allowed us to exactly retrieve the curve from a few measurements, it was not practical for natural images. In this section, we will propose alternative continuous-domain mask functions and efficient algorithms to extract their parameters. We will utilize a simple local linear model (Section 6.3.1). Subsequently, these local edge models are combined to build a global one, which is adapted to the local geometries of image edges (Section 6.3.2).

<sup>1</sup>Note that (6.3) is nothing more than the Fourier transform of (6.1).



**Figure 6.3:** Approximate image edges as a straight line in a localized block.

### 6.3.1 Block-wise linear edge model

Observe that if we magnify an image by a large enough factor, then we can well approximate image edges locally as a straight line (Figure 6.3):  $ax + by + c = 0$ , where  $a, b$  and  $c$  are parameters to be determined and  $x, y$  are horizontal and vertical coordinates, respectively. In the framework of FRI curve annihilation, this amounts to considering a mask function:

$$\mu(x, y) = ax + by + c,$$

whose roots define the curve. Based on the spatial domain interpretation in the previous section, the optimal parameters are therefore the ones that minimize (6.4). Such a simple linear edge model only makes sense when we consider the problem in a local area. Specifically, the edge parameters  $(a, b, c)$  in a block centered around  $(x_0, y_0)$  are

$$\begin{aligned} \min_{a,b,c} \quad & \iint |\mu(x - x_0, y - y_0) \cdot \partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy \\ \text{subject to} \quad & a^2 + b^2 = 1, \end{aligned} \quad (6.5)$$

where  $w(\cdot, \cdot)$  is a localized window, e.g., a Gaussian with a certain standard deviation  $\sigma$ ,  $\frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right)$  and  $\partial I(x, y)$  are the image gradients. The constraint is imposed to avoid the trivial solution where all three coefficients are zero. In practice, we may choose the block centers as the Canny edge positions [23] and repeat the same process for all blocks in the entire image (see Appendix 6.A for a non-iterative efficient implementation).

### Reconstruction of line parameters

Note that (6.5) is a simple quadratic minimization subject to one quadratic constraint. The associated Lagrangian of the constrained optimization is  $\mathcal{L}(\mathbf{d}, \lambda) = \mathbf{d}^T \mathbf{A} \mathbf{d} + \lambda(1 - \mathbf{d}^T \mathbf{B} \mathbf{d})$ , where

$$\mathbf{d} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Here  $\mathbf{A}$  is a  $3 \times 3$  Hermitian symmetric matrix, whose entries can be obtained directly from the derivative image  $\partial I(x, y)$ . By setting the derivatives of the Lagrangian to zero, (6.5) reduces to

an eigen-decomposition problem, which has a closed form solution:

$$\mathbf{d} = \begin{bmatrix} A_{1,3}(A_{2,2} - \lambda) - A_{1,2}A_{2,3} \\ (A_{1,1} - \lambda)A_{2,3} - A_{1,3}A_{1,2} \\ A_{1,2}^2 - (A_{1,1} - \lambda)(A_{2,2} - \lambda) \end{bmatrix}. \quad (6.6)$$

Here  $A_{i,j}$  for  $i, j = 1, 2, 3$  is the  $i$ -th row and  $j$ -th column entry of  $\mathbf{A}$ . The Lagrange multiplier  $\lambda$  is a solution of a quadratic equation<sup>2</sup>

$$\alpha_1 \lambda^2 + \alpha_2 \lambda + \alpha_3 = 0, \quad (6.7)$$

where  $\alpha_i$  for  $i = 1, 2, 3$  are constants that can be computed directly from  $\mathbf{A}$ 's entries as

$$\begin{aligned} \alpha_1 &= A_{3,3}, \quad \alpha_2 = A_{1,3}^2 + A_{2,3}^2 - A_{2,2}A_{3,3} - A_{1,1}A_{3,3}, \\ \alpha_3 &= A_{1,1}A_{2,2}A_{3,3} - A_{1,1}A_{2,3}^2 - A_{1,2}^2A_{3,3} - A_{1,3}^2A_{2,2} + 2A_{1,3}A_{1,2}A_{2,3}. \end{aligned}$$

Once we have reconstructed the line parameters, we can further define the *edge point*  $(x_e, y_e)$  for each block, which should be located on the reconstructed line while having the shortest distance to the block centre:

$$\begin{aligned} (x_e, y_e) &= \arg \min_{x,y} (x - x_0)^2 + (y - y_0)^2 \\ &\text{subject to } \mu(x - x_0, y - y_0) = 0, \end{aligned} \quad (6.8)$$

i.e.,  $x_e = x_0 - \frac{ac}{a^2+b^2}$ , and  $y_e = y_0 - \frac{bc}{a^2+b^2}$ . In comparison with conventional edge detection algorithms, e.g., the Canny edge detector [23], the reconstructed edge points are no longer restricted to be located on pixel centers.

Notice that the edge model (and hence the edge positions) within each block is directly related to the entries of  $\mathbf{A}$ , which is different from block to block. A naive implementation is to evaluate the matrix for each block, which is computationally inefficient. It turns out that we can reinterpret the block-wise process as convolutions [77], which is very fast. Consequently, the edge models for the entire image are reconstructed by applying a few *linear filtering*. Please refer to Appendix 6.A for a detailed discussion of the implementation.

### Further refinements with anisotropic windows

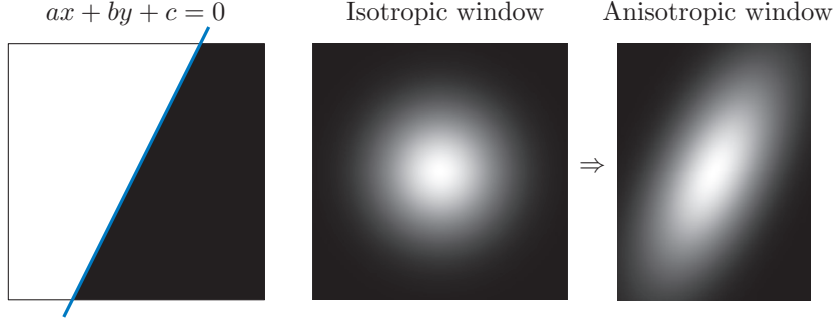
Based on the reconstructed linear edge model, we can adapt the window  $w(x, y)$  to further refine the estimated line parameters. For instance, we may scale the size (the standard deviations) of the Gaussian window and align it along the (estimated) edge direction (Figure 6.4). Another possible refinement is that we can now place the window around the reconstructed edge point  $(x_e, y_e)$  rather than the Canny edge point as was the case in the first round estimation.

Specifically, if the estimated linear edge model is  $ax + by + c = 0$ , then the Gaussian window is adapted as

$$\sigma \rightarrow \begin{cases} \sigma_x \\ \sigma_y \end{cases}, \quad \text{and} \quad \begin{bmatrix} x \\ y \end{bmatrix} \rightarrow \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix},$$

<sup>2</sup>The two solutions  $\lambda_{\min}$  and  $\lambda_{\max}$  correspond to cases where the annihilation errors are minimized and maximized, respectively.





**Figure 6.4:** Adapt window shape based on the reconstructed edge model.

i.e.,

$$w(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left(-\frac{(x \cos \theta - y \sin \theta)^2}{2\sigma_x^2} - \frac{(x \sin \theta + y \cos \theta)^2}{2\sigma_y^2}\right). \quad (6.9)$$

Here  $\sigma_x$  and  $\sigma_y$  are rescaled based on the ratio of the maximum and minimum eigenvalues from (6.7), e.g.,  $\sigma_x = \sigma \left(\frac{\lambda_{\max}}{\lambda_{\min}}\right)^\kappa$ ,  $\sigma_y = \sigma \left(\frac{\lambda_{\min}}{\lambda_{\max}}\right)^\kappa$  for some  $\kappa$ . The rotation angle  $\theta$  is determined by the slope of the linear edge model:  $\sin \theta = \frac{a}{\sqrt{a^2+b^2}}$  and  $\cos \theta = \frac{b}{\sqrt{a^2+b^2}}$ . Experimentally, with the adapted window, more reliable estimations of the edge parameters are obtained (see Section 6.5.1).

### 6.3.2 A locally adaptive global mask function

The block-wise linear edge model gives reliable estimations of edge orientation and sub-pixel edge position within each localized area. It may already be useful in its own right (see examples in Section 6.5.1). But is it possible to build a global edge model for the whole image by combing these local models? We may then use the global mask function to assign position-dependent weights (to differentiate between the edge and the smooth areas) in the subsequent processing and better preserve sharp edges [95].

One intuitive way to build a global model is to take a weighted sum of the local linear edge models over all blocks:

$$\mu(x, y) = \sum_{n=1}^N (a_n(x - x_n) + b_n(y - y_n) + c_n) w_n(x - x_n, y - y_n), \quad (6.10)$$

where  $(x_n, y_n)$  are the edge positions from (6.8);  $(a_n, b_n, c_n)$  are the reconstructed line parameters for the  $n$ -th block; and  $w_n(\cdot, \cdot)$  is a localization window (e.g., an isotropic Gaussian or the adapted one (6.9)). However, the straightforward approach has two noticeable drawbacks:

- Because of the localization windows, the global mask function (6.10) vanishes in areas that are not in the proximity of any block centers  $(x_n, y_n)$ . Typically, these areas correspond to the smooth part in an image. More concretely, consider an image up-sampling problem, where we use the annihilation error  $\|\mu \cdot \partial I\|$  to better preserve edges. Large gradients (correspondingly sharp edges) are encouraged wherever  $\mu$  is small, which includes both the actual edges and the smooth area had we used a mask function (6.10).

- The parameters  $(a_n, b_n, c_n)$ , which minimize the annihilation error within *each block*, fail to take inter-block interferences into account and hence are suboptimal in minimizing the annihilation error for the entire image.

We address these two issues by considering a modified edge model, whose parameters are reconstructed by minimizing the annihilation error globally.

### Mask parametrization

Since we do not expect to find any edges in smooth areas, we may choose the line parameters explicitly as  $a = b = 0$  and  $c = 1$  for blocks that are located in the smooth areas:

$$\mu(x, y) = \sum_{n=1}^N (a_n(x - x_n) + b_n(y - y_n) + c_n)w_n(x - x_n, y - y_n) + \underbrace{\sum_{m=1}^M \tilde{w}_m(x - \tilde{x}_m, y - \tilde{y}_m)}_{\mu_{\text{smooth}}(x, y)}, \quad (6.11)$$

where  $w_n$  and  $\tilde{w}_m$  are Gaussian windows with certain standard deviations centered around edges or in the smooth area, respectively. In the implementation, we may choose the block centers  $(x_n, y_n)$  as the sub-pixel accurate edge positions from the block-wise estimation (6.8) and  $(\tilde{x}_m, \tilde{y}_m)$  as the vertices of a sufficiently dense uniform grid in the smooth region. Notice that here  $\mu_{\text{smooth}}$  is independent of the unknown parameters  $\{a_n, b_n, c_n\}_{n=1}^N$  and the total degrees of freedom of the global mask function is still determined by the number of edge blocks  $N$ .


### Global optimization

Similarly to the block-wise edge modeling, the global edge model is reconstructed by minimizing the  $\ell_2$  norm of the annihilation errors:  $\|\mu \cdot \partial I\|^2$ . If we follow a similar formulation as in (6.5), we would like to minimize the quadratic function subject to several quadratic constraints  $a_n^2 + b_n^2 = 1$  for  $n = 1, \dots, N$ . For a natural image,  $N$  is typically very large, e.g., in the order of  $10^3$ , which precludes having a closed form solution for the quadratic constrained quadratic program (QCQP) in general.

Alternatively, we may solve the quadratic minimization subject to a set of *linear* constraints, e.g.,  $c_n = \text{const.}$  for  $n = 1, \dots, N$ . But what would be a good choice for this constant? Remember that  $(x_n, y_n)$  in (6.11) are the reconstructed edge positions from the block-wise estimation. Therefore, it is reasonable to assume that it should also locate approximately on the global edge:  $\mu(x, y)|_{x=x_n, y=y_n} \approx 0$ . Experimentally, we have observed that the reconstructed mask function vanishes around the edge points  $(x_n, y_n)$  with a simple choice as  $c_n = -1$ . The global edge model is reconstructed as:

$$\begin{aligned} & \min_{\substack{a_1, \dots, a_N \\ b_1, \dots, b_N \\ c_1, \dots, c_N}} \iint |\mu(x, y) \cdot \partial I(x, y)|^2 dx dy \\ & \text{subject to } c_n = -1 \quad \text{for } n = 1, \dots, N, \end{aligned}$$

which has a closed form solution and can be implemented efficiently.



**Figure 6.5:** Image up-sampling aims at reconstructing a high resolution image from a single low resolution image  $I_0$ , which is obtained from the high resolution image  $I_{\text{high res}}$  via lowpass filtering followed by down sampling.

## 6.4 Application on edge-preserving image up-sampling

It is challenging to obtain a high resolution image from a single low resolution one. This problem is different from the conventional super-resolution problem that aims at fusing multiple low resolution images to form a high resolution one [51, 127, 132]. In the up-sampling setup (Figure 6.5), the image reconstruction is formulated as a severely underdetermined inverse problem. A typical strategy to better-condition the restoration process is to incorporate prior knowledge on the reconstructed high resolution image, e.g., image smoothness.

As an example, the standard (thin plate spline) interpolation amounts to minimizing the Laplacian of the up-sampled image subject to the data-fidelity constraint, which requires that the up-sampled image should reproduce the low resolution image after going through the same data-acquisition process:

$$\begin{aligned} \min_{\mathbf{I}} \quad & \|\Delta \mathbf{I}\|_2^2 \\ \text{subject to} \quad & \Phi \mathbf{I} = \mathbf{I}_0. \end{aligned} \quad (6.12)$$


Here  $\mathbf{I}_0$  and  $\mathbf{I}$  are the given low resolution image and the up-sampled image, respectively (rearranged as column vectors);  $\Phi$  models the sampling process (lowpass filtering and down-sampling);  $\Delta$  is the Laplacian operator, which is commonly approximated as a discrete filter (see an alternative Fourier domain implementation in Section 6.4.1):

$$\frac{1}{4} \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}.$$

Despite of the simplicity, the up-sampled images from interpolation based approaches tend to be overly smooth and suffer from annoying ringing artifacts (Figure 6.6). However, as we have shown in the previous section, we can construct a continuous domain edge model from a discrete image. We may then enforce the edge model in the up-sampling process and better preserve sharp edges.

### 6.4.1 Problem formulation

Following [95], the annihilation-driven image up-sampling considers a modified objective function than in (6.12) by utilizing the  $\ell_1$ -norm of the annihilation error with the up-sampled image as



**Figure 6.6:** Interpolation based image up-sampling (6.12) blurs sharp edges and suffers from ringing artifacts (up-sampling factor: 3).

an edge preserving constraint:

$$\begin{aligned} \min_{\mathbf{I}} \quad & \|\Delta \mathbf{I}\|_2^2 + \ell \overbrace{\|\mathbf{M} \mathbf{D} \mathbf{I}\|_1}^{\text{edge model}} \\ \text{subject to} \quad & \Phi \mathbf{I} = \mathbf{I}_0, \end{aligned} \quad (6.13)$$

where  $\mathbf{M}$  is a diagonal matrix with entries specified by the reconstructed mask function (6.11); and  $\mathbf{D}$  is a derivative operator such that  $\mathbf{D} \mathbf{I}$  corresponds to the derivative image  $\partial \mathbf{I} \stackrel{\text{def}}{=} \left( \frac{\partial}{\partial x} + j \frac{\partial}{\partial y} \right) I(x, y)$ . We may either use a discrete filter, e.g., a Sobel filter, to approximate the complex differentiation in spatial domain or we may directly implement the differentiations in Fourier domain. In particular, let  $\hat{I}(\omega_x, \omega_y)$  be the Fourier transform of the up-sampled image  $\mathbf{I}$ , then<sup>3</sup>  $\mathbf{D} \mathbf{I} = \mathcal{F}^{-1} \{ j(\omega_x + j\omega_y) \hat{I}(\omega_x, \omega_y) \}$ .

More specifically, in a finite dimensional setting,  $\hat{I}$  is made of the DFT coefficients of the up-sampled image and  $\omega_x, \omega_y$  are the corresponding discrete frequencies between  $-\pi$  and  $\pi$ . The diagonal entries of  $\mathbf{M}_{l,l}$  are the  $l$ -th pixel value of a *discrete* image obtained by evaluating the continuous domain mask function  $\mu(x, y)$  on a finite grid that is determined by the target resolution. We restrict ourselves to a finite dimensional setting in the following sections for ease of explanation.

## 6.4.2 Solve (6.13) with ADMM algorithm

Observe that the objective function in (6.13) is a composition of two functions of different natures: a smooth  $\ell_2$  norm of the Laplacian and a non-smooth  $\ell_1$  norm of the annihilation error. It fits naturally in the framework of *alternating direction method of multipliers* (ADMM), which has been shown to be an efficient optimization tool in solving various image restoration problems [1, 2, 27, 137]. Under very mild conditions, the ADMM iteration is guaranteed to converge to the optimal solution [49]. Please refer to [17, 32, 50] and references therein for comprehensive reviews of the ADMM algorithm and its connections with other similar optimization techniques, including Douglas-Rachford splitting, split-Bregman methods, and proximal splitting method,

<sup>3</sup>Similarly, we have the Fourier domain implementation of the Laplacian  $\Delta \mathbf{I}$  as  $\mathcal{F}^{-1} \{ (\omega_x^2 + \omega_y^2) \hat{I}(\omega_x, \omega_y) \}$ .

etc.

For a given constrained optimization, there are many equivalent ADMM formulations by using different strategies of variable splittings [38]. The rule of thumb is that the resultant optimization problem becomes separable with respect to both the original optimization variables and the newly introduced auxiliary variables and hence is easy to solve. An augmented Lagrangian method is then used to solve the equivalent constrained problem. In particular, we have considered an ADMM formulation for the annihilation-driven up-sampling problem (6.13) as:

$$\begin{aligned} & \min_{\mathbf{I} \in \mathbb{R}^L, \mathbf{u} \in \mathbb{C}^L, \mathbf{w} \in \mathbb{R}^L} \|\Delta \mathbf{I}\|_2^2 + \ell \|\mathbf{M}\mathbf{u}\|_1 \\ & \text{subject to} \quad \Phi \mathbf{w} = \mathbf{I}_0 \\ & \quad \quad \quad \mathbf{u} = \mathbf{D}\mathbf{I} \\ & \quad \quad \quad \mathbf{w} = \mathbf{I}. \end{aligned}$$

The associated augmented Lagrangian is

$$\begin{aligned} \mathcal{L}_{\lambda, \rho, \gamma}(\mathbf{I}, \mathbf{u}, \mathbf{w}, \tilde{\mathbf{z}}, \tilde{\mathbf{v}}, \tilde{\mathbf{r}}) &= \frac{1}{2} \|\Delta \mathbf{I}\|_2^2 + \frac{\ell}{2} \|\mathbf{M}\mathbf{u}\|_1 + \tilde{\mathbf{z}}^H (\Phi \mathbf{w} - \mathbf{I}_0) + \frac{\lambda}{2} \|\Phi \mathbf{w} - \mathbf{I}_0\|_2^2 + \tilde{\mathbf{v}}^H (\mathbf{u} - \mathbf{D}\mathbf{I}) \\ & \quad + \frac{\rho}{2} \|\mathbf{u} - \mathbf{D}\mathbf{I}\|_2^2 + \tilde{\mathbf{r}}^H (\mathbf{w} - \mathbf{I}) + \frac{\gamma}{2} \|\mathbf{w} - \mathbf{I}\|_2^2, \end{aligned} \quad (6.14)$$

where  $\lambda, \rho, \gamma$  are scalar Lagrange multipliers and  $\tilde{\mathbf{z}}, \tilde{\mathbf{v}}, \tilde{\mathbf{r}}$  are vector Lagrange multipliers. We can rewrite (6.14) concisely as (in addition to constant terms that are independent of  $\tilde{\mathbf{z}}, \tilde{\mathbf{v}}$  and  $\tilde{\mathbf{r}}$ ):

$$\mathcal{L}_{\lambda, \rho, \gamma}(\mathbf{I}, \mathbf{u}, \mathbf{z}, \mathbf{v}, \mathbf{w}) = \frac{1}{2} \|\Delta \mathbf{I}\|_2^2 + \frac{\ell}{2} \|\mathbf{M}\mathbf{u}\|_1 + \frac{\lambda}{2} \|\Phi \mathbf{w} - \mathbf{I}_0 + \mathbf{z}\|_2^2 + \frac{\rho}{2} \|\mathbf{u} - \mathbf{D}\mathbf{I} + \mathbf{v}\|_2^2 + \frac{\gamma}{2} \|\mathbf{w} - \mathbf{I} + \mathbf{r}\|_2^2 \quad (6.15)$$

via the change of variable:  $\mathbf{z} = \frac{1}{\lambda} \tilde{\mathbf{z}}, \mathbf{v} = \frac{1}{\rho} \tilde{\mathbf{v}}$  and  $\mathbf{r} = \frac{1}{\gamma} \tilde{\mathbf{r}}$ .

The ADMM approach then minimizes (6.15) iteratively by solving three sub-problems and update the Lagrange multipliers  $\mathbf{z}, \mathbf{v}$  and  $\mathbf{r}$  based on the errors made on each constraint with the current iterate:

$$\begin{aligned} \mathbf{I}^{(n+1)} &= \arg \min_{\mathbf{I} \in \mathbb{R}^L} \mathcal{L}_{\lambda, \rho, \gamma}(\mathbf{I}, \mathbf{u}^{(n)}, \mathbf{w}^{(n)}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{r}^{(n)}), \\ \mathbf{u}^{(n+1)} &= \arg \min_{\mathbf{u} \in \mathbb{C}^L} \mathcal{L}_{\lambda, \rho, \gamma}(\mathbf{I}^{(n+1)}, \mathbf{u}, \mathbf{w}^{(n)}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{w}^{(n)}), \\ \mathbf{w}^{(n+1)} &= \arg \min_{\mathbf{w} \in \mathbb{R}^L} \mathcal{L}_{\lambda, \rho, \gamma}(\mathbf{I}^{(n+1)}, \mathbf{u}^{(n+1)}, \mathbf{w}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{w}^{(n)}), \\ \mathbf{z}^{(n+1)} &= \mathbf{z}^{(n)} + \Phi \mathbf{w}^{(n+1)} - \mathbf{I}_0, \\ \mathbf{v}^{(n+1)} &= \mathbf{v}^{(n)} + \mathbf{u}^{(n+1)} - \mathbf{D}\mathbf{I}^{(n+1)}, \\ \mathbf{r}^{(n+1)} &= \mathbf{r}^{(n)} + \mathbf{w}^{(n+1)} - \mathbf{I}^{(n+1)}. \end{aligned}$$

Here the superscript  $(\cdot)^{(n)}$  denotes the  $n$ -th iteration. In general, each sub-problem involved in the ADMM iteration is easier to solve, which is why we would like to introduce the auxiliary variables in the first place.

### ***I* update**

With  $\mathbf{u}^{(n)}$ ,  $\mathbf{w}^{(n)}$ ,  $\mathbf{z}^{(n)}$ ,  $\mathbf{v}^{(n)}$ ,  $\mathbf{r}^{(n)}$  fixed, the *I*-update subproblem is a simple quadratic minimization. By setting the derivative of the Lagrangian with respect to *I* to zero, we have

$$\mathbf{I} = (\mathbf{\Delta}^H \mathbf{\Delta} + \rho \mathbf{D}^H \mathbf{D} + \gamma \mathbf{I})^{-1} (\rho \mathbf{D}^H (\mathbf{u} + \mathbf{v}) + \gamma (\mathbf{w} + \mathbf{r})), \quad (6.16)$$

where  $\mathbf{I}$  is the identity matrix. Notice that since the Laplacian  $\mathbf{\Delta}$  and the derivative  $\mathbf{D}$  are easily diagonalized with DFT, we can implement the matrix inverse in (6.16) as element-wise division in the Fourier domain. Had we opted not to split the variable *I* with the extra constraint  $\mathbf{w} = \mathbf{I}$ , the solution of the corresponding *I* update subproblem requires solving a linear system of equations that involves  $(\mathbf{\Delta}^H \mathbf{\Delta} + \rho \mathbf{D}^H \mathbf{D} + \gamma \mathbf{\Phi}^H \mathbf{\Phi})^{-1}$ , which cannot be diagonalised in Fourier domain any more (see also Section 6.4.2).

### ***u* update**

By dropping terms that are independent of  $\mathbf{u}$ , we have the *u*-subproblem as

$$\min_{\mathbf{u} \in \mathbb{C}^L} \frac{1}{2} \|\mathbf{u} - \mathbf{d}\|_2^2 + \frac{\ell}{2\rho} \|\mathbf{M}\mathbf{u}\|_1,$$

where  $\mathbf{d} = \mathbf{D}\mathbf{I} - \mathbf{v}$  for notational brevity. Since  $\mathbf{M}$  is a diagonal matrix, the *u*-update subproblem becomes separable with respect to each component of the *complex-valued* variable  $\mathbf{u}$ . From the optimality condition, we have

$$u_l - d_l + \frac{\ell}{2\rho} |\mu_l| \cdot \frac{u_l}{\|u_l\|_2} = 0, \quad (6.17)$$

with  $\mu_l = \mathbf{M}_{l,l}$  and  $u_l$ ,  $d_l$  being the *l*-th components of  $\mathbf{u}$  and  $\mathbf{d}$ , respectively. Here  $\|\cdot\|_2$  is the 2-norm of a complex variable. The solution of (6.17) is given by the shrinkage formula [137]:

$$u_l = \max \left\{ \|d_l\|_2 - \frac{\ell}{2\rho} |\mu_l|, 0 \right\} \frac{d_l}{\|d_l\|_2} \text{ for } l = 1, \dots, L.$$

Here we have assumed that  $0 \cdot \frac{0}{0} = 0$ . In the actual implementation, we may simply modify the denominator as  $\max\{\|d_l\|_2, \varepsilon\}$  with a small  $\varepsilon$ , e.g.,  $10^{-12}$  for numerical stability.

### ***w* update**


Similar to the *I*-subproblem, we need to solve a simple quadratic minimization:

$$\min_{\mathbf{w}} \frac{\lambda}{2} \|\mathbf{\Phi}\mathbf{w} - \mathbf{I}_0 + \mathbf{z}\|_2^2 + \frac{\gamma}{2} \|\mathbf{w} - \mathbf{I} + \mathbf{r}\|_2^2,$$


whose closed-form solution is

$$\mathbf{w} = (\lambda \mathbf{\Phi}^H \mathbf{\Phi} + \gamma \mathbf{I})^{-1} (\lambda \mathbf{\Phi}^H (\mathbf{I}_0 - \mathbf{z}) + \gamma (\mathbf{I} - \mathbf{r})). \quad (6.18)$$

The challenge here is how to implement (6.18) efficiently: in the case of image up-sampling,  $\mathbf{\Phi}$  models both the anti-aliasing blurring kernel and down-sampling. Consequently,  $\mathbf{\Phi}^H \mathbf{\Phi}$  can no



**Figure 6.7:** Comparisons of edge detection results on a test image in (a) with (b) Canny edge detector and (c) the annihilation-driven approaches with isotropic windows (white diamond) and anisotropic ones (yellow dots).



**Figure 6.8:** Stack of light-field images and a slice on the  $x$ - $\alpha$  plane in the epipolar domain.


longer be diagonalized in the Fourier domain. Thanks to the matrix inversion lemma, we have:

$$(\lambda \Phi^H \Phi + \gamma \mathbf{I})^{-1} = \frac{1}{\gamma} \mathbf{I} - \frac{\lambda}{\gamma} \Phi^H (\gamma \mathbf{I} + \lambda \Phi \Phi^H)^{-1} \Phi. \quad (6.19)$$

On the other hand, it can be shown that  $\Phi \Phi^H$  is simply another convolution matrix associated with a certain filter, which again allows us to have an efficient implementation of (6.18) in the Fourier domain (see details in Appendix 6.B).

## 6.5 Experimental results

In this section we exemplify the proposed edge modeling algorithm (Section 6.5.1) and evaluate the performance of the edge-friendly single image up-sampling by comparing with state of the art (Section 6.5.2). We have run all experiments with Matlab 8.5 in OS X 10.10.3 on a MacBook Pro laptop with 2.5GHz Intel Core i7 CPU and 16GB RAM. Computational time is averaged over ten runs.



**Figure 6.9:** Reconstructed line models (yellow lines) for a light-field image in the epipolar domain with (a) uniform-sized Gaussian windows; and (b) windows that are adapted to the local geometries (image size:  $17 \times 307$ ).

### 6.5.1 Edge modelling

We apply the sub-pixel accurate edge detection algorithms of Section 6.3.1 to a natural image using isotropic and anisotropic Gaussian windows, respectively. The block size is  $21 \times 21$  pixels and standard deviation  $\sigma = 2$  (with respect to a pixel size) for cases with isotropic Gaussian windows. We may further refine the estimated edge models by using anisotropic windows, where the window shapes are adapted based on the estimated edge models (6.9). In particular, the standard deviations along the horizontal and vertical directions are rescaled as  $\sigma_x = \sigma(\lambda_{\max}/\lambda_{\min})^{1/4}$  and  $\sigma_y = \sigma(\lambda_{\min}\lambda_{\max})^{1/4}$ , where  $\lambda_{\min}$  and  $\lambda_{\max}$  are eigenvalues computed from (6.7). The reconstructed edge positions (6.8) are compared with the classic Canny edge detector in Figure 6.7 for a test image *pepper* (size:  $512 \times 512$ ). Thanks to the efficient implementation detailed in Appendix 6.A, it takes a mere 0.11 sec to reconstruct the edge models for the whole image when isotropic Gaussian windows are used.

Another potential application of the annihilation-driven edge detection algorithm is on light-field images with the epipolar-image (EPI) representation [16]. In the EPI domain, the horizontal axis corresponds to the pixel coordinates on the camera plane; while the vertical coordinates are associated with different camera view angles (Figure 6.8). One major advantage of the EPI representation is that the depth information of a 3D scene is naturally encoded as slopes of the line structures. Combined with matching algorithms, we can obtain more accurate estimations of camera translations, which are no longer limited by pixel-level accuracies. Considering the specific line structures, we have used a narrow Gaussian window with width of 9 pixels and spanning all rows of the epipolar image. The standard deviations along horizontal and vertical directions are  $\sigma_x = 2$  and  $\sigma_y = 5$ , respectively. The line models are then obtained by sliding Gaussian windows horizontally around edge points specified by the Canny edge detector. Note that in the algorithm we have assumed that there is only one line within each window. It may not be satisfied in all cases when a uniform-sized Gaussian window is applied to the whole image. By adapting window shapes (rescaling and rotating as in (6.9)) based on the estimated edge models, more reliable reconstructions are obtained (Figure 6.9).

### 6.5.2 Single image up-sampling

In this part, we evaluate the performance of the proposed image up-sampling algorithm applied to several test images in Figure 6.10. The ground truth high resolution image is low pass filtered with a bicubic kernel and down-sampled by a factor of 3. If the original image size is not a





**Figure 6.10:** Ground truth high resolution images considered in the single image up-sampling experiments. From left to right: *butterfly* ( $256 \times 256$ ), *chip* ( $200 \times 244$ ), *flowers* ( $362 \times 500$ ), *house* ( $256 \times 256$ ), *mandrill* ( $512 \times 512$ ), *mit* ( $256 \times 256$ ), *peppers* ( $512 \times 512$ ), *ppt* ( $656 \times 529$ ), *temple* ( $370 \times 370$ ), and *zebra* ( $391 \times 586$ ).

multiple of 3, we crop the image by removing the last few rows and/or columns so that the total number of rows and columns are integer multiples of the down-sampling ratio. The image up-sampling results are evaluated with both peak signal to noise ratio (PSNR) and structural similarity index (SSIM) between the up-sampled and ground truth image. For color images, we first transform the RGB channels to YCbCr color space. The up-sampling algorithm is applied to the illumination channel alone, on which the performance metrics are subsequently evaluated.

### Reconstruction of the global edge model

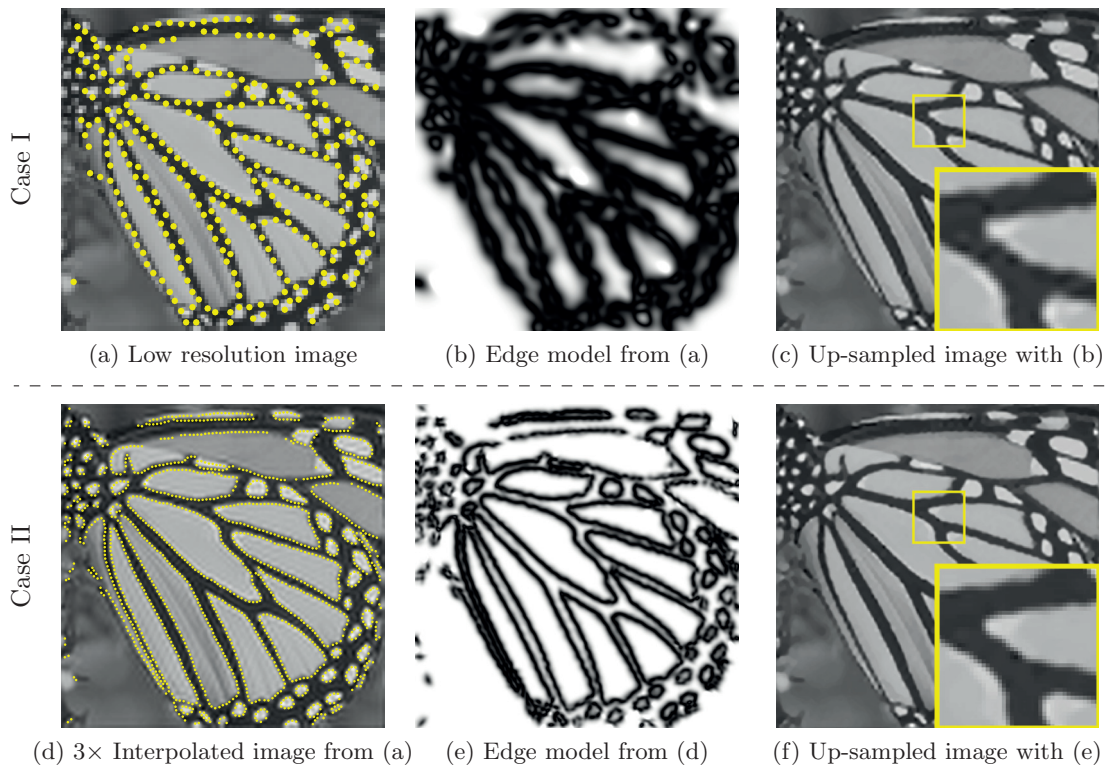
In experiments, we first reconstruct a global edge model (6.11) from the given low resolution image. We have centred edge blocks  $\{w_n\}_{n=1}^N$ , which are tailored to the local edge geometry (6.9), at sub-pixel accurate positions  $(x_n, y_n)$  given by the block-wise estimation. In the smooth area, isotropic Gaussian windows with  $\sigma = 2$  are placed on a uniform grid with intersections that are  $\sigma$ -apart on both vertical and horizontal directions. Empirically, we have observed that a more accurate edge model is obtained when we apply the algorithm to an image that is interpolated from the low resolution one, e.g., the solution of (6.12). One explanation might be that more edge control points, i.e.,  $N$  in (6.11), are allowed when we estimate the global edge model from the interpolated image. Take the test image *butterfly* as an example: the total number of edge blocks is 316 and 1547 for a global edge model estimated from the low resolution image directly or from the interpolated image, respectively. With increased degrees of freedom, the reconstructed model is better localized around edges, which facilitates an improved up-sampling result (Figure 6.11). The reconstructed edge model is incorporated to regularize the up-sampling process (6.13) with a certain regularization weight. Experimentally, we found that  $\ell = 15$  is a good compromise between maintaining the overall smoothness and preserving sharp edges.

### Stopping criterion

The up-sampled images are reconstructed by running the ADMM algorithm in Section 6.4.2 with a *fixed* number of 2000 iterations for all cases. The reason that we have used this criterion is that the solution obtained at the 2000-th iteration is sufficiently close to the optimal solution of (6.13) in practice. In Figure 6.12, we plot the evolution of the distance (which is a normalized MSE) between the intermediate solution at the  $n$ -th iteration  $\mathbf{I}^{(n)}$  and the optimal solution<sup>4</sup>  $\mathbf{I}^*$

$$\text{dist}(\mathbf{I}^{(n)}, \mathbf{I}^*) \stackrel{\text{def}}{=} 10 \log_{10} \left( \frac{255^2}{\text{MSE}(\mathbf{I}^{(n)}, \mathbf{I}^*)} \right) \quad (6.20)$$

<sup>4</sup>We ran the ADMM algorithm with  $10^6$  iterations and verified that the optimality conditions are satisfied within computational accuracy.



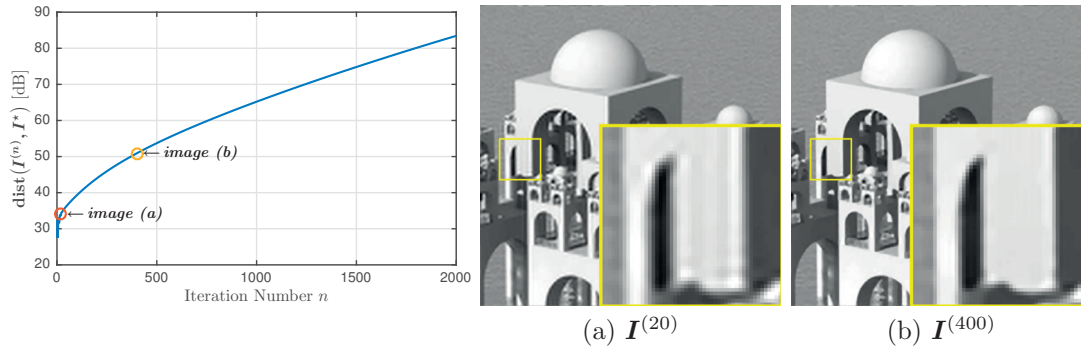
**Figure 6.11:** It is better to work on an interpolated image (*Case II*) than the given raw low resolution image directly (*Case I*): PSNRs are 26.47 dB and 28.09 dB for the up-sampled images in (c) and (f), respectively. Number of edge points (yellow dots): 316 in (a) and 1547 in (d); Continuously defined edge model plot size:  $1000 \times 1000$ .

for the test image *temple*. There is little difference between the up-sampled images obtained at 400-th iteration and 2000-th iteration visually (at a fifth of the computational cost), albeit a slight improvement in PSNR (less than 0.2 dB). For applications where the accuracy is less of a concern, we could potentially reduce the iteration numbers and hence the computation time significantly. Nevertheless, in our current Matlab implementation with the specified experimental setups (sampling kernel, maximum 2000 iterations, etc.), it takes 24.17 seconds to up-sample an  $85 \times 85$  image to  $255 \times 255$ . In terms of computational complexity, the iterative algorithm is  $\mathcal{O}(L \log L)$  because of the Fourier domain implementation, where  $L$  is the total number of pixels in the up-sampled image.

### Comparisons with state of the art

Recent developments in computer vision gave rise to several promising single image up-sampling algorithms by exploiting sparse representations for both low and high resolution images in the patch space [64, 73, 81, 100, 136, 144, 147]. We compare our proposed approach<sup>5</sup> with several state of the art algorithms and summarize the reconstruction performance (PSNR and SSIM) in

<sup>5</sup>The Matlab implementation will be made available online: <http://1cav.epfl.ch/people/hanjie.pan>



**Figure 6.12:** Evolution of the distance (6.20) between the reconstructed image at the  $n$ -th ADMM iteration  $I^{(n)}$  and the one at convergence  $I^*$  (test image: *temple*); The up-sampled images at (a) 20-th iteration and (b) 400-th iteration.

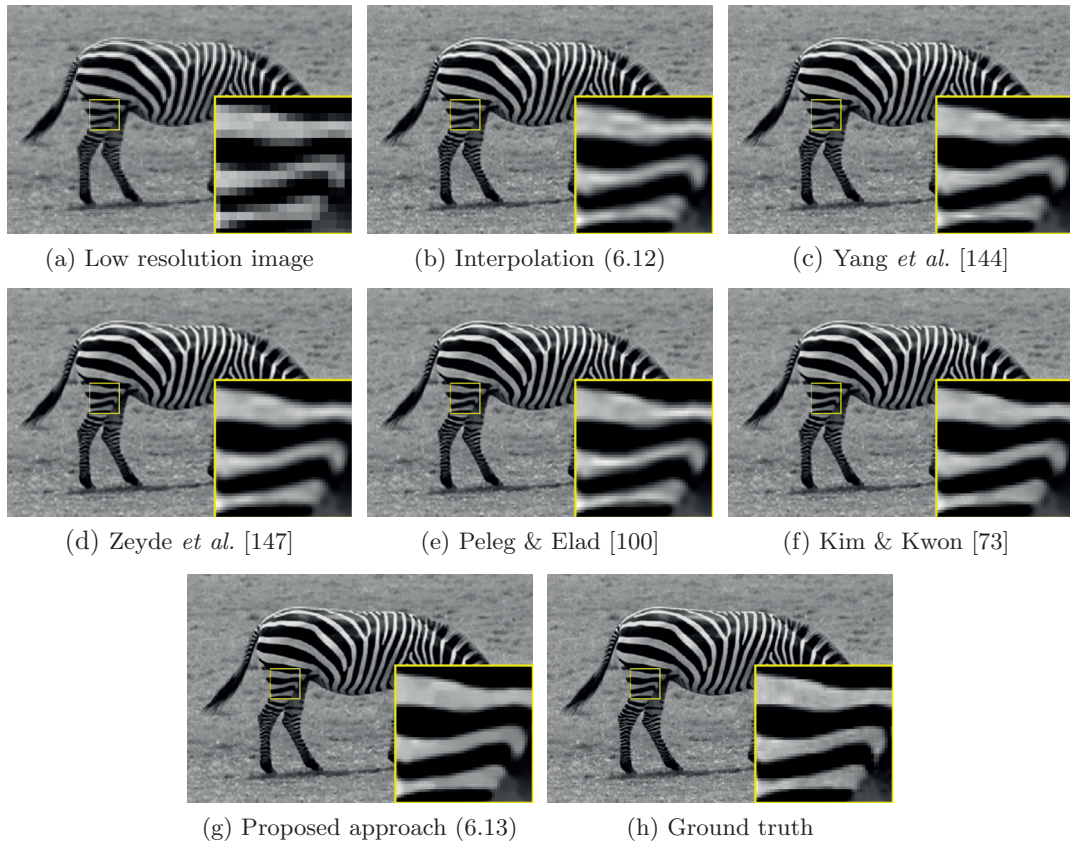
**Table 6.1:** Comparisons of Single Image Up-Sampling (bicubic kernel, up-sampling factor: 3)

Image	<i>butterfly</i>	<i>chip</i>	<i>flowers</i>	<i>house</i>	<i>mandrill</i>	<i>mit</i>	<i>peppers</i>	<i>ppt</i>	<i>temple</i>	<i>zebra</i>
Bicubic Interpolation	<u>24.04</u> 0.8220	<u>28.69</u> 0.9164	<u>27.23</u> 0.8017	<u>29.80</u> 0.8363	<u>22.43</u> 0.7960	<u>21.40</u> 0.7653	<u>32.41</u> 0.9688	<u>23.70</u> 0.9416	<u>25.26</u> 0.8130	<u>26.64</u> 0.9117
Thin-plate Spline Interpolation (6.12)	<u>25.08</u> 0.8258	<u>29.86</u> 0.9199	<u>28.18</u> 0.8312	<u>30.81</u> 0.8448	<u>22.77</u> 0.8440	<u>22.49</u> 0.7864	<u>33.26</u> 0.9744	<u>24.69</u> 0.9527	<u>26.29</u> 0.8263	<u>28.06</u> 0.9423
Yang <i>et al.</i> [144]	<u>26.30</u> 0.8853	<u>30.55</u> 0.9403	<u>28.52</u> 0.8424	<u>31.33</u> 0.8595	<u>22.80</u> 0.8461	<u>22.61</u> 0.8187	<u>33.86</u> 0.9782	<u>25.02</u> 0.9648	<u>26.62</u> 0.8508	<u>28.45</u> 0.9433
Zeyde <i>et al.</i> [147]	<u>25.98</u> 0.8775	<u>30.60</u> 0.9435	<u>28.43</u> 0.8384	<u>31.56</u> 0.8619	<u>22.77</u> 0.8387	<u>22.87</u> 0.8196	<u>34.09</u> 0.9782	<u>25.18</u> 0.9647	<u>26.72</u> 0.8523	<u>28.51</u> 0.9408
Peleg & Elad [100]	<u>26.75</u> 0.8999	<u>30.54</u> 0.9427	<u>28.82</u> 0.8474	<u>31.67</u> 0.8643	<b><u>22.87</u></b> <b>0.8480</b>	<u>23.19</u> 0.8337	<u>34.38</u> 0.9791	<u>25.72</u> 0.9708	<u>27.21</u> 0.8603	<u>28.67</u> 0.9432
Kim & Kwon [73]	<u>27.34</u> 0.9081	<u>31.14</u> 0.9529	<u>28.82</u> 0.8472	<u>31.87</u> 0.8619	<u>22.83</u> <b>0.8482</b>	<u>23.14</u> 0.8361	<u>34.38</u> 0.9781	<u>25.58</u> 0.9682	<u>27.13</u> 0.8586	<u>28.87</u> 0.9431
Proposed	<b><u>28.09</u></b> <b>0.9201</b>	<b><u>31.86</u></b> <b>0.9585</b>	<b><u>29.06</u></b> <b>0.8533</b>	<b><u>32.11</u></b> <b>0.8688</b>	<u>22.67</u> 0.8422	<b><u>23.54</u></b> <b>0.8478</b>	<b><u>34.47</u></b> <b>0.9793</b>	<b><u>26.49</u></b> <b>0.9770</b>	<b><u>27.85</u></b> <b>0.8766</b>	<b><u>29.19</u></b> <b>0.9447</b>

Performance metrics in each cell: PSNR [dB] (upper) and SSIM (lower).

Table 6.1. We have kept the default parameter settings in the distributed softwares and used the provided dictionaries for [100, 147]. Since the provided dictionary in [144] is for up-sampling by a factor of 2, we re-trained the dictionary based on the provided training data set included in the software. In most cases, the annihilation-driven image up-sampling algorithm outperforms Yang *et al.* [144], Zeyde *et al.* [147], Peleg & Elad [100] and Kim & Kwon [73] and in many cases the improvements are quite substantial. An example with the test image *zebra* is included for visual comparisons in Figure 6.13. The annihilation-driven approach not only gives sharper edges but also reduces ringing artifacts significantly in the smooth area.

Notice that the proposed algorithm does not work well on *mandrill*, which consists of a large portion of high frequency texture area. Such behavior, however, is not unexpected. A reminder that what differs between (6.13) and the standard interpolation result (6.12) is essentially the edge-friendly regularization. The algorithm has difficulty in distinguishing between edges and textures, both of which contribute to high frequency components in an image. Consequently, sharp “edges” are encouraged both around actual edges *and* in the texture areas leading to undesirable artifacts (Figure 6.14). To confirm our conjecture, we evaluate the local

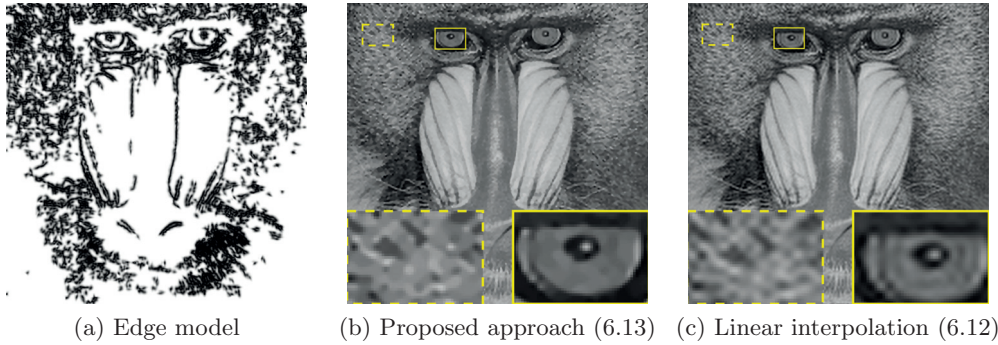


**Figure 6.13:** Visual comparisons of various up-sampling algorithms (sampling kernel: bicubic; up-sampling ratio: 3). (a) The given low resolution image (size:  $130 \times 195$ ); (b)–(g) Up-sampled image with thin-plate spline interpolation (6.12), Yang *et al.* [144], Zeyde *et al.* [147], Peleg & Elad [100], Kim & Kwon [73], and the proposed approach (6.13), respectively; (h) The ground truth high resolution image. See the last column of Table 6.1 for PSNRs.

PSNRs within two rectangles placed around edges (solid line) and in the texture area (dashed line):  $\text{PSNR}_{\text{edge}} = 25.51 \text{ dB}$  and  $\text{PSNR}_{\text{texture}} = 19.22 \text{ dB}$  with the proposed method; while  $\text{PSNR}_{\text{edge}} = 24.17 \text{ dB}$  and  $\text{PSNR}_{\text{texture}} = 19.37 \text{ dB}$  with thin plate spline interpolation (6.12). We may potentially improve the up-sampling results further with a better edge model by excluding edge windows (i.e.,  $w_n$  in (6.11)) located in the texture areas via segmentations, e.g., based on the level of consistency of edge orientations in a local area.

### Example with actual image up-samplings

In addition to the simulated experiments in the previous section, we also considered more practical cases where we do not have access to the ground truth of the high resolution image. With the proposed approach, we need to estimate the data acquisition model that generates the low resolution image. For simplicity, here we may assume that the low resolution image is obtained by low pass filtering a high resolution image with certain anti-aliasing filter, e.g., a bicubic kernel,



**Figure 6.14:** Annihilation-driven image up-sampling algorithm encourages sharp edges around actual edges as well as textures, which produces artificial edges in texture areas. (a) Estimated global edge model (plot size:  $1000 \times 1000$ ); (b) Up-sampled image with the proposed method (6.13). Local PSNRs within an edge block (solid line) and a texture block (dashed line):  $\text{PSNR}_{\text{edge}} = 25.51$  dB and  $\text{PSNR}_{\text{texture}} = 19.22$  dB; (c) Up-sampled image with thin plate spline interpolation (6.12). Local PSNRs:  $\text{PSNR}_{\text{edge}} = 24.17$  dB and  $\text{PSNR}_{\text{texture}} = 19.37$  dB.

and down sampling to the target resolution. Then, we can apply a similar approach to up-sample the given image:

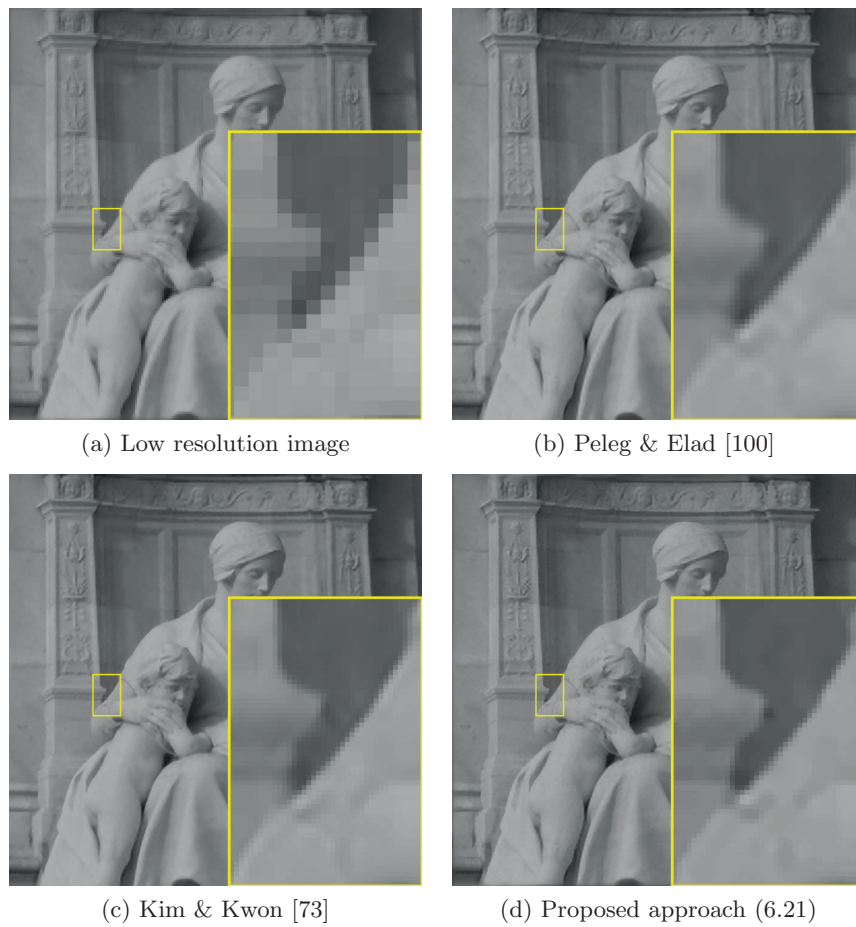
$$\begin{aligned} \min_{\mathbf{I}} \quad & \|\Delta \mathbf{I}\|_2^2 + \ell \|\mathbf{MDI}\|_1 \\ \text{subject to} \quad & \|\Phi \mathbf{I} - \mathbf{I}_0\|_2^2 \leq \varepsilon^2. \end{aligned} \quad (6.21)$$

Here we no longer enforce the exact data-fidelity constraint due to errors in the estimated data acquisition process  $\Phi$  and/or noise in the image. Instead, we allow certain level of discrepancies (controlled by  $\varepsilon > 0$ ) in the up-sampling process. (6.21) can be solved efficiently with a slight modification to the ADMM algorithm presented in Section 6.4.2 (see details in Appendix 6.C). Whether this assumption is a good approximation or not requires further investigation<sup>6</sup>. Since we no longer have access to the ground truth, we summarize the up-sampled images in Figure 6.15 for visual comparisons. Here we have chosen  $\varepsilon$  such that it corresponds to a 35 dB SNR in  $\mathbf{I}_0$ . Compared with [100] and [73], the up-sampled image using our approach retains sharp edges and is visually appealing.

## 6.6 Conclusion


We have presented a novel image edge modeling framework based on a local edge model following the principle of curve annihilation: instead of directly addressing the edge modeling problem by matching a function with image gradients, we reconstruct a mask (which is our reconstructed edge model) such that its product with the gradients is small. We show that accurate edge models are estimated by assuming a simple local line model. The same idea is generalized to reconstruct a global edge model for the whole image. We demonstrate the effectiveness of this model in preserving sharp edges for image up-sampling, and the proposed algorithm achieves state of the art performance. Future work shall focus on more practical cases where we have no

<sup>6</sup>One possibility is to utilize the *blur-SURE* criterion [143], which is an unbiased estimate of the blur MSE, for the anti-aliasing kernel estimation.



**Figure 6.15:** Visual comparisons for an actual image up-sampling. (a) The low resolution image (size:  $200 \times 200$ ); Up sampled images (size:  $600 \times 600$ ) with (b) Peleg & Elad [100], (c) Kim & Kwon [73] and (d) the proposed approach (6.21), respectively.

direct access to the data-acquisition process that generates the low resolution image. It would also be interesting to investigate how a multi-stage up-sampling scheme (similar to [60]) may affect the performance.

Figure 6.16: Block diagram of the operator  $\Phi\Phi^H$ .

## 6.A Efficient implementation of block-wise edge modeling

If we expand the quadratic objective function in (6.5), we have the entries of the Hermitian symmetric matrix  $\mathbf{A}$  as:

$$\begin{aligned} A_{1,1} &= \iint (x - x_0)^2 |\partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy, \\ A_{2,2} &= \iint (y - y_0)^2 |\partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy, \\ A_{3,3} &= \iint |\partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy, \\ A_{1,2} = A_{2,1} &= \iint (x - x_0)(y - y_0) |\partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy, \\ A_{1,3} = A_{3,1} &= \iint (x - x_0) |\partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy, \\ A_{2,3} = A_{3,2} &= \iint (y - y_0) |\partial I(x, y)|^2 w(x - x_0, y - y_0) dx dy. \end{aligned}$$

We can treat these integrations as convolutions between the square of the derivative image  $|\partial I(x, y)|^2$  and certain filters evaluated at  $(x_0, y_0)$ :

$$\begin{aligned} A_{1,1} &= (x^2 w') * |\partial I|^2 \Big|_{\substack{x=x_0 \\ y=y_0}}, & A_{2,2} &= (y^2 w') * |\partial I|^2 \Big|_{\substack{x=x_0 \\ y=y_0}}, \\ A_{3,3} &= w' * |\partial I|^2 \Big|_{\substack{x=x_0 \\ y=y_0}}, & A_{1,2} = A_{2,1} &= (xy w') * |\partial I|^2 \Big|_{\substack{x=x_0 \\ y=y_0}}, \\ A_{1,3} = A_{3,1} &= -(x w') * |\partial I|^2 \Big|_{\substack{x=x_0 \\ y=y_0}}, & A_{2,3} = A_{3,2} &= -(y w') * |\partial I|^2 \Big|_{\substack{x=x_0 \\ y=y_0}}, \end{aligned}$$

where  $w'(x, y) = w(-x, -y)$ . Hence, instead of computing matrix  $\mathbf{A}$  for each block, we can simply apply 6 filtering to the whole image. These filtered images then serve as “look-up tables”, based on which the edge models (6.6) are reconstructed simultaneously for all the blocks.

## 6.B Implementation of (6.18) in the DFT domain

The key to have an efficient implementation of the  $\mathbf{w}$  update lies on finding the matrix inverse  $(\gamma \mathbf{I} + \lambda \Phi \Phi^H)^{-1}$ . Let  $\varphi_l$  be the antialiasing filter used in the sampling process  $\Phi$  and  $s$  the down-sampling ratio (Figure 6.16). Denote  $\psi_l = \varphi_{-l}^* * \varphi_l$ , then the output  $X_2(z)$  is related with the input  $X_1(z)$  as

$$X_2(z) = \left( X_1(z^s) \hat{\psi}(z) \right) \downarrow s = \frac{1}{s} \sum_{i=0}^{s-1} X_1 \left( (W_s^i z^{1/s})^s \right) \hat{\psi} \left( W_s^i z^{1/s} \right) = X_1(z) \frac{1}{s} \sum_{i=0}^{s-1} \hat{\psi} \left( W_s^i z^{1/s} \right),$$

where  $\hat{\psi}$  is the  $z$ -transform of  $\psi$  and  $W_s = e^{-j\frac{2\pi}{s}}$ . Therefore, we may treat the operator  $\Phi\Phi^H$  as applying a filter, which is the down-sampled version of  $\psi_l = \varphi_{-l}^* * \varphi_l$  by a factor  $s$ . Combined with the matrix inversion lemma, we can implement the  $\mathbf{w}$  updates (6.18) efficiently in the DFT domain.

## 6.C Solve the noisy up-sampling (6.21) with ADMM

We address the inequality constraint in (6.21) by introducing a set indicator function [2, 17]  $g(\mathbf{p})$ :

$$g(\mathbf{p}) = \begin{cases} 0 & \text{if } \|\mathbf{p}\|_2^2 \leq \varepsilon^2, \\ +\infty & \text{otherwise.} \end{cases}$$

An equivalent ADMM formulation for (6.21) is:

$$\begin{aligned} \min_{\mathbf{I}, \mathbf{u}, \mathbf{w}, \mathbf{p}} \quad & \|\Delta\mathbf{I}\|_2^2 + \ell\|\mathbf{M}\mathbf{u}\|_1 + g(\mathbf{p}) \\ \text{subject to} \quad & \Phi\mathbf{w} - \mathbf{I}_0 = \mathbf{p} \\ & \mathbf{u} = \mathbf{D}\mathbf{I} \\ & \mathbf{w} = \mathbf{I}. \end{aligned}$$

The associated augmented Lagrangian function of the ADMM constrained formulation is:

$$\begin{aligned} \mathcal{L}_{\lambda, \rho, \gamma}(\mathbf{I}, \mathbf{u}, \mathbf{w}, \mathbf{p}, \mathbf{z}, \mathbf{v}, \mathbf{r}) = & \frac{1}{2}\|\Delta\mathbf{I}\|_2^2 + \frac{\ell}{2}\|\mathbf{M}\mathbf{u}\|_1 + g(\mathbf{p}) + \frac{\lambda}{2}\|\Phi\mathbf{w} - \mathbf{I}_0 - \mathbf{p} + \mathbf{z}\|_2^2 \\ & + \frac{\rho}{2}\|\mathbf{u} - \mathbf{D}\mathbf{I} + \mathbf{v}\|_2^2 + \frac{\gamma}{2}\|\mathbf{w} - \mathbf{I} + \mathbf{r}\|_2^2. \end{aligned}$$

We need to solve four subproblems and update the vector Lagrange multipliers  $\mathbf{z}$ ,  $\mathbf{v}$ , and  $\mathbf{r}$  based on the reconstruction at the current iterate:

$$\begin{aligned} \mathbf{I}^{(n+1)} &= \arg \min_{\mathbf{I}} \mathcal{L}(\mathbf{I}, \mathbf{u}^{(n)}, \mathbf{w}^{(n)}, \mathbf{p}^{(n)}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{r}^{(n)}), \\ \mathbf{u}^{(n+1)} &= \arg \min_{\mathbf{u}} \mathcal{L}(\mathbf{I}^{(n+1)}, \mathbf{u}, \mathbf{w}^{(n)}, \mathbf{p}^{(n)}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{r}^{(n)}), \\ \mathbf{w}^{(n+1)} &= \arg \min_{\mathbf{w}} \mathcal{L}(\mathbf{I}^{(n+1)}, \mathbf{u}^{(n+1)}, \mathbf{w}, \mathbf{p}^{(n)}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{r}^{(n)}), \\ \mathbf{p}^{(n+1)} &= \arg \min_{\mathbf{p}} \mathcal{L}(\mathbf{I}^{(n+1)}, \mathbf{u}^{(n+1)}, \mathbf{w}^{(n+1)}, \mathbf{p}, \mathbf{z}^{(n)}, \mathbf{v}^{(n)}, \mathbf{r}^{(n)}), \\ \mathbf{z}^{(n+1)} &= \mathbf{z}^{(n)} + \Phi\mathbf{w}^{(n+1)} - \mathbf{I}_0 - \mathbf{p}^{(n+1)}, \\ \mathbf{v}^{(n+1)} &= \mathbf{v}^{(n)} + \mathbf{u}^{(n+1)} - \mathbf{D}\mathbf{I}^{(n+1)}, \\ \mathbf{r}^{(n+1)} &= \mathbf{r}^{(n)} + \mathbf{w}^{(n+1)} - \mathbf{I}^{(n+1)}. \end{aligned}$$

In fact, we have exactly the same  $\mathbf{I}$ - and  $\mathbf{u}$ -subproblems as in Section 6.4.2. While  $\mathbf{w}$  is updated by solving a similar quadratic minimization, whose closed form solution is

$$\mathbf{w}^{(n+1)} = (\lambda\Phi^H\Phi + \gamma\mathbf{I})^{-1}(\lambda\Phi^H(\mathbf{I}_0 + \mathbf{p} - \mathbf{z}) + \gamma(\mathbf{I} - \mathbf{r})).$$



It can be implemented efficiently with the aid of matrix inversion lemma (6.19) and the trick in Appendix 6.B.

Now we focus on the  $\mathbf{p}$ -subproblem:

$$\min_{\mathbf{p}} \frac{\lambda}{2} \left\| \mathbf{p} - \left( \Phi \mathbf{w}^{(n+1)} + \mathbf{z}^{(n)} - \mathbf{I}_0 \right) \right\|_2^2 + g(\mathbf{p}),$$

whose solution is a projection of  $\Phi \mathbf{w}^{(n+1)} + \mathbf{z}^{(n)} - \mathbf{I}_0$  onto the convex set  $\mathcal{C} = \{\mathbf{p} \mid \|\mathbf{p}\|_2^2 \leq \varepsilon^2\}$ , i.e.,

$$\mathbf{p}^{(n+1)} = \begin{cases} \Phi \mathbf{w}^{(n+1)} + \mathbf{z}^{(n)} - \mathbf{I}_0 & \text{if } \|\Phi \mathbf{w}^{(n+1)} + \mathbf{z}^{(n)} - \mathbf{I}_0\|_2^2 \leq \varepsilon^2, \\ \varepsilon \frac{\Phi \mathbf{w}^{(n+1)} + \mathbf{z}^{(n)} - \mathbf{I}_0}{\|\Phi \mathbf{w}^{(n+1)} + \mathbf{z}^{(n)} - \mathbf{I}_0\|_2} & \text{otherwise.} \end{cases}$$



## Chapter 7

# Conclusion

众里寻他千百度，蓦然回首，那人却在，灯火阑珊处。

In the crowd once and again, I look for her in vain. When all at once I turn my head, I find her there where lantern light is dimly shed.

---

青玉案·元宵 *Qing Yu'an · Lantern Festival*  
辛弃疾 XIN, QIJI

The thesis investigated the finite rate of innovation sampling framework on both the theoretical and algorithmic fronts. We generalized the FRI technique to cases with non-uniform sampling. The generalization makes FRI applicable to many practical sparse recovery problems. Further, we proposed an efficient sampling framework with a linear sample complexity for multi-dimensional Diracs. The new multi-dimensional sampling framework is a significant improvement over previous methods, which have a sample complexity that increases exponentially with dimension. We validated the applicability of the FRI-based continuous sparse recovery technique in three practical problems, namely (i) radio astronomy point source estimation, (ii) direction of arrival estimation in acoustics, and (iii) single image up-sampling.

Reflecting back on the four-year journey, the path we took is much clearer now: For a given sparse reconstruction problem, as soon as (i) we identify a set of uniform samples of sinusoids (which are related to the unknown sparse signal parameters); and (ii) establish the connection from these sinusoidal samples to the measurements, then the FRI-based sparse recovery technique developed in this thesis can be applied to solve the problem. In a sense, we have invented a hammer and all at once the world appears to be full of nails. Here, we outline a few interesting problems, where the FRI technique may provide an alternative and potentially significant improvement over existing methods.

**Inverse gravimetric problem** The goal of inverse gravimetric problem is to infer the mass distributions within the Earth from measurements of the gravitational potentials outside the Earth. For instance, the Gravity Recovery and Climate Experiment (GRACE) is an on-going NASA project that measures the gravitational field with a pair of satellites orbiting the Earth. One frequently used model for the mass distribution is the point mass model. Then the original problem boils down to the estimation of the point source locations and the associated masses. This turns out to be very similar to the 3D DOA estimation problem that we have considered in Chapter 5 with the distinction that sources are in a “near-field” setting: we would also like to infer the depths the the point masses. With a slight modification, we can apply the same FRI-based sparse recovery technique to the inverse gravimetric problem.

**Single molecule localization microscopy** The resolution power of a microscope is limited by the light wavelength, typically around a few hundred nanometers. Single-molecule localization microscopy (SMLM) [108] is an imaging technique that achieves a resolution down to around 20 nm. In SMLM, the samples to be imaged are covered by fluorescent molecules. At each time instance, a subset of the fluorescences are excited. By estimating the few (often isolated) fluorescent molecules over different frames, SMLM achieves super-resolution. The FRI-based sparse recovery technique is directly applicable here: What the FRI technique can bring is the possibility to have more fluorescent molecules excited within each frame (and hence fewer frames to image the whole sample). The challenge facing existing methods in these scenarios is that they often have difficulty in estimating the molecule locations reliably when images of adjacent fluorescences overlap, a problem that FRI is likely to solve.

**Single-photon emission computed tomography** Single-photon emission computed tomography (SPECT) is an imaging technique of biological activities, e.g., metabolism, by detecting gamma radiation of radioactive tracer. SPECT is similar to a better known technique, Positron Emission Tomography (PET). The challenge with SPECT is that typically the scans give lower resolution compared with PET and usually take a long time for each acquisition. A robust technique, like FRI, is therefore potentially very useful in SPECT to provide better localization of biological activities and reduce the number of measurements needed.

**Crystallography** Crystallography, or more specifically X-ray crystallography [89], aims at determining the structure of molecules (i.e., the locations of atoms) from measured diffraction patterns. The molecule is first crystalized, which in essence forms a periodic stream of Diracs in 3D (where the electron density of each atom is modeled as a Dirac). The crystal is then lit by X-ray, which has a comparable wavelength to a typical atom size. It is not surprising that the scattering patterns of the crystal are related to the Fourier series (FS) coefficients of the Dirac stream. The challenge here is that the phases of the complex-valued FS coefficients are lost — only the amplitudes are measured from the scattering pattern. In terms of FRI-based recovery, this amounts to solving a “blind” estimation problem, where the transformation from the (complex-valued) FS coefficients to the measured scattering pattern (i.e., the moduli of them) is only partially known.

# Bibliography

- [1] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, “Fast image recovery using variable splitting and constrained optimization,” *IEEE Transactions on Image Processing*, vol. 19, no. 9, pp. 2345–2356, 2010.
- [2] —, “An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems,” *IEEE Transactions on Image Processing*, vol. 20, no. 3, pp. 681–695, 2011.
- [3] A. Aldroubi and K. Gröchenig, “Nonuniform sampling and reconstruction in shift-invariant spaces,” *SIAM review*, vol. 43, no. 4, pp. 585–620, 2001.
- [4] H. Aly and E. Dubois, “Image up-sampling using total-variation regularization with a new observation model,” *IEEE Transactions on Image Processing*, vol. 14, no. 10, pp. 1647–1659, 2005.
- [5] F. Andersson, M. Carlsson, J.-Y. Tournier, and H. Wendt, “A new frequency estimation method for equally and unequally spaced data,” *IEEE Transactions on Signal Processing*, vol. 62, no. 21, pp. 5761–5774, 2014.
- [6] C. Aubel, D. Stotz, and H. Bölcskei, “A theory of super-resolution from short-time fourier transform measurements,” *Journal of Fourier Analysis and Applications*, 2017, to appear.
- [7] J. Azcarreta Ortiz, “Pyramic array: An FPGA based platform for many-channel audio acquisition,” Master’s thesis, EPFL, Lausanne, Switzerland, Aug. 2016.
- [8] L. Baboulaz and P. L. Dragotti, “Exact feature extraction using finite rate of innovation principles with an application to image super-resolution,” *IEEE Transactions on Image Processing*, vol. 18, no. 2, pp. 281–298, 2009.
- [9] A. Barabell, “Improving the resolution performance of eigenstructure-based direction-finding algorithms,” in *Proc. Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP’83.*, vol. 8, pp. 336–339. IEEE, 1983.
- [10] F. Belloni, A. Richter, and V. Koivunen, “DoA estimation via manifold separation for arbitrary array structures,” *IEEE Transactions on Signal Processing*, vol. 55, no. 10, pp. 4800–4810, 2007.
- [11] J. Berent, P. L. Dragotti, and T. Blu, “Sampling piecewise sinusoidal signals with finite rate of innovation methods,” *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 613–625, 2010.

- 
- [12] S. Bhatnagar and T. Cornwell, "Scale sensitive deconvolution of interferometric images-I. adaptive scale pixel (Asp) decomposition," *Astronomy & Astrophysics*, vol. 426, no. 2, pp. 747–754, 2004.
- [13] T. Blu, H. Bay, and M. Unser, "A new high-resolution processing method for the deconvolution of optical coherence tomography signals," in Proc. *Proceedings of the First IEEE International Symposium on Biomedical Imaging: Macro to Nano (ISBI'02)*, vol. III, pp. 777–780, July 7–10, 2002.
- [14] T. Blu, P. L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 31–40, 2008.
- [15] T. Blu and M. Unser, "Quantitative Fourier analysis of approximation techniques: Part I—Interpolators and projectors," *IEEE Transactions on Signal Processing*, vol. 47, no. 10, pp. 2783–2795, October 1999.
- [16] R. Bolles, H. Baker, and D. Marimont, "Epipolar-plane image analysis: An approach to determining structure from motion," *International Journal of Computer Vision*, vol. 1, no. 1, pp. 7–55, 1987.
- [17] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, January 2011.
- [18] Y. Bresler and A. Macovski, "Exact maximum likelihood parameter estimation of superimposed exponential signals in noise," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 34, no. 5, pp. 1081–1089, 1986.
- [19] T. P. Bronez, "Sector interpolation of non-uniform arrays for efficient high resolution bearing estimation," in Proc. *Acoustics, Speech, and Signal Processing, 1988. ICASSP-88., 1988 International Conference on*, pp. 2885–2888. IEEE, 1988.
- [20] M. Brüggen, R. van Weeren, and H. Röttgering, "Simulating the toothbrush: evidence for a triple merger of galaxy clusters," *Monthly Notices of the Royal Astronomical Society: Letters*, vol. 425, no. 1, pp. L76–L80, 2012.
- [21] J. A. Cadzow, "Signal enhancement—A composite property mapping algorithm," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 36, no. 1, pp. 49–62, 1988.
- [22] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Communications on Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, 2014.
- [23] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [24] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, 1969.
- [25] R. E. Carrillo, J. D. McEwen, and Y. Wiaux, "PURIFY: a new approach to radio-interferometric imaging," *Monthly Notices of the Royal Astronomical Society*, vol. 439, no. 4, pp. 3591–3604, 2014.

- 
- [26] A. Cayley, "Note sur la méthode d'élimination de Bezout." *Journal für die reine und angewandte Mathematik*, vol. 53, pp. 366–367, 1857.
- [27] S. H. Chan, R. Khoshabeh, K. B. Gibson, P. E. Gill, and T. Q. Nguyen, "An augmented Lagrangian method for total variation video restoration," *IEEE Transactions on Image Processing*, vol. 20, no. 11, pp. 3097–3111, 2011.
- [28] H. Chang, D.-Y. Yeung, and Y. Xiong, "Super-resolution through neighbor embedding," in Proc. *2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, pp. I–I. IEEE, 2004.
- [29] C. Chen, P. Marziliano, and A. C. Kot, "2D finite rate of innovation reconstruction method for step edge and polygon signals in the presence of noise," *IEEE Transactions on Signal Processing*, vol. 60, no. 6, pp. 2851–2859, 2012.
- [30] Y. Chen and Y. Chi, "Robust spectral compressed sensing via structured matrix completion," *IEEE Transactions on Information Theory*, vol. 60, no. 10, pp. 6576–6601, October 2014.
- [31] D. Colton and R. Kress, *Inverse acoustic and electromagnetic scattering theory*. Springer Science & Business Media, 2012, vol. 93.
- [32] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-point algorithms for inverse problems in science and engineering*. Springer, 2011, pp. 185–212.
- [33] L. Condat and A. Hirabayashi, "Cadzow denoising upgraded: A new projection method for the recovery of Dirac pulses from noisy linear measurements," *Sampling Theory in Signal and Image Processing*, vol. 14, no. 1, pp. p–17, 2015.
- [34] J. Condon, W. Cotton, E. Greisen, Q. Yin, R. Perley, G. Taylor, and J. Broderick, "The NRAO VLA sky survey," *The Astronomical Journal*, vol. 115, no. 5, p. 1693, 1998.
- [35] J. Conway, T. Cornwell, and P. Wilkinson, "Multi-frequency synthesis—a new technique in radio interferometric imaging," *Monthly Notices of the Royal Astronomical Society*, vol. 246, p. 490, 1990.
- [36] J. L. Coolidge, *A treatise on algebraic plane curves*. Courier Corporation, 2004, vol. 1.
- [37] T. J. Cornwell, K. Golap, and S. Bhatnagar, "The noncoplanar baselines effect in radio interferometry: The W-projection algorithm," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 5, pp. 647–657, 2008.
- [38] R. Courant *et al.*, "Variational methods for the solution of problems of equilibrium and vibrations," *Bulletin of the American Mathematical Society*, vol. 49, no. 1, pp. 1–23, 1943.
- [39] A. Dabbech, C. Ferrari, D. Mary, E. Slezak, O. Smirnov, and J. S. Kenyon, "MORESANE: MOdel REconstruction by Synthesis-ANalysis Estimators—A sparse deconvolution algorithm for radio interferometric imaging," *Astronomy & Astrophysics*, vol. 576, p. A7, 2015.
- [40] S. Dai, M. Han, W. Xu, Y. Wu, and Y. Gong, "Soft edge smoothness prior for alpha channel super resolution," in Proc. *2007 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8. IEEE, 2007.

- 
- [41] B. De Moor, "Total least squares for affinely structured matrices and the noisy realization problem," *IEEE Transactions on Signal Processing*, vol. 42, no. 11, pp. 3104–3113, 1994.
- [42] S. Deslauriers-Gauthier and P. Marziliano, "Sampling signals with a finite rate of innovation on the sphere," *IEEE Transactions on Signal Processing*, vol. 61, no. 18, pp. 4552–4561, September 2013.
- [43] E. D. Di Claudio and R. Parisi, "WAVES: Weighted average of signal subspaces for robust wideband direction finding," *IEEE Transactions on Signal Processing*, vol. 49, no. 10, pp. 2179–2191, 2001.
- [44] J. H. DiBiase, *A high-accuracy, low-latency technique for talker localization in reverberant environments using microphone arrays*. Brown University Providence, 2000.
- [45] I. Dokmanić and Y. M. Lu, "Sampling sparse signals on the sphere: Algorithms and applications," *IEEE Transactions on Signal Processing*, vol. 64, no. 1, pp. 189–202, January 2016.
- [46] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, April 2006.
- [47] Z. Doğan, C. Gilliam, T. Blu, and D. Van De Ville, "Reconstruction of finite rate of innovation signals with model-fitting approach," *IEEE Transactions on Signal Processing*, vol. 63, no. 22, pp. 6024–6036, November 2015.
- [48] P. L. Dragotti, M. Vetterli, and T. Blu, "Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang-Fix," *IEEE Transactions on Signal Processing*, vol. 55, no. 5, pp. 1741–1757, May 2007.
- [49] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1-3, pp. 293–318, 1992.
- [50] E. Esser, "Applications of Lagrangian-based alternating direction methods and connections to split Bregman," *CAM report*, vol. 9, p. 31, 2009.
- [51] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advances and challenges in super-resolution," *International Journal of Imaging Systems and Technology*, vol. 14, no. 2, pp. 47–57, 2004.
- [52] M. Fatemi, A. Amini, and M. Vetterli, "Sampling and reconstruction of shapes with algebraic boundaries," *IEEE Transactions on Signal Processing*, vol. 64, no. 22, pp. 5807–5818, 2016.
- [53] R. Fattal, "Image upsampling via imposed edge statistics," in *Proc. ACM Transactions on Graphics*, vol. 26, no. 3, p. 95. ACM, 2007.
- [54] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *International journal of computer vision*, vol. 40, no. 1, pp. 25–47, 2000.
- [55] B. Friedlander, "The root-MUSIC algorithm for direction finding with interpolated arrays," *Signal Processing*, vol. 30, no. 1, pp. 15–29, 1993.



- 
- [56] D. Gabor, "A new microscopic principle," *Nature*, vol. 161, no. 4098, pp. 777–778, 1948.
- [57] H. Garsden, J. Girard, J.-L. Starck, S. Corbel, C. Tasse, A. Woiselle, J. McKean, A. S. Van Amesfoort, J. Anderson, I. Avruch, *et al.*, "LOFAR sparse image reconstruction," *Astronomy & Astrophysics*, vol. 575, p. A90, 2015.
- [58] C. Gilliam and T. Blu, "Fitting instead of annihilation: Improved recovery of noisy FRI signals," in Proc. *Proceedings of the Thirty-ninth IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'14)*, pp. 51–55, May 4–9, 2014.
- [59] —, "Finding the minimum rate of innovation in the presence of noise," in Proc. *Proceedings of the Forty-first IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'16)*, pp. 4019–4023, March 20–25, 2016.
- [60] D. Glasner, S. Bagon, and M. Irani, "Super-resolution from a single image," in Proc. *2009 IEEE International Conference on Computer Vision*, pp. 349–356. IEEE, 2009.
- [61] G. H. Golub and C. F. Van Loan, *Matrix computations*. Johns Hopkins University Press, 2012, vol. 3.
- [62] S. Haykin and J. A. Cadzow, eds., *Special Issue on Spectral Estimation*, vol. 70, no. 9, September 1982.
- [63] P. J. Hayuningtyas and P. Marziliano, "Finite rate of innovation method for DOA estimation of multiple sinusoidal signals with unknown frequency components," in Proc. *2012 9th European Radar Conference (EuRAD)*, pp. 115–118. IEEE, 2012.
- [64] L. He, H. Qi, and R. Zaretzki, "Beta process joint dictionary learning for coupled feature spaces with application to single image super-resolution," in Proc. *2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 345–352. IEEE, 2013.
- [65] F.-M. Hoffmann, F. M. Fazi, and P. Nelson, "Plane wave identification with circular arrays by means of a finite rate of innovation approach," in Proc. *Audio Engineering Society Convention 140*. Audio Engineering Society, 2016.
- [66] J. Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines," *Astronomy & Astrophysics Supplement Series*, vol. 15, p. 417, 1974.
- [67] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 38, no. 5, pp. 814–824, 1990.
- [68] M. M. Hyder and K. Mahata, "Direction-of-arrival estimation using a mixed  $ell_{2,0}$  norm approximation," *IEEE Transactions on Signal Processing*, vol. 58, no. 9, pp. 4646–4655, 2010.
- [69] H. Intema, P. Jagannathan, K. Mooley, and D. Frail, "The GMRT 150 MHz all-sky radio survey – first alternative data release TGSS ADR1," *Astronomy & Astrophysics*, vol. 598, p. A78, 2017.
- [70] T. Jiang, N. D. Sidiropoulos, and J. M. ten Berge, "Almost-sure identifiability of multi-dimensional harmonic retrieval," *IEEE Transactions on Signal Processing*, vol. 49, no. 9, pp. 1849–1859, 2001.

- 
- [71] B. Keinert, M. Innmann, M. Sanger, and M. Stamminger, “Spherical Fibonacci mapping,” *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, p. 193, 2015.
- [72] R. G. Keys, “Cubic convolution interpolation for digital image processing,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 29, no. 6, pp. 1153–1160, 1981.
- [73] K. I. Kim and Y. Kwon, “Single-image super-resolution using sparse regression and natural image prior,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 6, pp. 1127–1133, 2010.
- [74] S. Kunis, T. Peter, T. Romer, and U. von der Ohe, “A multivariate generalization of Prony’s method,” *Linear Algebra and its Applications*, vol. 490, pp. 31–47, 2016.
- [75] P. C. Lauterbur, “Image formation by induced local interactions: examples employing nuclear magnetic resonance,” *Nature*, vol. 242, pp. 190–191, 1973.
- [76] C. L. Lawson and R. J. Hanson, eds., *Solving least squares problems*. SIAM, 1995.
- [77] H. F. Lee, “A novel sub-pixel edge detection algorithm: with applications to super-resolution and edge sharpening,” Master’s thesis, The Chinese University of Hong Kong, 2013.
- [78] X. Li and M. T. Orchard, “New edge-directed interpolation,” *IEEE Transactions on Image Processing*, vol. 10, no. 10, pp. 1521–1527, 2001.
- [79] Y. Li and Y. Chi, “Off-the-grid line spectrum denoising and estimation with multiple measurement vectors,” *IEEE Transactions on Signal Processing*, vol. 64, no. 5, pp. 1257–1269, 2016.
- [80] C. J. Lonsdale, R. J. Cappallo, M. F. Morales, F. H. Briggs, L. Benkevitch, J. D. Bowman, J. D. Bunton, S. Burns, B. E. Corey, S. S. Doeleman, *et al.*, “The Murchison widefield array: Design overview,” *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1497–1506, 2009.
- [81] X. Lu, H. Yuan, P. Yan, Y. Yuan, and X. Li, “Geometry constrained sparse coding for single image super-resolution,” in *Proc. 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1648–1655. IEEE, 2012.
- [82] W.-K. Ma, T.-H. Hsieh, and C.-Y. Chi, “DOA estimation of quasi-stationary signals with less sensors than sources and unknown spatial noise covariance: a Khatri-Rao subspace approach,” *IEEE Transactions on Signal Processing*, vol. 58, no. 4, pp. 2168–2180, 2010.
- [83] D. Malioutov, M. Cetin, and A. S. Willsky, “A sparse signal reconstruction perspective for source localization with sensor arrays,” *IEEE transactions on signal processing*, vol. 53, no. 8, pp. 3010–3022, 2005.
- [84] I. Maravić and M. Vetterli, “Exact sampling results for some classes of parametric nonbandlimited 2-D signals,” *IEEE Transactions on Signal Processing*, vol. 52, no. 1, pp. 175–189, 2004.
- [85] —, “Sampling and reconstruction of signals with finite rate of innovation in the presence of noise,” *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 2788–2805, 2005.

- 
- [86] I. Markovsky, “How effective is the nuclear norm heuristic in solving data approximation problems?” in Proc. *16th IFAC Symposium on System Identification (Sysid 2012)*, vol. 45, no. 16, pp. 316–321. Elsevier, 2012.
- [87] A. Marquina and S. J. Osher, “Image super-resolution by TV-regularization and Bregman iteration,” *Journal of Scientific Computing*, vol. 37, no. 3, pp. 367–382, 2008.
- [88] P. Milanfar, G. C. Verghese, W. C. Karl, and A. S. Willsky, “Reconstructing polygons from moments with connections to array processing,” *IEEE Transactions on Signal Processing*, vol. 43, no. 2, pp. 432–443, 1995.
- [89] R. P. Millane, “Phase retrieval in crystallography and optics,” *Journal of the Optical Society of America A*, vol. 7, no. 3, pp. 394–411, 1990.
- [90] S. Mulleti and C. S. Seelamantula, “Ellipse fitting using the finite rate of innovation sampling principle,” *IEEE Transactions on Image Processing*, vol. 25, no. 3, pp. 1451–1464, 2016.
- [91] Y. Nakatsukasa, V. Noferini, and A. Townsend, “Computing the common zeros of two bivariate functions via Bézout resultants,” *Numerische Mathematik*, vol. 129, no. 1, pp. 181–209, 2015.
- [92] A. Offringa, B. McKinley, N. Hurley-Walker, F. Briggs, R. Wayth, D. Kaplan, M. Bell, L. Feng, A. Neben, J. Hughes, *et al.*, “WSClean: an implementation of a fast, generic wide-field imager for radio astronomy,” *Monthly Notices of the Royal Astronomical Society*, vol. 444, no. 1, pp. 606–619, 2014.
- [93] G. Ongie, S. Biswas, and M. Jacob, “Structured low-rank recovery of piecewise constant signals with performance guarantees,” in Proc. *2016 IEEE International Conference on Image Processing (ICIP)*, pp. 963–967, Sept 2016.
- [94] G. Ongie and M. Jacob, “Off-the-grid recovery of piecewise constant images from few fourier samples,” *SIAM Journal on Imaging Sciences*, vol. 9, no. 3, pp. 1004–1041, 2016.
- [95] H. Pan, T. Blu, and P. L. Dragotti, “Sampling curves with finite rate of innovation,” *IEEE Transactions on Signal Processing*, vol. 62, no. 2, pp. 458–471, January 2014.
- [96] H. Pan, T. Blu, and M. Vetterli, “Towards generalized FRI sampling with an application to source resolution in radioastronomy,” *IEEE Transactions on Signal Processing*, vol. 65, no. 4, pp. 821–835, 2017.
- [97] H. Pan, R. Scheibler, E. Bezzam, I. Dokmanić, and M. Vetterli, “FRIDA: FRI-based DOA estimation for arbitrary array layouts,” in Proc. *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3186–3190. IEEE, 2017.
- [98] H. Pan, M. Simeoni, P. Hurley, T. Blu, and M. Vetterli, “LEAP: Looking beyond pixels with continuous-space Estimation of Point sources,” *Astronomy & Astrophysics*, vol. 608, p. A136, 2017.
- [99] H. Pan, T. Blu, and M. Vetterli, “Efficient multi-dimensional Diracs estimation with linear sample complexity,” *IEEE Transactions on Signal Processing*, 2018, submitted to.

- 
- [100] T. Peleg and M. Elad, “A statistical prediction model based on sparse representations for single image super-resolution.” *IEEE Transactions on Image Processing*, vol. 23, no. 6, pp. 2569–2582, 2014.
- [101] R. Perley, F. Schwab, and A. Bridle, *Synthesis imaging in radio astronomy*. San Francisco, CA (US); Astronomical Society of the Pacific, January 1989.
- [102] T. Peter, G. Plonka, and R. Schaback, “Reconstruction of multivariate signals via Prony’s method,” *Proc. Appl. Math. Mech.*, to appear, 2017.
- [103] R. Prony, “Essai expérimental et analytique sur les lois de la dilatabilité des fluides élastiques et sur celles de la force expansive de la vapeur de l’alkool, à différentes températures,” *Journal de l’Ecole Polytechnique*, vol. 1, no. 2, pp. 24–76, 1795.
- [104] B. Rafaely, *Fundamentals of spherical array processing*. Springer, 2015, vol. 8.
- [105] F. Roemer, M. Haardt, and G. Del Galdo, “Analytical performance assessment of multi-dimensional matrix and tensor-based ESPRIT-type algorithms,” *IEEE Transactions on Signal Processing*, vol. 62, no. 10, pp. 2611–2625, 2014.
- [106] S. Rouquette and M. Najim, “Estimation of frequencies and damping factors by two-dimensional ESPRIT type methods,” *IEEE Transactions on signal processing*, vol. 49, no. 1, pp. 237–245, 2001.
- [107] M. Rubsamen and A. B. Gershman, “Direction-of-arrival estimation for nonuniform sensor arrays: from manifold separation to Fourier domain MUSIC methods,” *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 588–599, 2009.
- [108] M. J. Rust, M. Bates, and X. Zhuang, “Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM),” *Nature methods*, vol. 3, no. 10, pp. 793–796, 2006.
- [109] S. Sahnoun, K. Usevich, and P. Comon, “Multidimensional ESPRIT for damped and undamped signals: Algorithm, computations, and perturbation analysis,” *IEEE Transactions on Signal Processing*, vol. 65, no. 22, pp. 5897–5910, 2017.
- [110] R. Sault and M. Wieringa, “Multi-frequency synthesis techniques in radio interferometric imaging.” *Astronomy & Astrophysics Supplement Series*, vol. 108, 1994.
- [111] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE transactions on antennas and propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [112] F. Schwab, *Optimal gridding of visibility data in radio interferometry*. Cambridge University Press, 1984.
- [113] P. Serra, T. Westmeier, N. Giese, R. Jurek, L. Flöer, A. Popping, B. Winkel, T. van der Hulst, M. Meyer, B. S. Koribalski, *et al.*, “SoFiA: a flexible source finder for 3D spectral line data,” *Monthly Notices of the Royal Astronomical Society*, vol. 448, no. 2, pp. 1922–1929, 2015.
- [114] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.

- 
- [115] P. Shukla and P. L. Dragotti, "Sampling schemes for multidimensional signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3670–3686, 2007.
- [116] M. M. J.-A. Simeoni, "Towards more accurate and efficient beamformed radio interferometry imaging," Ecole Polytechnique Fédérale de Lausanne, Tech. Rep., 2015.
- [117] J.-L. Starck, F. Murtagh, and J. M. Fadili, *Sparse image and signal processing: wavelets, curvelets, morphological diversity*. Cambridge university press, 2010.
- [118] P. Stoica and R. L. Moses, *Spectral analysis of signals*. Pearson/Prentice Hall Upper Saddle River, NJ, 2005.
- [119] J. Sun, J. Sun, Z. Xu, and H.-Y. Shum, "Image super-resolution using gradient profile prior," in Proc. *2008 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–8. IEEE, 2008.
- [120] J. J. Sylvester, "On a theory of the syzygetic relations of two rational integral functions, comprising an application to the theory of Sturm's functions, and that of the greatest algebraical common measure," *Philosophical transactions of the Royal Society of London*, vol. 143, pp. 407–548, 1853.
- [121] H. Takeda, S. Farsiu, and P. Milanfar, "Kernel regression for image processing and reconstruction," *IEEE Transactions on Image Processing*, vol. 16, no. 2, pp. 349–366, 2007.
- [122] G. Tang, B. N. Bhaskar, and B. Recht, "Near minimax line spectral estimation," *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 499–512, 2015.
- [123] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht, "Compressed sensing off the grid," *IEEE Transactions on Information Theory*, vol. 59, no. 11, pp. 7465–7490, 2013.
- [124] C. Tasse, S. van der Tol, J. van Zwieten, G. van Diepen, and S. Bhatnagar, "Applying full polarization A-Projection to very wide field of view instruments: An imager for LOFAR," *Astronomy & Astrophysics*, vol. 553, p. A105, 2013.
- [125] G. B. Taylor, C. L. Carilli, and R. A. Perley, "Synthesis imaging in radio astronomy II," in Proc. *Synthesis Imaging in Radio Astronomy II*, vol. 180, 1999.
- [126] A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and synthesis in radio astronomy*. John Wiley & Sons, 2001.
- [127] R. Tsai and T. S. Huang, "Multiframe image restoration and registration," *Advances in computer vision and Image Processing*, vol. 1, no. 2, pp. 317–339, 1984.
- [128] M. Unser, "Sampling—50 years after Shannon," *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, 2000.
- [129] J. Urigüen, T. Blu, and P. L. Dragotti, "FRI sampling with arbitrary kernels," *IEEE Transactions on Signal Processing*, vol. 61, no. 21, pp. 5310–5323, November 2013.
- [130] A.-J. van der Veen and S. J. Wijnholds, "Signal processing tools for radio astronomy," in *Handbook of Signal Processing Systems*. Springer, 2013, pp. 421–463.

- 
- [131] M. Van Haarlem, M. Wise, A. Gunst, G. Heald, J. McKean, J. Hessels, A. De Bruyn, R. Nijboer, J. Swinbank, R. Fallows, *et al.*, “LOFAR: The low-frequency array,” *Astronomy & Astrophysics*, vol. 556, p. A2, 2013.
- [132] P. Vandewalle, “Super-resolution from unregistered aliased images,” Ph.D. dissertation, EPFL, 2006.
- [133] M. Vetterli, P. Marziliano, and T. Blu, “Sampling signals with finite rate of innovation,” *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [134] M. Vetterli, J. Kovačević, and V. K. Goyal, *Foundations of signal processing*. Cambridge University Press, 2014.
- [135] H. Wang and M. Kaveh, “Coherent signal-subspace processing for the detection and estimation of angles of arrival of multiple wide-band sources,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 4, pp. 823–831, 1985.
- [136] S. Wang, D. Zhang, Y. Liang, and Q. Pan, “Semi-coupled dictionary learning with applications to image super-resolution and photo-sketch synthesis,” in *Proc. 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2216–2223. IEEE, 2012.
- [137] Y. Wang, J. Yang, W. Yin, and Y. Zhang, “A new alternating minimization algorithm for total variation image reconstruction,” *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 248–272, 2008.
- [138] S. Wenger, M. Magnor, Y. Pihlström, S. Bhatnagar, and U. Rau, “SparseRI: A compressed sensing framework for aperture synthesis imaging in radio astronomy,” *Publications of the Astronomical Society of the Pacific (PASP)*, vol. 122, no. 897, pp. 1367–1374, Oct. 2010.
- [139] M. T. Whiting, “Duchamp: a 3D source finder for spectral-line data,” *Monthly Notices of the Royal Astronomical Society*, vol. 421, no. 4, pp. 3242–3256, 2012.
- [140] Y. Wiaux, L. Jacques, G. Puy, A. Scaife, and P. Vandergheynst, “Compressed sensing imaging techniques for radio interferometry,” *Monthly Notices of the Royal Astronomical Society*, vol. 395, no. 3, pp. 1733–1742, 2009.
- [141] Y. Wiaux, G. Puy, and P. Vandergheynst, “Compressed sensing reconstruction of a string signal from interferometric observations of the cosmic microwave background,” *Monthly Notices of the Royal Astronomical Society*, vol. 402, no. 4, pp. 2626–2636, 2010.
- [142] W. Williams, R. Van Weeren, H. Röttgering, P. Best, T. Dijkema, F. de Gasperin, M. Hardcastle, G. Heald, I. Prandoni, J. Sabater, *et al.*, “LOFAR 150-MHz observations of the Boötes field: catalogue and source counts,” *Monthly Notices of the Royal Astronomical Society*, vol. 460, no. 3, pp. 2385–2412, 2016.
- [143] F. Xue and T. Blu, “A novel SURE-based criterion for parametric PSF estimation,” *IEEE Transactions on Image Processing*, vol. 24, no. 2, pp. 595–607, February 2015.
- [144] J. Yang, J. Wright, T. S. Huang, and Y. Ma, “Image super-resolution via sparse representation,” *IEEE Transactions on Image Processing*, vol. 19, no. 11, pp. 2861–2873, 2010.

- 
- [145] J. C. Ye, J. M. Kim, K. H. Jin, and K. Lee, “Compressive sampling using annihilating filter-based low-rank interpolation,” *IEEE Transactions on Information Theory*, vol. PP, no. 99, pp. 1–1, 2016.
- [146] Y.-S. Yoon, L. M. Kaplan, and J. H. McClellan, “TOPS: New DOA estimator for wideband signals,” *IEEE Transactions on Signal Processing*, vol. 54, no. 6, pp. 1977–1989, 2006.
- [147] R. Zeyde, M. Elad, and M. Protter, “On single image scale-up using sparse-representations,” in *Curves and Surfaces*. Springer, 2012, pp. 711–730.





# Hanjie Pan

---

Route Cantonale 33  
1025 St-Sulpice (VD), Switzerland  
Tel: +41 78 879 0742  
hanjie.pan@epfl.ch  
<http://lcav.epfl.ch/people/hanjie.pan>

## RESEARCH INTERESTS

---

Continuous-domain sparse recovery (algorithms and applications in radio astronomy, direction of arrival estimation), image restorations with sparsity constraints

## EDUCATION

---

- 2013–2018**    **Ph.D. candidate in Computer and Communication Sciences**  
Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland  
Advisors: Prof. Martin Vetterli and Prof. Thierry Blu
- 2010–2013**    **M.Phil. in Electronic Engineering**  
The Chinese University of Hong Kong (CUHK), Hong Kong
- 2006–2010**    **B.Sc. in Electronic Engineering**  
The Chinese University of Hong Kong (CUHK), Hong Kong

## TEACHING AND SUPERVISION

---

Teaching assistant for

Advanced Engineering Mathematics  
Digital Image Processing  
Signal Processing for Communications  
Mathematical Foundations of Signal Processing

Co-advisor of undergraduate and M.Sc. students

## AWARDS AND HONORS

---

- 2013**            EPFL School of Computer and Communication Sciences fellowship
- 2010–2012**    CUHK Graduate School studentship
- 2010**            Prof. Charles Kao research exchange scholarship
- 2007, 2008**    CUHK Electronic Engineering Alumni Association Scholarship

## INDUSTRY

---

**IBM Research – Zürich**, Switzerland  
**November 2016 – April 2017**, Research Engineer  
Point source estimation from observations with the LOFAR radio telescope

## SKILLS

---

**Technical** Python,  $\LaTeX$ , Matlab, C/C++

**Languages** Mandarin Chinese, English, Cantonese (conversational)

## PROFESSIONAL ACTIVITIES

---

Reviewer for

IEEE Transactions on Image Processing (TIP)

IEEE Transactions on Signal Processing (TSP)

IEEE Transactions on Very Large Scale Integration (VLSI) Systems

IEEE Signal Processing Letters (SPL)

IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

Sampling Theory and Applications (SampTA)

Member of the IEEE Signal Processing Society

## PUBLICATIONS

---

### Journal Articles

- [8] H. Pan, M. Simeoni, P. Hurley, T. Blu, and M. Vetterli, "LEAP: Looking beyond pixels with continuous-space Estimation of Point sources", *Astronomy & Astrophysics*, vol. 608, A136, 2017.
- [7] H. Pan, T. Blu, and M. Vetterli, "Towards generalized FRI sampling with an application to source resolution in radioastronomy", *IEEE Transactions on Signal Processing*, vol. 65, no. 4, pp. 821–835, 2017.
- [6] H. Pan, T. Blu, and P.-L. Dragotti, "Sampling curves with finite rate of innovation", *IEEE Transactions on Signal Processing*, vol. 2, pp. 458–471, 2014.
- [5] H. Pan and T. Blu, "An iterated linear expansion of thresholds for  $\ell_1$ -based image restoration", *IEEE Transactions on Image Processing*, vol. 22, no. 9, pp. 3715–3728, 2013.

### Conference Papers

- [4] H. Pan, R. Scheibler, E. Bezzam, I. Dokmanić, and M. Vetterli, "FRIDA: FRI-based DOA estimation for arbitrary array layouts", in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 3186–3190.
- [3] H. Pan, T. Blu, and M. Vetterli, "Annihilation-driven localised image edge models", in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane: IEEE, 2015.
- [2] H. Pan and T. Blu, "Sparse image restoration using iterated linear expansion of thresholds", in *2011 IEEE International Conference on Image Processing (ICIP)*, Brussels, Belgium, 2011, pp. 1905–1908.
- [1] H. Pan, T. Blu, and P.-L. Dragotti, "Sampling curves with finite rate of innovation", in *2011 Sampling Theory and Applications (SampTA2011)*, Singapore, 2011.

