



Sparse Models and Convex Optimisation for Convolutional Blind Source Separation

Prasad Sudhakara Murthy

► To cite this version:

Prasad Sudhakara Murthy. Sparse Models and Convex Optimisation for Convolutional Blind Source Separation. Signal and Image processing. Université Rennes 1, 2011. English. NNT : . tel-00586610

HAL Id: tel-00586610

<https://theses.hal.science/tel-00586610>

Submitted on 18 Apr 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE / UNIVERSITÉ DE RENNES 1
sous le sceau de l'Université Européenne de Bretagne

pour le grade de
DOCTEUR DE L'UNIVERSITÉ DE RENNES 1

Mention : Traitement du Signal et Télécommunications

Ecole doctorale Matisse

présentée par

Prasad Sudhakara Murthy

préparée à l'unité de recherche l'INRIA
Institut National de Recherche en Informatique et Automatique

**Sparse Models and
Convex Optimisation
for Convolutional
Blind Source
Separation**

Thèse soutenue à Rennes

le 21 Février 2011

devant le jury composé de :

Jean-Jacques FUCHS

Professeur, Université de Rennes 1 /
président

Mark PLUMBLEY

Professor, Queen Mary, University of
London, UK / rapporteur

Michael ZIBULEVSKY

Senior Research Associate, Technion,
Israel / rapporteur

Abdeldjalil AISSA-EL-BEY

Associate Professor, Telecom Bretagne,
Brest, France / examinateur

Matthieu KOWALSKI

Assistant Professor, University Paris-Sud
11, France / examinateur

Rémi GRIBONVAL

Directeur de Recherche, INRIA Rennes /
directeur de thèse

To my parents and my sister.

$$0.1 \times 0.1 = 0.01$$

Acknowledgements

My deepest and heartfelt thanks to Rémi Gribonval for being a wonderful mentor during my PhD years and who still continues to inspire me. His sharp insights and clarity on scientific matters played a big role in successfully finishing my thesis. Despite my numerous shortcomings, he was always kind and patient enough to stand by me. Apart from being a brilliant researcher, he is also a wonderful human being and it is not an exaggeration if I say that I was simply lucky to be his student.

Many thanks to the members of the jury, Mark, Michael, J-J. Fuchs, Abdeljalil and Matthieu for their constructive feedback on my work, and also being helpful during the thesis defence.

My special thanks to Simon Arberet, my friend, ex-colleague and collaborator who helped me a great deal in understanding the subject and for bringing in his expertise into my work.

Life in INRIA was made very comfortable and memorable by the great company of METISS members. Thanks to Frédéric, Emmanuel, Guillaume for providing one of the best work environments. Many thanks to Armando for being a nice office-mate and for all the technical and non-technical discussions in the cafétéria. Thank you Nancy and Valentin for all the chiselling, especially during the defence. Thanks are due to all my other team members: Jules, Stefan, Sangnam, Ronan, Gabriel, Ngoc, Kamil, Alexey, Alexis, Boris and Stéphanie.

My stay in Rennes was made beautiful by all my other friends: Kumar, Hari, Ajay, Kamal, Pathi, Ramona, Sagar, Vijay, Basker, Yogesh, Charles, Kishore, and many more. Thank you guys, and I will cherish the memories of weekend cricket sessions in Rennes. A special thanks to my *local* friends: Basu and family, Neil and family, Claudine, Clément, Amélie and Tangui.

Thanks to Arjun and his team at the Apple care centre in Bangalore for their timely help in restoring my work from a crashed computer.

This endeavour was impossible if not for the constant encouragement and support that I got from my friends back in India. Many many thanks to all the L-boys, including the bearded variety..!

My boundless thanks to my parents and sister, who gave up their comforts to let me pursue a Ph.D, for their unconditional support and love. Lastly, thanks to my fiancée for patiently waiting for me..!

Contents

Acknowledgements	iii
List of figures	x
Notations and glossary	xi
0 Résumé en Français	1
1 Introduction	11
1.1 Source localisation and separation	12
1.2 Abstraction of the problem	14
1.2.1 Notations	15
1.2.2 A standard architecture for source separation systems	16
1.2.3 Types of mixing filters	17
1.2.4 Types of sources	18
1.3 Exploitation of sparsity	19
1.3.1 Filter sparsity	19
1.3.2 Source sparsity	20
1.4 The big picture	21
1.5 Plan of the thesis	22
1.6 Publications related to thesis contributions	24
I State of the art	27
2 The source separation problem	29
2.1 Introduction	29
2.1.1 Trivial setting	29
2.1.2 Plan of the chapter	30
2.2 Independent Component Analysis (ICA)	30
2.2.1 Hypothesis and principle of ICA	30
2.2.2 Algorithms for ICA	32
2.2.3 Limitations of ICA	32
2.2.4 ICA for convolutive mixtures	33

CONTENTS

2.2.4.1	Time domain methods	33
2.2.4.2	Frequency domain methods	33
2.3	Sparse Component Analysis (SCA)	35
2.3.1	Sparse representations	37
2.3.1.1	Sparsity of audio signals	38
2.3.1.2	Disjointness of the time-frequency supports	38
2.3.1.3	General principle of mixing parameter estimation	40
2.3.2	Linear instantaneous mixtures	42
2.3.2.1	Estimation of intensity parameter using global scatter plots	42
2.3.2.2	Estimation of intensity parameter using local scatter plots	43
2.3.3	Anechoic mixtures	44
2.3.4	Convulsive mixtures	45
2.3.4.1	Brief review of existing methods for blind estimation of $\tilde{\mathbf{S}}$	46
2.3.4.2	Estimation of $\tilde{\mathbf{S}}$ when \mathbf{A} is known	47
2.4	Summary	47
3	Solutions of underdetermined linear systems	49
3.1	Introduction	49
3.2	Sparse recovery problem	50
3.2.1	Algorithmic approaches	50
3.2.2	Recovery guarantees	52
3.3	Summary	52
4	Role of sparsity in channel estimation	55
4.1	Introduction	55
4.2	Methods for channel equalisation	56
4.3	Cross relation method for channel estimation	57
4.4	Sparse filter estimation	59
4.4.1	Sparse channel estimation: Single source setting	59
4.4.2	Sparse channel estimation: Multiple source setting	60
4.5	Summary	61
II	Contributions	63
5	Convulsive source localisation: Could we exploit filter sparsity?	67
5.1	Permutation and scaling problems: State of the art	68
5.1.1	Notations	69
5.1.2	Existing approaches for permutation correction	71
5.2	Assessing the impact of permutations on ℓ^1 norm of filters	73
5.3	Combinatorial ℓ^1 minimisation for solving permutation ambiguity	77
5.3.1	Algorithm	77
5.3.2	Complexity of the algorithm	78

5.3.3	Experimental results	79
5.3.3.1	Recovery in noiseless condition	79
5.3.3.2	Robustness to noise	81
5.4	Summary	82
6	Permutations and Sparsity: Theory	85
6.1	Notations	85
6.2	Main results	85
6.3	Proofs	86
6.3.1	Proof of Theorem 6.1	86
6.3.2	Proof of Theorem 6.2	87
6.4	Summary	90
7	A convex optimisation framework for single source localisation in sparse convolutive setting	91
7.1	Time-frequency domain cross-relation	92
7.1.1	Narrowband approximation	92
7.1.2	Time-frequency domain cross-relation	92
7.1.3	Filter estimation by convex optimisation	93
7.2	Filter recovery in single source setting: Experiments	94
7.2.1	Data models and performance measure	95
7.2.1.1	Sparse filter generation	95
7.2.1.2	Source generation	95
7.2.1.3	Performance measure	97
7.2.2	Recovery with full set of observations	97
7.2.3	Recovery with limited set of observations	99
7.3	Summary	100
8	A convex optimisation framework for multiple source localisation in sparse convolutive setting	103
8.1	General idea	104
8.2	A two stage framework	105
8.3	Multiple filter recovery using oracle clustering	106
8.3.1	Experiments with sinusoid sources	107
8.3.2	Experiments with Gabor sources	109
8.4	From narrowband to wideband cross-relation	114
8.4.1	Time-frequency wideband formulation	114
8.4.2	Experiments with Gabor sources	117
8.5	Blind clustering and filter estimation in a special convolutive setting	120
8.5.1	Setting and approach	120
8.5.2	Experiments with audio sources	121
8.6	Summary	123

CONTENTS

III Conclusion and perspectives	125
9 Conclusions and perspectives	127
Bibliography	132

List of Figures

1.1	Depiction of a typical musical concert scenario. What is actually heard by the ears is the combination of sounds that are generated by the different instruments.	13
1.2	Abstraction of the mixing process.	15
1.3	A typical two-stage source separation system.	16
1.4	Types of mixing filters	17
1.5	An example of time-frequency domain sparse signal	19
1.6	The landscape of approaches for source localisation and the placement of our contributions of this thesis.	22
2.1	An example of time-frequency domain sparse signal	39
2.2	Spectrograms of two sources	39
2.3	Geometric structure of scatter plot	41
2.4	Workflow of a typical mixing parameter estimation system based on the scatter plot of the mixtures.	42
2.5	Workflow of a typical frequency domain based mixing parameter estimation for convolutive mixtures.	46
4.1	An example of time-frequency domain sparse signal	62
4.2	Our contribution in the thesis.	66
4.3	Workflow of the filter estimation system based on the sparsity of sources and mixing filters.	66
5.1	Plots showing the increase of ℓ^1 norm of filter matrices due to permutations.	75
5.2	Box plots showing the sensitivity of ℓ^1 norm of filter matrices due to a single permutation of two sources.	76
5.3	Performance of Algorithm 1 in terms of average output SNR and number of sweeps.	80
5.4	Performance of Algorithm 1 in terms of average output SNR and number of sweeps under noisy settings.	83
7.1	Average output SNR versus sparsity k	98
7.2	Illustration of a sparse source.	99

LIST OF FIGURES

7.3	Transition diagram for sparse filter recovery as a function of $ \Omega /2L$ and sparsity k	101
8.1	Illustration of the time-frequency domain sparsity and disjointness of sources.	105
8.2	Block diagram of the proposed framework.	106
8.3	Number of source frames required for the output SNR to reach 20dB in the single source setting $N = 1$	108
8.4	Recovery performance of the narrowband CR approach for Type 2 sources.	110
8.5	Effect of STFT window size on narrowband CR approach.	112
8.6	Effect of clustering threshold ν on narrowband CR approach.	113
8.7	Effect of STFT window size on wideband CR approach.	118
8.8	Effect of clustering threshold ν on wideband CR approach.	119
8.9	Spectrograms of two sources	121
8.11	Performance of filter recovery using narrowband and wideband CR approach for $\theta_1 = 0.2$ radian.	123

Notations and glossary

Standard mathematical symbols

\mathbb{Z}	Set of integers
\mathbb{R}	Set of real numbers
\mathbb{C}	Set of complex numbers
\mathbf{i}	$\sqrt{-1}$

Vectors and matrices

\mathbf{a}	Vector
\mathbf{A}	Matrix
$\mathbf{A}^{\mathbf{T}}$	Matrix transpose
$\text{diag}(\mathbf{a})$	A diagonal matrix with \mathbf{a} as its diagonal
$\mathbf{I}_{N \times N}$	An identity matrix of size $N \times N$

Indices

t	Time index
f	Frequency index
j	Source index
i	Channel index
τ	Frame index

Sizes

N	Number of sources
M	Number of channels
L	Filter length
F	Size of short-time Fourier transform window
k	Sparsity

LIST OF FIGURES

Mixing filters

$\mathbf{A} \in \mathbb{R}^{M \times N \times L}$	Matrix of filters
$a_{ij}(t) \in \mathbb{R}$	Filter coefficient at t^{th} time index
$\mathbf{a}_{ij} \in \mathbb{R}^L$	Time domain filter vector
$\hat{\mathbf{a}}_{ij} \in \mathbb{C}^L$	Frequency domain filter vector
$\hat{a}_{ij}(f) \in \mathbb{C}$	Filter coefficient at f^{th} frequency index

Sources

$s_j(t) \in \mathbb{R}$	Source j at t^{th} time index
$\mathbf{s}_j \in \mathbb{R}^T$	Time-domain source vector of length T
$\hat{\mathbf{s}}_j \in \mathbb{C}^T$	(Time)-Frequency domain source vector of length T

Mixtures

$x_i(t) \in \mathbb{R}$	Mixture i at t^{th} time index
\mathbf{x}_j	Time-domain mixture vector
$\hat{\mathbf{x}}_j$	(Time)-Frequency domain mixture vector

Abbreviations

DFT	Discrete Fourier transform
CR	Cross relation
STFT	Short-Time Fourier Transform
dB	deciBel

Chapter 0

Résumé en Français

L'homme est souvent s'inspire souvent de la nature pour construire des systèmes cognitifs automatiques, comme un système auditif, dont les applications couvrent bien au-delà de ce qui est suggéré par la nature. Cette thèse est centrée sur un problème d'ingénierie, tels qui fait partie d'un système artificiel qui a son origine dans un problème biologique cognitive. Une caractéristique importante des systèmes auditifs biologiques est la capacité à *localiser* l'origine des sons. Autrement dit, les organismes ont un moyen de réaliser d'où provient le son dns l'espace de se rendre compte de l'endroit d'où provient le son. Illustrons cela avec l'exemple suivant.

Fig. 1 représente le scénario d'un concert de musique classique, où un certain nombre d'instruments de musique sont joués et le public arrive à écouter le son produit par l'ensemble. Dans un tel contexte, ce qui est réellement entendu par les oreilles est un *mélange* de sons différents qui sont générés par les différentes *sources*, les instruments. Les coordonnées physiques des sources sont différentes et, partant, les sons provenant de différents instruments empreuntent différents chemins multiples pour atteindre les oreilles de l'auditeur. Ils lesattergnent non seulement par la voie directe, mais aussi par de multiples réflexions depuis toutes les directions, comme représenté sur la figure.

Les systèmes humains-cognitifs sont capables de séparer les différentes sources mélange. En outre, l'emplacement physique des sources peu également être identifié. Dans le contexte biologique, le terme *localisation de source* signifie en grande partie la capacité d'extraire à partir des mélanges des informations sur les emplacements physiques réels des sources, directement ou indirectement, et *séparation de source* est la capacité à distinguer ou à séparer les différentes sons à partir des mélanges. Les systèmes biologiques sont capables de localiser les différentes sources et de séparer les sources, en dépit des interférences et du bruit ambiant.

Le problème d'ingénierie que nous traitons dans cette thèse est celui de la localisation de source. Dans le contexte de l'ingénierie, le terme localisation s'étend bien au-delà de la notion de simple spécification de l'emplacement physique des sources. En général, cela signifie la spécification de caractéristiques acoustiques et / ou physiques entre les sources et les oreilles (ou les capteurs, plus généralement). En outre, les problèmes de localisation et de séparation sont intrinsèquement liés et souvent la localisation des

0. RÉSUMÉ EN FRANÇAIS




Figure 1: Représentation d'un scénario de concert de musique classique. Ce qui est réellement entendu par les oreilles est la combinaison de sons qui sont générés par les différents instruments.

source précède la séparation des sources. Par conséquent, chaque fois que nous nous référons au terme *separation de sources* à partir de maintenant, il faut comprendre qu'il inclut le problème de la localisation des sources.

Au cours des deux dernières décennies, des efforts ont été entrepris par des communautés scientifiques, de divers horizons tels que traitement du signal, mathématiques, statistiques, réseaux de neurones, apprentissage, etc, pour construire des systèmes de localisation et de séparation de sources. Le lecteur pourra se référer à [1] pour un historique des recherches en séparation de sources. En dépit d'efforts considérables, la question n'est pas complètement encore résolue et il reste encore beaucoup de problèmes difficiles et sous-problèmes à résoudre afin de parvenir à une réponse satisfaisante.

Les systèmes de séparation de sources trouvent des applications dans une grande variété de domaines tels que le traitement du signal de musique, le traitement des signaux biomédicaux, de la radio-astronomie, de traitement d'image, etc. Les systèmes de séparation de sources ont une importance significative sur les différents aspects de la vie quotidienne et sur le développement de la science elle-même. Par conséquent, il nous a motivés et incités à entreprendre l'étude du problème de séparation de sources.

0.1 Le problème mathématique

Prenons N signaux sources (signaux numériques) et notons-les $s_j(t)$, $1 \leq j \leq N$. Considérons également M mélanges et notons-les $x_i(t)$, $1 \leq i \leq M$. Lorsque $M = 2$, on parle du cas des mélanges *stéréo*. Les trajets multiples dus aux réflexions sont modélisés par des réponses impulsionnelles discrètes de longueur L entre chaque paire source-microphone, désignées par $a_{ij}(t)$.

Le processus de mélange est modélisé comme un système linéaire, et à chaque micro i les observations ne sont rien d'autre qu'une somme de toutes les sources $s_j(t)$ qui ont été convoluées avec les filtres correspondants de mélange $a_{ij}(t)$. Mathématiquement, on peut abstraire le processus de mélange par :

$$x_i(t) = \sum_{j=1}^N (a_{ij} \star s_j)(t) + v_i(t), \quad (1)$$

où $a_{ij}(t)$ est un filtre de longueur L , et $v_i(t)$ est le bruit au i^{me} enregistreur. Par souci de concision, nous noterons les sources, les filtres, le bruit et les mélanges de s_j, a_{ij}, v_i et x_i respectivement, sans expliciter l'indice de temps.

Les réponses impulsionnelles a_{ij} , $1 \leq i \leq M$, $1 \leq j \leq N$ dépendent de divers facteurs tels que les positions relatives des sources et des micros, l'acoustique de l'environnement, les propriétés physiques des appareils, etc. Si les caractéristiques physiques et acoustiques de l'environnement sont fixes, alors les seuls facteurs qui influent sur les réponses impulsionnelles sont les positions relatives et ils caractérisent entièrement la localisation des sources.

En outre, le problème de localisation et de séparation des sources doit être résolu sans la connaissance explicite du processus de mélange. C'est-à-dire qu'aucune information sur les sources ou les filtres n'est à notre disposition, et donc le contexte est appelé *aveugle*. Avec cette notion, nous pouvons maintenant poser la question qui est au cœur de toutes les recherches en *localisation aveugle de sources*:

Question centrale en localisation de source: *Etant donné les mélanges $x_i(t)$, $i = 1 \dots M$, est-il possible d'obtenir des informations explicites sur les filtres a_{ij} , $1 \leq i \leq M$, $1 \leq j \leq N$?*

Techniquement, la définition du modèle de mélange est générique par nature et la nature exacte des signaux (l'audio, image, etc) est sans importance pour la modélisation des processus de mélange en soi. Cependant, il jouera un rôle crucial dans l'élaboration des algorithmes et des systèmes de séparation, qui le plus souvent exploitent les propriétés particulières des types de sources spécifiques. Cette thèse se situe dans le contexte de signaux audio, et certaines propriétés spécifiques des signaux audio seront exploitées dans nos contributions. En outre, le terme localisation de source fait physiquement sens lorsque les sources sous-jacentes sont des sources sonores.

0. RÉSUMÉ EN FRANÇAIS

0.2 Notations

Avant d'aller plus loin dans le problème, définissons d'abord quelques notations. Nous considérons des signaux source de longueur T : soient $\mathbf{s}_j, j = 1, \dots, N$ les vecteurs source, chacun de taille $T \times 1$.

Soit $\mathbf{a}_{ij} \in \mathbb{R}^L$ le vecteur correspondant au j^{me} filtre a_{ij} , et définissons $\mathbf{A}(t) \in \mathbb{R}^{M \times N}, 1 \leq t \leq L$ la matrice des coefficients des filtres à l'indice de temps t . Le coefficient (ij) de la matrice contient le vecteur filtre \mathbf{a}_{ij} de longueur L , le long de la troisième dimension.

De même, soient $\mathbf{x}_i, i = 1, \dots, M$ les vecteurs de longueur $T + L - 1$ correspondant à la somme de mélanges $x_i, i = 1, \dots, M$.

Soit $\mathbf{x}(t) = [x_1(t), \dots, x_M(t)]^T \in \mathbb{R}^{M \times 1}$ et $\mathbf{s}(t) = [s_1(t), \dots, s_N(t)]^T \in \mathbb{R}^{N \times 1}$ les vecteurs des mélanges et les sources, respectivement à l'indice de temps t . Ensuite, on peut écrire (1) comme

$$\mathbf{x}(t) = \sum_{\ell=0}^{L-1} \mathbf{A}(\ell) \mathbf{s}(t - \ell) + \mathbf{v}(t), \quad (2)$$

où $\mathbf{v}(t) \in \mathbb{R}^{M \times 1}$ est un vecteur d'échantillons de bruit à l'indice de temps t . Par souci de concision, nous utilisons aussi parfois les notations matricielles suivantes. Soit $\mathbf{S} \in \mathbb{R}^{N \times T}$ la matrice des vecteurs source définie comme $\mathbf{S} = [\mathbf{s}_1 \dots \mathbf{s}_N]^T$ et $\mathbf{X} \in \mathbb{R}^{M \times (T+L-1)}$ la matrice des vecteurs mélange défini par $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_M]^T$.

En notation abrégée, le processus de mélange convolutif est souvent écrit comme:

$$\mathbf{X} = \mathbf{A} \star \mathbf{S} + \mathbf{V}. \quad (3)$$

où \star représente l'opération de convolution.

0.3 Une architecture standard pour les systèmes de séparation de source

Une architecture typique pour les systèmes de séparation de source se compose de deux étapes. Dans la première étape, une estimation des filtres de mélange \tilde{a}_{ij} est d'abord obtenue à partir de la somme de mélanges x_i , avec une hypothèse appropriée sur les sources et / ou sur les filtres eux-mêmes. Une fois obtenue une estimation des filtres, une estimation \tilde{s}_j des sources réelles est obtenue en utilisant les filtres estimés. Fig. 2 représente un système de séparation de sources typique à deux étages.

Il est clair que la première étape de cette architecture standard est la localisation de source. Ainsi, en référence à l'architecture standard, l'objectif de cette thèse est la première étape.

Si l'on regarde l'équation de mélange (1), si les sources sont numérotées de 1 à N , il est clair que l'ordre des sources et des filtres correspondants n'a pas d'importance dans le processus de mélange. Les sources sont ordonnées uniquement pour la commodité de la modélisation. De même, les sources et les filtres peuvent être redimensionnés par des

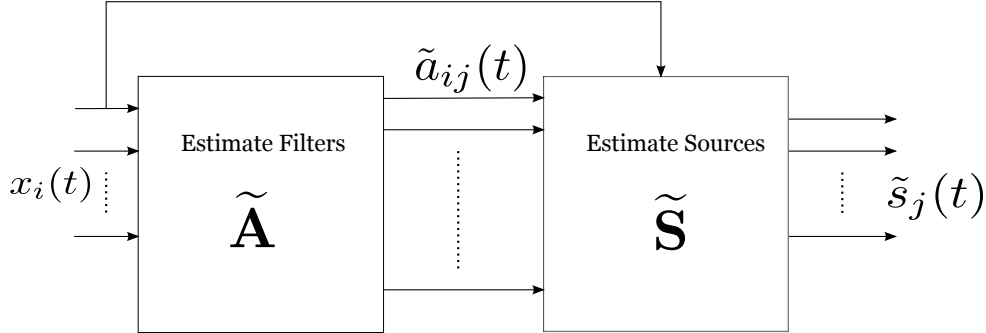


Figure 2: Un système typique à deux étages de séparation de sources.

facteurs arbitraires réciproque sans changer les mélanges. Par conséquent, les systèmes de séparation de sources permettent d'estimer les filtres, et les sources, seulement à une permutation et facteur d'échelle près. Ces ambiguïtés sont désignées par les termes ambiguïtés de *permutation* et *échelle*. Elles sont inhérentes au modèle lui-même et par conséquent, sont indépendantes du type d'algorithme ou de technique que nous utilisons pour effectuer la séparation des sources.

Chaque étape d'un système de séparation aveugle de sources repose sur certaines hypothèses sur les filtres de mélange, les sources ou les deux. Dans la section suivante, nous décrivons l'hypothèse utilisée dans ces travaux de thèse sur les filtres de mélange et sur les sources.

0.4 Filtres et sources parcimonieux

Lorsque l'environnement de mélange n'est pas réverbérant et quand il n'y a que quelques chemins de réflexion entre les sources et les micros, ce genre de mélange est modélisé par des filtres qui contiennent des pics très peu nombreux par rapport à la longueur des filtres, comme le montre la Fig. 3, et ils sont appelés filtres *parcimonieux* (en Anglais: *sparse*). Selon le dictionnaire de langue anglaise de l'Université d'Oxford, le mot *sparse* signifie de *finement dispersé*, ce qui est approprié dans notre exemple.

On recontre des modèles parcimonieux dans plusieurs domaines applicatifs tels que l'acoustique sous-marine, la géo-acoustique, les communications sans fil, etc. Dans cette thèse, les filtres de mélange parcimonieux sont notre intérêt premier, et la notion de filtre parcimonieux sera formalisée dans le chapitre 4.

Dans la plupart des approches de séparation de sources, les sources sont généralement modélisées en fonction de leurs propriétés statistiques. Les sources sont caractérisées en fonction de leurs propriétés telles que l'indépendance statistique, la stationnarité, la non-négativité, etc et en fonction des différentes distributions qu'elles peuvent suivre.

Outre les types ci-dessus mentionnés, les sources sont souvent modélisées comme parcimonieuses dans un certain domaine transformé, comme le domaine temps-fréquence (par exemple, la Transformée de Fourier à court terme (TFCT)), les ondelettes, etc.

0. RÉSUMÉ EN FRANÇAIS




Figure 3: Un exemple d'un filtre clairsemée.

Intuitivement, une source est dite parcimonieuse dans un certain domaine de la transformée si sa représentation dans le domaine transformé n'a que très peu de coefficients significatifs.

Fig 4 montre un exemple d'une source parcimonieuse. Fig 4(a) montre la représentation dans le domaine temporel d'un son de flûte et la figure 4(b) montre la courbe d'amplitude du même signal dans le domaine TFCT, où les régions des plus claires indiquent les coefficients de très faible amplitude. Notez que la source a un très petit nombre coefficients TFCT importants et par conséquent, elle est parcimonieuse dans le domaine TFCT.




Figure 4: Une exemple d'une source parcimonieuse dans le domaine temps-fréquence

Dans cette thèse, nous nous appuyons sur l'hypothèse de la parcimonie des sources et les détails techniques de cette hypothèse seront discutés plus tard dans le chapitre 2.

0.5 Exploitation de la parcimonie

Dans cette section, nous examinons comment la parcimonie des filtres et des sources peut être utilisée pour résoudre certains problèmes de séparation de sources. Tout d'abord, nous examinerons le cas où les sources sont connues et où nous cherchons à obtenir les filtres, qui sont parcimonieux dans le domaine temporel. C'est le problème de l'estimation de filtre (localisation de source) dans un cadre non-aveugle. Ensuite, nous examinerons de la situation inverse où les filtres sont connus et où nous voulons obtenir les sources, qui sont supposés être parcimonieux le domaine de la transformée temps-fréquence.

0.5.1 Filtre parcimonieux

Supposons que nous observons x , la convolutions d'un signal s par un filtre a :

$$x = a \star s. \quad (4)$$

En supposant que le signal s et le filtre a sont de longueur finie, nous pouvons les vectoriser en \mathbf{s} et \mathbf{a} . Ensuite, nous pouvons écrire l'équation (4) comme:

$$\mathbf{x} = \mathcal{T}[\mathbf{s}] \cdot \mathbf{a}, \quad (5)$$

où $\mathcal{T}[\mathbf{s}]$ est la matrice de Toeplitz formée à partir du vecteur source, et \mathbf{x} est le vecteur d'observation.

Le problème à résoudre est maintenant d'estimer le filtre \mathbf{a} à partir du vecteur d'observation \mathbf{x} , connaissant le vecteur source \mathbf{s} . Il s'agit d'un problème souvent rencontré dans les systèmes de communication sans fil. Dans un tel système, les caractéristiques du canal sont modélisées par un filtre à réponse impulsionnelle finie \mathbf{a} et le signal reçu \mathbf{x} est donné par l'équation (4). Dans le jargon des communications, \mathbf{a} est appelé *réponse impulsionnelle du canal*.

Généralement, la réponse impulsionnelle d'un canal est estimée par l'envoi de signaux d'étalonnage, dont la connaissance est disponible au niveau du récepteur, à intervalles intermittents. La réponse du canal estimé impulsion est ensuite utilisée pour estimer les signaux transmis.

Mathématiquement, l'équation (5) est une instance particulière d'une classe générale de problèmes linéaires inverses. Dans un problème linéaire inverse, nous sommes intéressés à trouver un vecteur \mathbf{a} en utilisant \mathbf{x} , de sorte que

$$\mathbf{x} = \Phi \mathbf{a}, \quad (6)$$

où Φ est une matrice. Eq. (5) est simplement un cas particulier de l'équation (6). avec $\Phi = \mathcal{T}[\mathbf{s}]$.

En outre, si la taille du vecteur \mathbf{x} est inférieure à celle du vecteur inconnu \mathbf{a} , alors nous avons un système sous-déterminé d'équations linéaires. Trouver des solutions aux problèmes linéaires inverses sous-déterminés est un domaine plus vaste de la recherche dont les implications sont bien au-delà de l'estimation de canal.

0. RÉSUMÉ EN FRANÇAIS

Les problèmes sous-déterminés sont mal posés, dans le sens où ils n'ont pas de solution unique et donc nous avons besoin de quelques hypothèses sur le genre de solution que nous recherchons afin de les résoudre. Une hypothèse qui a reçu beaucoup d'attention ces dernières années est la parcimonie de la solution. En général, nous sommes intéressés à trouver une solution parcimonieuse au système sous-déterminé d'équations linéaires, et cette famille de problèmes est collectivement appelé *problèmes de reconstruction parcimonieuse*.

Une littérature abondante existe sur les problèmes de reconstruction parcimonieuse dans une grande variété de contextes, et un résultat central de le problème est que si Φ et \mathbf{a} satisfont certaines conditions, alors une solution de l'équation (6) peut être trouvée en résolvant un problème d'optimisation impliquant Φ et \mathbf{x} .

Dans de nombreux contextes tels que l'acoustique sous-marine, l'exploration géologique, les communications sans fil, etc. les réponses impulsionnelles des canaux sont parcimonieux dans la nature et donc elles peuvent être estimées en résolvant un problème d'optimisation approprié, utilisant la connaissance des sources.

0.5.2 Parcimonie des sources

Nous avons mentionné et illustré dans Sec. 0.4 que les sources audio sont parcimonieuses dans le domaine temps-fréquence. Quand les filtres sont connus, la propriété de la parcimonie des sources peut être utilisée pour les estimer. Dans le cadre simple d'une seule source, la formulation du problème d'estimation de la source reste la même que celle décrite dans la section précédente, mais les rôles du filtre et de la source sont intervertis. Par conséquent, les sources peuvent être estimées en résolvant un problème d'optimisation.

Dans le cas de sources multiples, il y a des approches dans la littérature [2] qui exploitent la structure parcimonieuse des sources pour les estimer en utilisant la connaissance des filtres. En effet, l'hypothèse de source parcimonieuses a été largement utilisée dans le cas des mélanges linéaire instantanés et anéchoïques pour estimer les paramètres de mélange [3].

En plus de l'hypothèse de parcimonie des sources, on suppose qu'elles sont disjointes dans le domaine temps-fréquence. Cela signifie qu'à chaque point temps-fréquence, une seule de ces sources est active et contribue aux mélanges. Dans un tel cas, le processus de mélange est trivialisé à chaque point temps-fréquence et cette propriété est exploitée pour estimer les paramètres de mélange pour chaque source [4]. C'est l'idée sous-jacente à l'exploitation de la parcimonie des sources pour la localisation, et plusieurs généralisations et des améliorations ont été proposées dans la littérature [5, 6, 7].

Nous allons maintenant décrire la proposition de cette thèse, dont la contribution centrale est le mariage de la parcimonie des filtres et de celle des sources pour la tâche d'estimation aveugle de filtres.

0.6 Proposition de la thèse

Dans Sec. 0.5.1, nous avons vu comment la connaissance explicite de la source et la parcimonie du filtre peuvent être utilisées pour estimer un filtre, avec une seule observation. Au lieu d'une seule observation, supposons que nous avons maintenant deux observations x_1 et x_2 , à travers les filtres a_1 et a_2 . Il s'agit du paradigme "Une entrée-Deux sorties" (SITO: Single-Input-Two-Output), et en raison de la propriété de commutativité et d'associativité de l'opérateur de convolution, on peut écrire:

$$(x_2 \star a_1 - x_1 \star a_2)(t) = 0, \quad \forall t. \quad (7)$$

Cela implique la relation croisée (CR: Cross-Relation) [8]:

$$(x_2 \star a_1)(t) = (x_1 \star a_2)(t) = (a_2 \star a_1 \star s)(t), \quad \forall t. \quad (8)$$

La CR nous fournit une contrainte qui ne nécessite pas la connaissance des sources explicitement. Ceci a été largement utilisé dans le domaine des communications pour estimer les filtres de manière aveugle. De plus, si les filtres sont parcimonieuses, un problème d'optimisation peut alors être formulé et résolu pour estimer a_1 et a_2 simultanément. Il s'agit d'un système d'estimation aveugle, et on se place toujours dans le cadre d'une unique source. La question que nous pouvons nous poser maintenant est: comment pouvons-nous exploiter cette CR pour estimer des filtres de mélange dans le cadre de sources multiples?

C'est à ce stade que la parcimonie des sources dans le domaine temps-fréquence joue un rôle. Nous avons mentionné dans Sec. 0.5.2 qu'une hypothèse couramment utilisée est que les sources sont parcimonieuses et disjointes dans le domaine temps-fréquence. Avec cette hypothèse, nous pouvons trouver des points temps-fréquence dans les mélanges, pour lesquels une seule des sources est active. Par conséquent, même si la CR de l'équation (8) peut ne pas être satisfaite dans le domaine temporel pour n'importe quelle source, en raison de la présence d'autres sources, une version adaptée au domaine temps-fréquence de la CR sera satisfaite à ces points temps-fréquence où une seule source est active. Ensuite, ces CR temps-fréquence pour chacun de ces points peuvent être utilisées pour formuler un problème d'optimisation pour l'estimation de filtre.

Cette thèse est axée sur l'élaboration d'un cadre qui repose sur la parcimonie dans le domaine temps-fréquence et la disjonction des sources, ainsi que sur la version temps-fréquence de la CR. Fig 5 représente un schéma symbolique de notre contribution.

Nous proposons un cadre qui exploite la parcimonie des filtres et des sources pour la tâche d'estimation de multiples filtres parcimonieux à partir de mélanges convolutifs. Les principales tâches qui sont impliquées dans un tel cadre sont les suivantes: 1) Identifier les points temps-fréquence dans les mélanges où une seule source est active, et les regrouper entre eux selon les sources et 2) estimer le filtre en exploitant la CR temps-fréquence CR dans chaque groupe.

Un tel système aurait un flux de travail tel que décrit dans la figure 6. Les étapes peuvent être résumées comme suit:

0. RÉSUMÉ EN FRANÇAIS

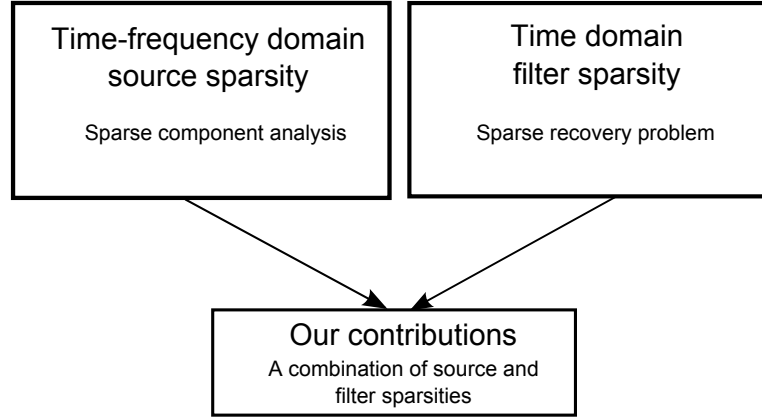


Figure 5: La contribution de notre la thèse.

1. Transformation des mélanges dans le domaine temps-fréquence, où la structure parcimonieuse des sources est rendue explicite;
2. Identification des points temps-fréquence où une seule source est active;
3. Regroupement des points et formulation du problème d'estimation de filtre;
4. Solution au problème d'estimation du filtre.




Figure 6: Flux de travail du système d'estimation de filtre basé sur la parcimonie des sources et des filtres de mélange.

Les chapitres 2 - 4 donnent un aperçu de l'état de l'art, les chapitres 5 - 8 présentant nos contributions originales, et la thèse se conclut sur les perspectives de travaux futurs dans le chapitre 9.

Chapter 1

Introduction

Survival instinct has been a key driving force behind the evolution of Life on earth. Nature has evolved efficient biological systems and subsystems in all the species that proliferate our planet, over billions of years. Human beings, with all their intelligence have been striving to mimic such perfect or near-perfect systems for their own benefit in their day-to-day lives. The first important eon in the history of engineering gave us the industrial revolution, where man built machines that mimicked the physical world. After that, the next important phase came in the form of information revolution. In the past few decades, scientists and engineers have been trying to build intelligent systems that can sense and process information in various physical forms.

Human beings are endowed with five information sensing systems and processing systems in the brain corresponding to the following information types: vision, sound, smell, taste and touch. Of all these, visual and sound information are generally processed continuously in a human brain. The very fact that the sensory organ for vision, the eye, has lids to cut off information, whereas the corresponding organ for sound, the ear, has no such facility, hints that sound is a very fundamental kind of information. This particular ability to process sound information efficiently has been the key to the very survival of many different species. For example, birds attract their mates by singing elaborate bird songs, they alert their own kind for some danger (predator birds) through distress calls, etc. Each purpose has a sound associated with it, and it is very important for a species to have an efficient sound sensing, processing and interpreting system.

Wherever there is information, then correspondingly associated with that is interference or noise. In a complex system such as our natural world, noise could originate from various sources and it is inevitable for the species to handle interference from unwanted sources. To that extent, Nature has done a fairly excellent job of evolving cognitive systems that can successfully separate the necessary information and unwanted interference. For example, we can easily follow a conversation in a crowded party hall, or follow the notes of a particular instrument in a music concert, etc. in spite of all the unwanted sources of sound that are present simultaneously.

It is natural that man is often inspired by Nature to build such cognitive systems artificially, and whose application scenarios span well beyond what is suggested by

1. INTRODUCTION

Nature. This thesis is focussed on one such engineering problem which forms a part of an artificial system that has its origins in a biological cognitive problem.

In the rest of this chapter, we shall informally introduce the problem that we will be dealing throughout this thesis, and present a bird's eye view of the standard approaches used to tackle it. Then, we briefly describe the central hypothesis on which this thesis is based on and we end this chapter by providing the layout of the thesis.

1.1 Source localisation and separation

As it was discussed previously, biological cognitive systems are highly sophisticated to differentiate between various sources of sound. That is, organisms are capable of *separating* various sounds that are heard simultaneously. This is achieved by the help of various cues that are associated with sounds such as the location of the sources, temporal and other structure of the sources, etc. One important feature of these biological systems is also the ability to *localise* the origin of the sounds. That is, the organisms have a way to realise the location from which the sound is originating. Let us illustrate this with the following example.

Fig. 1.1 depicts the scenario of a typical music concert, where a number of musical instruments are played and the audience gets to listen to the collective sound. In such a setting, what is actually heard by the ears is a *mixture* of various sounds that are generated by the various *sources*, the instruments. The physical co-ordinates of the sources are different and hence the sounds originating from different instruments take different multiple paths to reach the human ears. They not only reach the human ears via the direct path, but also through multiple reflections from all directions, as depicted in the figure.

In the biological context, the term *source localisation* largely means the ability to extract information from the mixtures about the actual physical locations of the sources directly or indirectly, and *source separation* is the ability to distinguish or separate multiple sounds from the mixtures. Biological systems are capable of localising the different sources and also separating the sources, in spite of interference and ambient noise.

Though the material presented in the first four chapters of the thesis is related to both source separation and localisation, the key engineering problem that we are dealing with is that of source localisation. In the engineering context, the term localisation extends well beyond the notion of simply specifying the physical locations of the sources. In general, it means the specification of some acoustic and/or physical characteristics between the sources and the ears (sensors, in general). Also, the problems of source localisation and separation are inherently connected and often source localisation precedes source separation. Hence, whenever we refer to the term source separation from now on, it should be understood that it includes the problem of source localisation.

Over the last couple of decades, efforts have been undertaken by the scientific community, from various backgrounds such as Signal Processing, Mathematics, Statistics, Neural Networks, Machine Learning, etc., to build source localisation and separation




Figure 1.1: Depiction of a typical musical concert scenario. What is actually heard by the ears is the combination of sounds that are generated by the different instruments.

systems. For historical notes on source separation research, please refer to [1]. Researchers have studied the source separation problem at various levels of complexity, and have come up with different approaches and systems. In spite of the huge efforts, the question is not completely answered yet and there remains a lot of difficult problems and sub-problems to be solved in order to arrive at a satisfactory answer.

In such a juncture, one might ponder about the significance of such systems. There are various reasons and use cases where source separation systems find applications, and here we will primarily quote the applications related to audio. We have already quoted the example of separating musical instruments from the recordings in a concert. Performing such a separation has various applications such as music transcription, selective enhancement, track extraction, etc. A related application is de-noising of general audio signals, where one wishes to separate out all the noise from the actual audio content.

Source separation also finds applications in biomedical signal processing. Suppose one wishes to measure the heart activity signal of an individual. The signals that are actually picked up by the sensors are a combination of the signals originating from the heart and other parts of the human body and noise. In such a scenario, it becomes increasingly important to separate out unwanted signals from the signal of interest. Hence, source separation is very relevant in this context.

In radio-astronomy, one studies the properties of stars and other astronomical objects by observing the radio signals that are emitted by them. It is very well known that the outer space consists of billions of stars and other bodies, and hence the signals

1. INTRODUCTION

that are sensed by the antennas on earth are highly corrupted. If one wishes to extract the signals of interest, then source separation is the way.

Source separation also has a number of applications in image processing, such as de-noising, etc.

With these representative applications, we can confidently say that source separation systems have a significant impact on the various aspects of day-to-day life, and more importantly, on the development of Science itself. Hence, it motivates and excites us to undertake the study of the source separation problem.

1.2 Abstraction of the problem

In all the above discussion, we have been very vague about the terms: sources, mixtures, mixing process, localisation and separation. Let us make them more precise in this section. Though in a real acoustic environment, the sound signals that are heard by human ears are nothing but mechanical waves that are continuous in nature, in practice the sound signals are digitized for the purposes of storing and processing them on digital computers. Hence, all the signals that we talk about in the source separation problem are digital signals.

Let us consider N source signals and denote them by $s_j(t)$, $1 \leq j \leq N$. Also, let us consider M mixtures and denote them by $x_i(t)$, $1 \leq i \leq M$. When $M = 2$, then it is referred to as *stereo* mixtures case. The multiple reflection paths are modelled by discrete impulse responses of length L between each pair of source and microphone, and this is denoted by $a_{ij}(t)$. An abstraction of the mixing scenario is illustrated in the Fig. 1.2. The circles on the left side of the diagram represent the microphones which observe the mixtures and the circles on the right side represent the sources. Between each pair of source and microphone is an impulse response that is represented by an arrow.

The mixing process is modelled as a linear system, and at each microphone i the observations are nothing but a summation of all the sources $s_j(t)$ that have been convolved with the corresponding mixing filters $a_{ij}(t)$. Mathematically, we can abstract the mixing process as

$$x_i(t) = \sum_{j=1}^N (a_{ij} \star s_j)(t) + v_i(t), \quad (1.1)$$

where $a_{ij}(t)$ is a filter of length L , $v_i(t)$ is the noise at the i^{th} recorder and \star is the convolution operation. For brevity, we denote the sources, filters, noise and mixtures by s_j, a_{ij}, v_i and x_i respectively, by dropping the time index.

The impulse responses a_{ij} , $1 \leq i \leq M$, $1 \leq j \leq N$ depend upon various factors such as the relative locations of the sources and the microphones, acoustics of the environment, physical properties of the devices, etc. If the physical and acoustic characteristics of the environment are fixed, then the only factor that influences the impulse responses




Figure 1.2: Abstraction of the mixing process.

are the relative physical locations and they completely characterise the localisation of the sources.

Further, source separation and localisation problems have to be solved without the explicit knowledge of the mixing process. That is, no information about either the sources or the filters are available to us, and hence the context is referred to as *blind*. With this notion, we can now pose the question that is at the heart of all the *blind source localisation* research:

Central question in source localisation: *Given the mixtures $x_i(t)$, $i = 1 \dots M$, is it possible to obtain explicit information about the filters a_{ij} , $1 \leq i \leq M$, $1 \leq j \leq N$?*

Technically, the mixing model definition is generic in nature and the exact nature of the signals, like audio, image, etc. is immaterial to the modeling of the mixing process per se. However, it will play a crucial role in devising the separation algorithms and systems, which more often than not exploit the special properties of specific source types. This thesis is set in the context of audio signals, and certain specific properties of audio signals will be exploited in our contributions. Also, the term source localisation makes a physical sense when the underlying source signals are audio.

1.2.1 Notations

Before we proceed further into the problem, let us first setup some notations. Let us consider sources of length T and let \mathbf{s}_j , $j = 1, \dots, N$ be the source vectors, each of size

1. INTRODUCTION

$T \times 1$. Let $\mathbf{a}_{ij} \in \mathbb{R}^L$ be the vector corresponding to the filter a_{ij} , and let us define $\mathbf{A}(t) \in \mathbb{R}^{M \times N}$, $1 \leq t \leq L$ to be the matrix of filter coefficients at time index t .

Similarly, let \mathbf{x}_i , $i = 1, \dots, M$ be the vectors of length $T + L - 1$ corresponding to the mixtures x_i , $i = 1, \dots, M$. Let $\mathbf{x}(t) = [x_1(t), \dots, x_M(t)]^T \in \mathbb{R}^{M \times 1}$ and $\mathbf{s}(t) = [s_1(t), \dots, s_N(t)]^T \in \mathbb{R}^{N \times 1}$ be vectors of mixtures and sources respectively at time index t . Then, we can write (1.1) as

$$\mathbf{x}(t) = \sum_{\ell=0}^{L-1} \mathbf{A}(\ell) \mathbf{s}(t - \ell) + \mathbf{v}(t), \quad (1.2)$$

where $\mathbf{v}(t) \in \mathbb{R}^{M \times 1}$ is a vector of noise samples at time index t . For the sake of brevity we shall also sometimes use the following matrix notations. Let $\mathbf{S} \in \mathbb{R}^{N \times T}$ be the matrix of source vectors defined as $\mathbf{S} = [\mathbf{s}_1 \dots \mathbf{s}_N]^T$ and let $\mathbf{X} \in \mathbb{R}^{M \times (T+L-1)}$ be the matrix of mixture vectors defined as $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_M]^T$.

As a shorthand notation, the convolutive mixing process is often written as

$$\mathbf{X} = \mathbf{A} \star \mathbf{S} + \mathbf{V}. \quad (1.3)$$

where \star represents the convolution operation.

1.2.2 A standard architecture for source separation systems

A typical system level architecture for source separation systems consists of two stages. In the first stage, an estimate of the mixing filters \tilde{a}_{ij} is first obtained starting from the mixtures x_i , with some suitable hypothesis on either the sources and/or the filters themselves. Once an estimate of the filters are obtained, then an estimate \tilde{s}_j of the actual sources are obtained using the filter estimates. Fig. 1.3 depicts a typical two-stage source separation system.

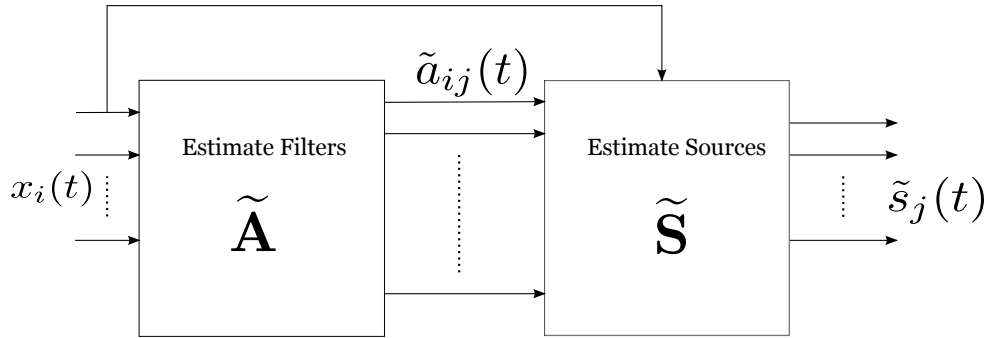


Figure 1.3: A typical two-stage source separation system.

It is clear that the first stage of this standard architecture is source localisation. So, with reference to the standard architecture, the focus of this thesis is the first stage.

If we look at the mixing equation (1.1), though the sources are numbered from 1 to N , it is clear that the order of the sources and corresponding filters do not matter

for the mixing process. The source ordering is just for the convenience of modelling. Similarly, the sources and filters can be scaled by arbitrary reciprocal factors and still the mixtures remain the same. As a consequence of this, source separation systems can estimate filters and sources only up to a permutation and scaling factor. These ambiguities are termed as the *permutation* and *scaling* ambiguities. These are inherent to the model itself and hence it is independent of the kind of algorithm or technique we use to perform source separation.

Each stage of a blind source separation system relies on certain hypotheses on either the mixing filters or the sources or both. In the next section, let us describe some of the common types of hypothesis that are assumed on the mixing filters and the sources.

1.2.3 Types of mixing filters

Several simplifications can be made to the mixing filter model and in each case one obtains a different kind of mixture. Fig. 1.4 shows the different possible types of the mixing filters a_{ij} . In the simplest of the cases, the filter could be a delta function $a_{ij} = c_{ij} \cdot \delta_{ij}(t)$, as illustrated in Fig. 1.4(a). In such a case, a mixture at a given time index t is essentially a weighted sum of the source signals at the same time instant. Hence, such mixtures are called *linear-instantaneous* mixtures.

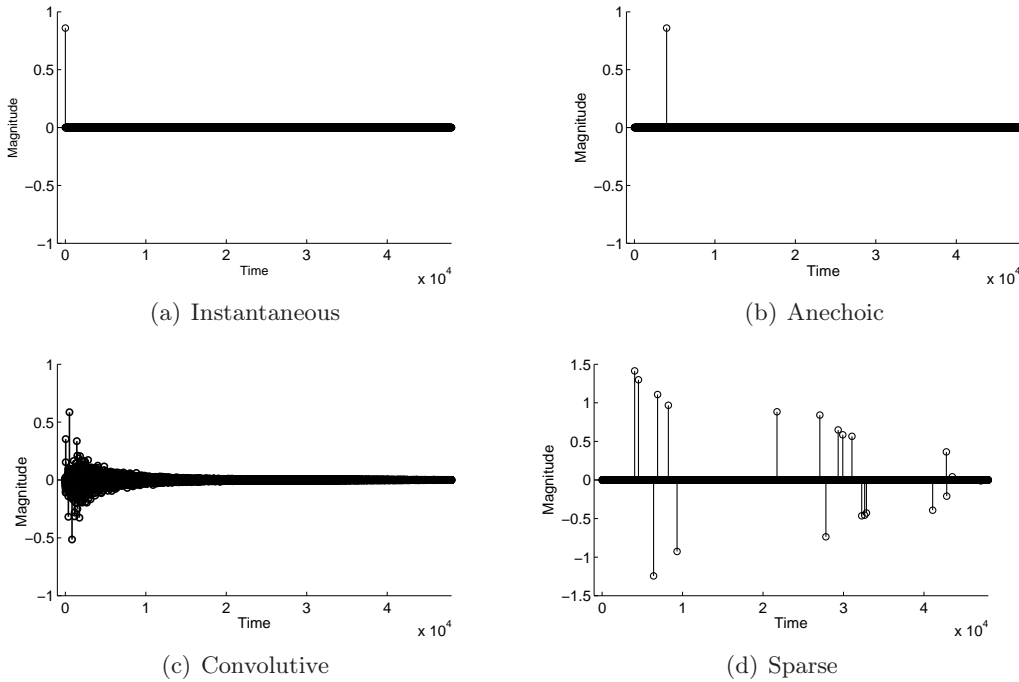


Figure 1.4: Types of mixing filters

The second kind of filters are shifted delta functions $a_{ij} = c_{ij} \cdot \delta_{ij}(t - t_{ij})$, as shown in Fig. 1.4(b). In this case, the mixtures are obtained by the summation of shifted and

1. INTRODUCTION

scaled sources and hence the mixtures are called *anechoic* mixtures.

In a realistic environment the source signals will go through echos and reverberations and hence the linear instantaneous and anechoic filter models are not capable of capturing the complexity of the mixing process. Hence, these kind of mixing is modelled using finite impulse filters as shown in Fig. 1.4(c). The mixtures are then called *convolutive* mixtures, and they are difficult to work with in general.

However, when the mixing environment is not reverberant and when there are only a few reflection paths between the sources and the microphones, the complexity of the mixing process lies in between the anechoic and fully convolutive cases. This kind of mixing is modelled by filters which contain very few peaks compared to the length of the filters as shown in Fig. 1.4(d), and they are called *sparse* filters. According to the English language dictionary of Oxford University, the word *sparse* means *thinly dispersed or scattered*, which is appropriate in our example.

Sparse filter models find application in several fields such as underwater acoustics, geo-acoustics, wireless communications, etc. In this thesis, sparse mixing filters are of our primary interest and the notion of filter sparsity will be formalised in chapter 4.

Over-determined, determined and underdetermined mixtures: Apart from the classification of mixtures based on the kind of mixing filters, we can also have a classification based on the relationship between the number of sources N and the number of mixtures M .

When $N < M$, then the number of observable variables are more than the unknown variables and hence it is referred to as an *over-determined* case. If $N = M$, then we have as many observable variables as unknowns, and hence this is a *determined* case. The most difficult case is when the number of unknowns are more than the number of observable variables, $N > M$, which is aptly called the *underdetermined* case.

1.2.4 Types of sources

In most of the approaches for source separation, the sources are generally modelled based on their statistical properties. The sources are characterised based on their properties such as statistical independence, stationarity, non-negativity, etc. and the distributions they assume in various cases.

Apart from the above mentioned types, sources are often modelled as sparse in a certain transform domain such as the time-frequency (e.g. Short Time Fourier Transform (STFT)), Wavelets, etc. Intuitively, a source is said to be sparse in a certain transform domain if its transform domain representation has only a very few significant coefficients.

Fig. 1.5 shows an example of a sparse source. Fig. 1.5(a) shows the time domain plot of a flute sound and Fig. 1.5(b) shows the magnitude plot of the same signal in the STFT domain, and the lighter regions indicate the very low magnitude coefficients. Notice that the source has only a very few significant STFT coefficients and hence it is sparse in the STFT domain.




Figure 1.5: An example of time-frequency domain sparse signal

In this thesis, we rely on the assumption of the sparsity of sources and the technical details of the hypothesis will be discussed later in chapter 2.

1.3 Exploitation of sparsity

As we have mentioned repeatedly, sparsity of mixing filters and sources is the central theme of this thesis. In this section, let us examine how the sparsity of these objects helps us to solve different problems. Firstly, we will look at the case where the sources are known to us and we are interested in obtaining the filters, which are assumed to be sparse in the time-domain. This is the problem of filter estimation (source localisation) in a non-blind setting. Then, we will look at the converse situation where the filters are known to us and we want to obtain the sources, which are assumed to be sparse in a time-frequency transform domain.

1.3.1 Filter sparsity

Suppose we observe x , a signal s which is convolved with a filter a :

$$x = a \star s. \quad (1.4)$$

Assuming finite lengths for the signal s and the filter a , we can vectorise them as \mathbf{s} and \mathbf{a} . Then we can write Eq. (1.4) as:

$$\mathbf{x} = \mathcal{T}[\mathbf{s}] \cdot \mathbf{a}, \quad (1.5)$$

where $\mathcal{T}[\mathbf{s}]$ is the Toeplitz matrix formed using the source vector, and \mathbf{x} is the observation vector.

The problem at hand now is to estimate the filter \mathbf{a} from the observation vector \mathbf{x} , with the knowledge of the source vector \mathbf{s} . This is a problem that is often encountered in wireless communication systems. In such a system, the channel characteristics are

1. INTRODUCTION

modelled by a finite impulse response filter \mathbf{a} and the received signal \mathbf{x} is given by Eq. (1.4). In the communications engineering parlance, \mathbf{a} is known as *channel impulse response*.

Typically, the impulse response of a channel is estimated by sending out pilot signals, whose knowledge is available at the receiver, at intermittent intervals. The estimated channel impulse response is then used to estimate the transmitted signals.

Mathematically, Eq. (1.5) is a particular instance of a general class of linear inverse problems. In a linear inverse problem, we are interested in finding a vector \mathbf{a} using \mathbf{x} , such that

$$\mathbf{x} = \Phi \mathbf{a}, \quad (1.6)$$

where Φ is a matrix. Eq. (1.5) is just a special case of Eq. (1.6) with $\Phi = \mathcal{T}[\mathbf{s}]$.

Further, if the size of the vector \mathbf{x} is smaller than the unknown vector \mathbf{a} , then we have an underdetermined system of linear equations. Finding solutions to underdetermined linear inverse problems is a broader area of research whose implications are much beyond channel estimation.

Underdetermined problems are ill-posed in the sense that they do not have unique solutions and hence we need some hypothesis about the kind of solution we are looking for in order to solve them. One such hypothesis that has received a lot of attention in the recent years is the sparsity of the solution. In general, we are interested in finding a sparse solution to underdetermined system of linear equations, and this family of problems is collectively called as *sparse recovery* problems.

A large body of literature exists concerning sparse recovery problems in a wide variety of settings, and a central result of sparse recovery is that if the matrix Φ and vector \mathbf{a} satisfy certain conditions, then a sparse solution of Eq. (1.6) can be found by solving an optimisation problem involving Φ and \mathbf{x} .

In many settings such as underwater acoustics [9], geological exploration [10], wireless communications [11], etc. the concerned channel impulse responses are sparse in the time-domain and hence they can be estimated by solving an appropriate optimisation problem, with the knowledge of the sources.

1.3.2 Source sparsity

We mentioned in Sec. 1.2.4 that audio sources are sparse in the time-frequency domain, and also saw an illustration. When the filters are known to us, then the sparsity property of the sources can be utilised to estimate them. In the simple setting of one source, the formulation of the source estimation problem remains the same as described in the previous section, but the roles of the filter and the source are interchanged. Hence, the sources can be estimated by solving an optimisation problem.

In the case of multiple sources, there are approaches in literature [2] that exploit the sparse structure of the sources to estimate them by using the knowledge of the filters. Indeed, sparse source hypothesis has been extensively used in the linear-instantaneous and anechoic mixtures cases to blindly estimate the mixing parameters [3].

Along with the hypothesis that the sources are sparse, they are also assumed to be disjoint in the time-frequency domain. This means that at each time-frequency point, only one of the sources is active and contributes to the mixtures. This property is then exploited to estimate the mixing parameters for each source [4]. This has been the basic idea behind exploiting source sparsity for source localisation, and several generalisations and improvements have been proposed in literature [5, 6, 7].

We now describe the proposition of this thesis, whose central contribution is the marriage of filter sparsity and source sparsity for the task of blind filter estimation.

1.4 The big picture

In this section, we shall take a bird's eye view of the approaches that we have presented till now, and motivate the contribution chapters.

Fig. 1.6 is a depiction of the landscape of approaches for source localisation and separation problems. The x-axis represents the possible hypothesis that one can assume about the mixing filters and y-axis represents the possible hypothesis about the sources. In chapter 2, we discuss briefly about the usage of different kinds of source hypotheses such as independence and sparseness that can be used for source separation. ICA based techniques were primarily developed to solve instantaneous mixtures problem, and they were extended to also anechoic and convolutive settings with appropriate modifications. ICA or the source sparsity based approaches do not assume anything about the sparsity of the mixing filters.

Sparsity based techniques were first developed to handle underdetermined settings for linear-instantaneous and anechoic mixtures, and then extended to convolutive mixtures primarily based on the narrowband approximation. As it was described in the previous section, narrowband approximation based approaches for convolutive source separation which exploit source sparsity, involve two steps: 1) a clustering step to identify the time-frequency points in the mixtures which belong to one of the sources and 2) a permutation solving step before the actual parameter estimation.

In chapter 4, we look at the problem of estimating channel impulse responses from observations in single source and multiple source settings in general. Further, assuming that the filters are sparse, we discussed a specific family of approaches based on the cross-correlation across channels for blind filter estimation in single source setting. Subsequently, a particular work on the extension of single source CR based filter estimation method to multiple source setting was discussed. In this discussion, no hypothesis about the time-frequency domain sparsity of sources was placed.

It should be noted that non-sparse methods for convolutive source separation suffer from arbitrary permutation and scaling of subbands. Also, source *separation* using SCA requires the knowledge of the mixing matrix \mathbf{A} . The objective of our effort is to develop an approach for estimating mixing filters blindly in the multiple source setting, which inherently handles the permutation and scaling problems. From this requirement point of view, we can summarise our discussion as follow:

1. On the one hand, methods relying on source sparsity for convolutive source locali-

1. INTRODUCTION

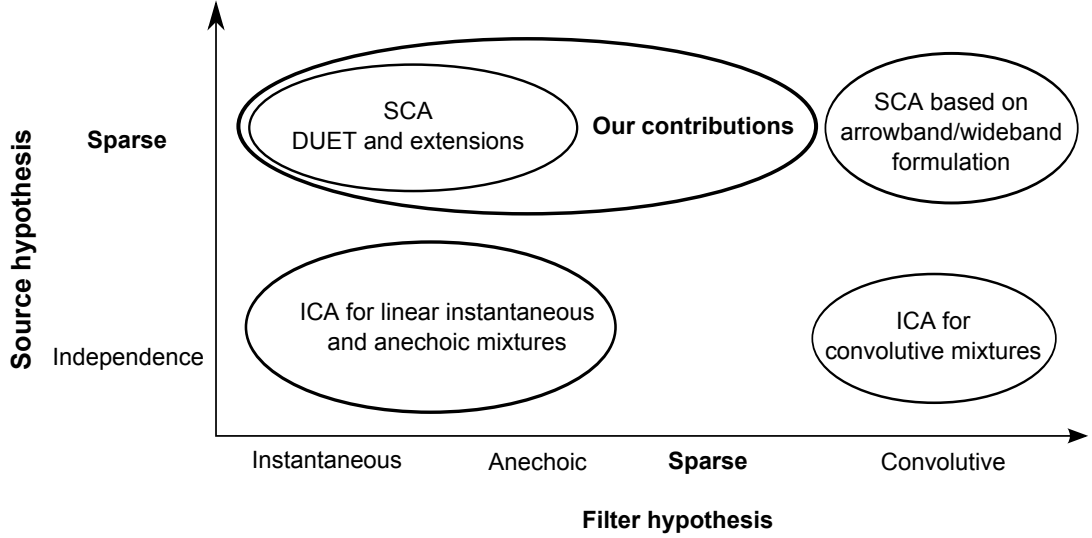


Figure 1.6: The landscape of approaches for source localisation and the placement of our contributions of this thesis.

sation have to perform clustering and permutation alignment in order to estimate the filters.

2. On the other hand, sparsity hypothesis of filters enables us to formulate an ℓ^1 minimisation problem for blind filter estimation but it is effective only in the single source setting.

The central idea in this thesis is to combine both the source and filter hypothesis and propose a filter estimation framework. As depicted in Fig. 1.6, our contribution falls in the region where both the sources and filters are assumed to be sparse. As a first contribution, we first show the effectiveness of filter sparsity to solve the permutation problem, when the filter coefficients are estimated using the narrowband approximation. Subsequently, we develop a framework which relies on the time-frequency domain source sparsity to formulate ℓ^1 minimisation problems for filter estimation, which inherently incorporates the filter sparsity to effectively mitigate the permutation and scaling problems.

1.5 Plan of the thesis

This thesis is divided into three parts:

- State of the art
- Contributions
- Conclusions and perspectives

The first part of the thesis focusses on the state of the art in sparse methods for source localisation. Firstly we present a formal introduction to the general problem of source localisation and separation, describe special cases of the problem and a survey of existing approaches and techniques in literature. Chapter 2 contains a mathematical introduction to the problem and the notations and terminologies used in the thesis. An overview of source separation methods based on different kinds of hypothesis about the sources (independent sources, sparsity, etc.) and separation criteria is presented.

The second half of chapter 2 exclusively focusses on approaches that have sparsity of sources or mixing filters as their main tool for source localisation. Firstly the use of sparsity of sources for mixing parameter estimation for linear-instantaneous and anechoic mixtures are described. These methods are based on the assumption that in the time-frequency domain, the sources have almost-disjoint support and this enables us to estimate the mixing parameters geometrically. We describe a very well known approach known as Degenerate Unmixing Estimation Technique (DUET) [4], which was proposed for anechoic stereo mixtures of fully disjoint sources. Subsequent improvements and generalisations are also presented.

The usage of sparsity to find solutions for underdetermined systems of linear equation is the topic of chapter 3. We present an introduction to the problem and a brief survey of grand families of approaches based on sparsity for approaching this problem.

Chapter 4 focusses on the problem of channel estimation in communication systems. The problem is introduced, and parallels are with the source localisation problem. A brief overview of non-blind methods that are developed by the communications engineering community is presented.

The second part of chapter 4 deals with the techniques that have been developed specifically for estimating sparse channels. Sparse channel estimation has flourished in the recent years, after an affluence of theoretical and algorithmic developments in the area of sparse recovery. We describe the relevant work in this context and point out their limitations.

The chapter is concluded with the positioning of the problem that is explored in this thesis within the big picture of various problem types and the corresponding standard approaches.

Chapters 5 to 8 contain the contributions of the thesis.

In chapter 5, our focus is on the assessment of the usage of sparsity criterion for solving a specific problem that arises in convolutive source localisation: the permutation ambiguity. We consider the permutation problem in isolation and we explore the possibility of using the sparsity of filters as a tool to correct the permutation ambiguity. Assuming that the mixing filters in the frequency domain are already estimated by an independent technique, with sub-bands that are randomly permuted, we aim to correct these permutations using the time-domain sparsity of the filters as a consistency measure. We first demonstrate how the sub-band permutations of the filters affect the overall time-domain sparsity of the filters, and substantiate the usage of sparsity as a consistency measure.

A combinatorial algorithm to correct the permutation ambiguity is then presented

1. INTRODUCTION

and the algorithm is experimentally evaluated for its performance.

Theoretical connections between filter sparsity and permutations are presented in chapter 6.

In chapter 7, we consider the problem of estimating sparse filters in a Single-Input-Two-Output (SITO) setting using cross-relation based approach. We introduce a time-frequency domain version of the time-domain CR, based on the narrowband approximation, and formulate a convex optimisation problem for estimating the filters. Experimental evaluation of the time-frequency domain narrowband CR based filter estimation approach is presented in the second half of the chapter.

In chapter 8, we present a two-stage framework for estimating multiple sparse filters from stereo mixtures, which generalises the approach for single filter estimation presented in chapter 7. The filters corresponding to each source are estimated by two steps: (i) a time-frequency points clustering step which exploits the source sparsity to gather points in the mixture which satisfy the CR and (ii) a sparse filter recovery step which uses the cluster obtained in stage (i). The clustering stage enables us to formulate a SITO problem to estimate filters corresponding to each of the sources. The experimental evaluation focusses on the filter estimation stage of the framework, by assuming that the clustering stage is solved using side information about the sources or true mixing filters.

In the second part of chapter 8 we develop a second time-frequency domain version of the CR that is not based on the narrowband approximation. This form of CR is more accurate than the narrowband CR and we refer to this as the wideband CR. The two-stage framework for multiple filter estimation is tailored for wideband CR and experimental assessment of the filter estimation stage is presented.

The last part of chapter 8 focusses on a specific type of mixtures which involves only two sources: one source which is mixed with convolutive filters and the remaining source is mixed with anechoic filters. In such a setting, the time-frequency points clustering can be performed blindly and they can be used to estimate the filters. Experiments with real audio sources to demonstrate the blind clustering and filter estimation stages are presented.

In chapter 9, we conclude the thesis and we propose some perspectives for the improvement and generalisation of the filter estimation framework that is presented in Chapter 8, especially to make the framework completely blind in a generic setting.

1.6 Publications related to thesis contributions

1. S. Arberet, P. Sudhakar and R. Gribonval, A wideband doubly-sparse approach for MIMO sparse filter estimation, *Accepted for publication in ICASSP 2011, Prague, Czech Republic*.
2. P. Sudhakar, S. Arberet and R. Gribonval, Double Sparsity: Towards blind estimation of multiple channels, *In Proc. of Latent Variable Analysis and Signal Separation 2010, St. Malo, France*.

1.6 Publications related to thesis contributions

3. P. Sudhakar and R. Gribonval, Sparse filter models for solving permutation indeterminacy in convolutive blind source separation, *In SPARS'09, St. Malo, France.*
4. P. Sudhakar and R. Gribonval, A sparsity-based method to solve the permutation indeterminacy in frequency domain convolutive blind source separation, *In Proc. of Independent Component Analysis and Signal Separation 2009, Paraty, Brazil.*

1. INTRODUCTION

Part I

State of the art

Chapter 2

The source separation problem

In this chapter, we will describe the state of the art mixing filter estimation problem in the source separation setting.

2.1 Introduction

As discussed in the introduction chapter, a standard architecture for source separation is the two step approach: a first step where the mixing filters are estimated and a second step where the estimated mixing filters are used to estimate the sources. An advantage of having a two stage approach is that it offers modularity and we can mix and match any mixing filter estimation technique with any source estimation technique. However, some of the earlier approaches to source separation aimed at directly recovering the sources, without explicitly estimating the mixing filters.

Although we are primarily interested in the mixing parameter estimation problem in convolutive setting, we will also briefly describe the classical families of approaches where either the sources are estimated directly or the mixing filters and the sources are estimated jointly.

There has been a proliferation of literature on the source separation methods, and in this chapter we provide a brief overview of the underlying principles of different families of approaches. For a detailed survey of methods, please refer to the Handbook of Blind Source Separation [1]. However, our focus in this thesis is on the methods which rely on the sparsity hypothesis of the sources and mixing filters, we will deal with that in detail in chapter 3.

2.1.1 Trivial setting

In the case of linear instantaneous mixing process, the mixtures are nothing but a sum of scaled versions of the sources, and the mixing matrix \mathbf{A} becomes a usual two-dimensional matrix of size $M \times N$. So, in the absence of noise the mixing equation (1.3) can be written as:

$$\mathbf{X} = \mathbf{A}\mathbf{S}. \tag{2.1}$$

2. THE SOURCE SEPARATION PROBLEM

Suppose that the source matrix \mathbf{S} is known to us explicitly, then we can estimate the mixing matrix by simply inverting the source matrix. That is, the estimate $\tilde{\mathbf{A}}$ of the filter matrix is given by

$$\tilde{\mathbf{A}} = \mathbf{X}\mathbf{S}^\dagger,$$

where \mathbf{S}^\dagger is the pseudo-inverse of the matrix \mathbf{S} . However, in a blind setting where no explicit knowledge about the sources is available, the problem of mixing parameter estimation becomes very challenging. In such cases, we need additional prior knowledge on the sources in order to estimate the mixing parameters. In what follows, we describe different kinds of prior knowledge that is assumed about the sources and the methods that exploit them to estimate the mixing parameters.

2.1.2 Plan of the chapter

We present two grand families of approaches for source localisation and separation in the following sections, which are based on two different types of hypothesis of sources and mixing processes. Firstly, we describe the principle behind Independent Component Analysis (ICA) approach, where the sources are assumed to be statistically independent, and describe its application to source separation.

Then, we introduce Sparse Component Analysis (SCA) where the sources are assumed to be sparse in the time-frequency domain. We begin with the general concept of redundant representations of data and recollect the definition of sparsity and then introduce the idea of disjointness of the time-frequency support of the sources and then we describe the general principle of SCA.

Subsequently, a brief survey of methods which exploit the time-frequency domain disjointness of the sources to estimate mixing parameters in case of linear instantaneous and anechoic settings is presented. We then look at the convolutive setting and describe the difficulties involved in directly extending the methods used for anechoic case.

2.2 Independent Component Analysis (ICA)

The first significant breakthrough in source separation for instantaneous mixtures came with the advent of *Independent Component Analysis* (ICA). ICA was first proposed for the separation of instantaneous mixtures in the (over)-determined setting. In such a setting, ICA aims at directly extracting the source components \mathbf{S} from the mixtures \mathbf{X} .

2.2.1 Hypothesis and principle of ICA

The rows of the matrix \mathbf{S} contains individual sources and each source is regarded as a *component* of the mixture. The hypothesis on which ICA is based is that the sources \mathbf{S} that combine to yield the mixtures \mathbf{X} are assumed to be statistically independent of each other. That is, the joint probability distribution of the sources $P_{\mathbf{S}}(\mathbf{S})$ is assumed to be the product of the probability distribution of individual sources $P_{s_j}(s_j)$:

2.2 Independent Component Analysis (ICA)

$$P_{\mathbf{S}}(\mathbf{S}) = \prod_j P_{s_j}(s_j). \quad (2.2)$$

Therefore, the principle of ICA is to achieve separation of sources by maximising the independence between the estimated sources [12], or equivalently, minimising the dependency between the estimated sources. Hence the name independent component analysis.

In a determined setting, if \mathbf{A} is the $M \times M$ mixing matrix which is unknown but invertible, we seek another matrix $\tilde{\mathbf{B}}$, known as *un-mixing matrix* such that, the estimated sources

$$\tilde{\mathbf{S}} = \tilde{\mathbf{B}}\mathbf{X}$$

has independent components. That is, we are interested in a matrix $\tilde{\mathbf{B}}$ that maximises the statistical independence between the rows of the matrix $\mathbf{B}\mathbf{X}$. Mathematically, the matrix $\tilde{\mathbf{B}}$ is obtained by solving the following optimisation problem:

$$\tilde{\mathbf{B}} = \arg \max_{\mathbf{B}} I(\mathbf{B}\mathbf{X}), \quad (\text{ICA})$$

where $I(\mathbf{B}\mathbf{X})$ measures the statistical independence between the rows of the matrix $\mathbf{B}\mathbf{X}$.

The exact definition of *statistical independence* to be used and the specific algorithm to maximise the same vary on the context of usage, and hence there exists several variants of ICA.

One of the popular measures of dependence is the *mutual information*, which measures the similarity of statistical distributions in terms of the information measure. Given a matrix \mathbf{Y} whose rows y_j are random vectors, the mutual information between these vectors is defined as [13]:

$$I(\mathbf{Y}) := K(P_{\mathbf{Y}} | \prod_j P_{y_j}), \quad (2.3)$$

where $P_{\mathbf{Y}}$ is the joint probability distribution of the rows of the matrix \mathbf{Y} , P_{y_j} is the marginal distribution of the j^{th} row of \mathbf{Y} and $K(f|g)$ is the Kullback-Leibler divergence. The Kullback-Leibler divergence between two probability distributions f and g is defined by:

$$K(f|g) := \int f(x) \log \left(\frac{f(x)}{g(x)} \right) dx. \quad (2.4)$$

Mutual information is a non-negative quantity (but can be $+\infty$) and it vanishes only if the sources are mutually independent. It can be also written as:

$$I(\mathbf{Y}) = \sum_{j=1}^N H(y_j) - H(\mathbf{Y}), \quad (2.5)$$

2. THE SOURCE SEPARATION PROBLEM

where $H(\mathbf{Y})$ and $H(y_j)$ are the *joint entropy* and *marginal entropies*. The entropy¹ of a probability distribution f is defined by:

$$H(f(x)) := - \int f(x) \log f(x) dx.$$

Mutual information, defined by Eqs. (2.3) and (2.5) is a theoretical quantity as it involves the integrals with probability distributions and hence approximations are used in practice [14].

2.2.2 Algorithms for ICA

ICA, as described in the previous section, is a separation principle which is based on minimising an objective function. However, researchers have proposed various algorithms to minimise the criteria, depending on the kind of probability distributions that the sources admit. The pioneering algorithm for ICA was first proposed by Jutten and Héroult [15] based on the neuro-mimetic architecture. Bell and Sejnowski proposed an algorithm based on the Infomax principle of neural networks [16]. Other popular algorithm are the fastICA [17], algorithms proposed by Comon based on the higher order cumulants [14], and the algorithm based on the higher order statistics [18], the algorithm JADE [19] and other notable algorithms based on the joint diagonalisation principle proposed by Pham [20] and Cardoso [21].

2.2.3 Limitations of ICA

The primary limitation of ICA is the fact that it assumes the mixing matrix \mathbf{A} is invertible. Moreover, ICA inherently suffers from the *permutation* and *scaling* ambiguities, which are described next.

Permutation and scaling indeterminacy

In general, the definition of statistical independence is immune to the scaling of sources, and also the order of the sources. For example, from Eq. (2.5) it is evident that the mutual information is invariant to the order of the sources. Also, mutual information is invariant to scaling of sources [14]. Hence, ICA based techniques can estimate the sources only up to a scaling factor on each of the sources and a permutation of the sources. Hence, the un-mixing matrix obtained using ICA will be of the form

$$\tilde{\mathbf{B}}\mathbf{A} = \mathbf{\Lambda}\mathbf{P},$$

where $\mathbf{\Lambda}$ is a diagonal scaling matrix and \mathbf{P} is a permutation matrix².

The permutation and scaling indeterminacies are in fact inherent to the source separation problem itself. Hence, irrespective of the techniques used (ICA and many other techniques), the permutation and scaling ambiguities occur.

¹Technically, the definition is for *differential entropy* of continuous random variables, which is not the limit of Shannon's entropy for discrete random variables. However, it is called entropy in short.

²An $N \times N$ permutation matrix is an $N \times N$ Identity matrix, with permuted columns. For example,

2.2.4 ICA for convolutive mixtures

The application of ICA for source separation problem was first shown for instantaneous mixtures, but the principle of ICA has been applied to separating convolutive mixtures as well with adequate adaptation. In this section, we present a short overview of methods that use the independence prior of the source to achieve source separation in convolutive settings. A survey of methods for convolutive source separation can be found in [22].

2.2.4.1 Time domain methods

The first efforts on source separation of convolved mixtures were made in the time domain, mainly inspired by the *blind deconvolution* methods. Torkolla [23] modelled the unmixing procedure as FIR filtering of the mixtures themselves. That is,

$$\tilde{s}_j = \sum_i b_{ji} \star x_i, \quad (2.6)$$

where b_{ji} are the unmixing filters. In order to estimate them, he used an information maximisation approach as described in [16].

Lee et al. [24] looked at modelling the unmixing procedure as an IIR filter and derived a solution for this problem, noting that the recording environment had to be *minimum phase*. The time domain techniques for convolutive mixtures are computation intensive, and one way to deal with the problem is to work in the frequency domain.

2.2.4.2 Frequency domain methods

One of the standard ways of dealing with the convolutive problem is to transform the mixtures into the time-frequency domain using a suitable transform such as the Short Time Fourier Transform (STFT). Consider a discrete signal $s(t)$ with $0 \leq t < T$, and let $w(t)$ be a discrete window function whose support is $[-F/2, F/2]$ with unit ℓ^2 norm $\|w\|_2 = 1$. Then the STFT of $s(t)$ is defined as

$$\hat{s}(\tau, f) = \sum_{t=0}^{T-1} s(t)w(\tau - t)e^{-2i\pi ft}, \quad (2.7)$$

where f is the frequency index and τ is the frame index. The STFT coefficients are computed on a discrete grid specified by: $\tau = qF/2$, $q \in \mathbb{Z}$ and $f = l/F$, $0 \leq l \leq F/2$. STFT maps the time domain representation of a signal into a two dimensional time-frequency representation. The choices of the window function $w(t)$ and the window length F are dependent on the signal one wishes to analyse, and the specific properties one is looking for in the signal. Please refer to [25] for further details on the choices of the window function and their properties.

$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ is a 3×3 permutation matrix. The identity matrix $\mathbf{I}_{N \times N}$ is also a (trivial) permutation matrix.

2. THE SOURCE SEPARATION PROBLEM

As a result of STFT, the convolutions in the time-domain are now transformed into complex multiplications in the frequency domain. By narrowband approximation, we can write the time-frequency domain coefficients of the mixtures as a product-and-sum of the frequency domain filter coefficients and the time-frequency domain coefficients of the sources [26].

Let $\hat{x}_i(\tau, f)$ be the time-frequency domain representation of the i^{th} mixture at frame index τ and frequency f , then by assuming that the filters are fixed over time and by *narrowband approximation*, we have:

$$\hat{x}_i(\tau, f) \approx \sum_{j=1}^N \hat{a}_{ij}(f) \hat{s}_j(\tau, f), \quad 1 \leq i \leq M, \quad (2.8)$$

where $\hat{s}_j(\tau, f)$ is the time-frequency domain representation of the j^{th} source, and $\hat{a}_{ij}(f)$ is the frequency domain representation of the filter between j^{th} source and i^{th} sensor.

With this transformation, a single real-valued convolutive source separation problem has been converted into several instantaneous mixtures problem, possibly complex valued, in each frequency bin. Hence, ICA algorithms for instantaneous mixtures can be used to achieve separation in each frequency bin f . However in such a case, each separated frequency bin of the sources will be arbitrarily scaled and permuted due to the inherent nature of the ICA principle. That is, if $\hat{\mathbf{s}}(\tau, f)$ and $\tilde{\mathbf{s}}(\tau, f)$ are vectors containing the STFT coefficients of all the true and estimated sources at the time-frequency index (τ, f) , then they are related in the following way:

$$\tilde{\mathbf{s}}(\tau, f) = \mathbf{P}(f) \mathbf{\Lambda}(f) \hat{\mathbf{s}}(\tau, f) \quad (2.9)$$

where $\mathbf{P}(f)$ is a permutation matrix and $\mathbf{\Lambda}(f)$ is a diagonal scaling matrix, which are frequency dependent. The permutation and scaling ambiguities have to be corrected before transforming the estimated sources back into the time-domain.

Some of the frequency domain methods for convolutive source separation is based on the non-stationarity assumption of the sources. Let $S_{\mathbf{S}}(\tau, f) = \hat{\mathbf{s}}(\tau, f) \hat{\mathbf{s}}(\tau, f)^T$ be the covariance matrix of STFT of the source signals at the time-frequency index (τ, f) , then the covariance matrix of the STFT of the mixtures at the time-frequency index (τ, f) is given by

$$S_{\mathbf{X}}(\tau, f) = \hat{\mathbf{A}}(f) S_{\mathbf{S}}(\tau, f) \hat{\mathbf{A}}^*(f) \quad (2.10)$$

where $*$ denotes the transpose conjugated, and $\hat{\mathbf{A}}(f)$ is the matrix of DFT coefficients of the filters at frequency f .

Pham et al. [27] proposed the idea of finding another set of matrices $\mathbf{B}(f)$ such that the covariance matrices $\mathbf{B}(f) \tilde{S}_{\mathbf{X}}(\tau, f) \mathbf{B}^*(f)$ at different time points τ are as close to diagonal as possible, where $\tilde{S}_{\mathbf{X}}(\tau, f)$ are estimates of $S_{\mathbf{X}}(\tau, f)$. For a single matrix (at a given τ and f), the diagonality measure is given by

$$\frac{1}{2} \left\{ \log \det \text{diag} \left[\mathbf{B}(f) \tilde{S}_{\mathbf{X}}(\tau, f) \mathbf{B}^*(f) \right] - \log \det \left[\mathbf{B}(f) S_{\mathbf{X}}(\tau, f) \mathbf{B}^*(f) \right] \right\}, \quad (2.11)$$

where $\text{diag}(\cdot)$ denotes the operator which builds a diagonal matrix from its argument.

Parra and Spence also use a similar kind of argument for diagonalisation of the time varying spectra of the observed mixtures, and they have come up with their own version of the diagonality measure [28].

Mitianoudis and Davies [29] propose a fast frequency domain ICA framework for the separation of convolutive audio mixtures. They assume frequency independent priors for the unmixing matrix $\tilde{\mathbf{B}}(f)$ and the separated sources $\tilde{\mathbf{s}}(\tau, f)$, the unmixing matrix $\tilde{\mathbf{B}}(f)$ is found for each frequency bin by maximising the following log-likelihood:

$$\log P(\hat{\mathbf{x}}(\tau, f) | \tilde{\mathbf{B}}(f)) = \mathbb{E} \left\{ \log P(\tilde{\mathbf{B}}(f) \hat{\mathbf{x}}(\tau, f)) \right\} + \log \det \tilde{\mathbf{B}}(f). \quad (2.12)$$

They also alter the natural gradient algorithm [30] to incorporate some information about the scaling of the signal with time, and hence they have been effectively able to handle the permutation problem as well simultaneously.

2.3 Sparse Component Analysis (SCA)

The structure of this section largely follows chapter 10 of [1] by Gribonval and Zibulevsky.

In Sec. 1.2.4, we introduced the intuitive notion of sparse sources and illustrated it with an example. Sparsity of signals such as audio and images has been exploited for a long time for coding or compression purposes. However, sparsity as a tool for source separation has been in use only in the last decade or so [31, 6, 4, 5, 32].

Unlike in the ICA setting, where the sources were directly estimated from the mixtures using the independence assumption, most of the methods based on source sparsity performs source separation by first estimating the mixing matrix, and then subsequently the sources. The sparsity hypothesis plays a role in both the stages of source localisation and source estimation. Sparsity hypothesis enables us to estimate the mixing matrix and perform source estimation even in the underdetermined mixtures cases.

Hypothesis and principle of SCA

As the name suggests, the approaches that constitute sparse component analysis assume that the sources are sparse. The sparsity hypothesis means that a lot of source coefficients are zero or close to zero. The identification of the mixing matrix \mathbf{X} and the source matrix \mathbf{S} using SCA is unique up to a permutation and a scaling factor, under certain conditions on the structure of \mathbf{A} and the sparsity of the sources \mathbf{S} [33]. This is true even for the underdetermined cases and hence SCA makes an attractive tool to deal with underdetermined problems.

In the simplest case of instantaneous mixtures, *when the mixing matrix is known*, the principle of source estimation is the following: estimate the sources by maximising the sparsity of the sources under the equality condition $\mathbf{X} = \mathbf{AS}$.

The sparsity of a given vector $\mathbf{s} = [s(0), \dots, s(T-1)]^T$ is measured by the number

2. THE SOURCE SEPARATION PROBLEM

of non-zero coefficients in the vector, which is called the ℓ^0 norm¹, defined as:

$$\|\mathbf{s}\|_0 := \#\{t, s(t) \neq 0\} = \sum_t |s(t)|^0. \quad (2.13)$$

Hence, by minimising the ℓ^0 norm of the sources under the condition $\mathbf{X} = \mathbf{A}\mathbf{S}$, one estimates the sources in SCA. However, the ℓ^0 norm is a discrete norm and its minimisation is combinatorial in nature and hence it is a NP-hard problem [34]. So in practice, relaxed versions of the ℓ^0 norm, such as the ℓ^p norm is generally used for this purpose. For $p > 0$, the ℓ^p norm of a vector \mathbf{s} is defined by

$$\|\mathbf{s}\|_p := \left[\sum_t |s(t)|^p \right]^{\frac{1}{p}}, \quad (2.14)$$

with the convention that $\|\mathbf{s}\|_\infty := \max_t |s(t)|$. Hence, to estimate the source, one solves the following optimisation problem in SCA:

$$\min_{\mathbf{s}} \|\mathbf{s}\|_p \quad \text{subject to} \quad \mathbf{X} = \mathbf{A}\mathbf{S}. \quad (\text{SCA})$$

For values of $p \geq 1$, the ℓ^p norm is convex and hence the optimisation problem is also convex. If $p < 1$ and $p \rightarrow 0$, the norm is closer to the true definition of sparsity, the ℓ^0 norm, but it is no more convex. Hence, the ℓ^1 norm is a norm that is closest to the ℓ^0 norm and also convex. Due to the ease in designing algorithms for convex problems, ℓ^1 norm is widely used when looking for sparse solutions.

Generally, source signals are not sparse in the time-domain, as it was assumed while introducing the concept of SCA in the beginning of this section. However, signals might be sparse in a different domain such as the time-frequency domain, wavelets, etc. Hence, most often the sources are not estimated directly in the time-domain but they are estimated in the transform domain in which the sources are sparse.

Further, the sources are often assumed to have disjoint or near disjoint supports in the transform domain. This hypothesis will also aid the estimation of the mixing parameters \mathbf{A} from the mixtures \mathbf{X} .

Hence, SCA is composed of the following steps:

1. Transformation of the mixtures \mathbf{X} into a domain in which the sources are assumed to be sparse, to yield $\widehat{\mathbf{X}}$;
2. Estimation of the mixing parameters $\widetilde{\mathbf{A}}$ from $\widehat{\mathbf{X}}$;
3. Estimation of sources in the transform domain by solving the problem in Eq. (SCA);

¹Technically, this is not a norm but a quasi-norm.

2.3.1 Sparse representations

In Sec. 1.2.4 we introduced the concept of time-frequency domain sparsity of sources informally with an illustration and in Sec. 2.3 we formalised the notion of source sparsity and briefly described the principle of sparse component analysis (SCA) for source localisation and separation. During the formalisation of the notion of source sparsity in Eqs. (2.13) and (2.14) and the description of sparse component analysis in (SCA), we inherently assumed that the sources are sparse in the time-domain.

While the assumption of source sparsity in the time-domain was sufficient to describe the principle of SCA, it is not often true that the sources are sparse in the time-domain representation. However, it might be possible to change the representation of the source signals such that it is sparse in the new representation. Such representations which allow signals to be sparse are called *sparse representations*. The role of sparse representations span much beyond sparse component analysis, and the problem of finding sparse solutions to underdetermined systems can be generalised with sparse representations [35].

The idea behind sparse representations of signals is this: Any general signal vector \mathbf{s} can be (approximately) represented by a linear combination of few elementary signals φ_n , possibly complex valued, which are called *atoms*. Though we retained the notation of a source signal to describe sparse representations, this holds good for any arbitrary signal. Therefore, mathematically this means that a signal vector $\mathbf{s} = [s(1), \dots, s(T)]^T$ can be written as:

$$\mathbf{s} = \sum_n c_n \varphi_n, \quad (2.15)$$

where c_n is the coefficient corresponding to φ_n . Such a collection of atoms is known as a *dictionary*, and if there are a finite number of atoms then we can collect all the atoms as the columns of a matrix Φ , called a *dictionary matrix*, and the coefficients in a vector \mathbf{c} called the *coefficient vector*. These atoms are usually normalised to have unit ℓ^2 norm. Then, Eq. (2.15) can be written as $\mathbf{s} = \Phi \mathbf{c}$.

If the number of atoms in the dictionary is a finite number N_ϕ then $\Phi \in \mathbb{C}^{T \times N_\phi}$ and $\mathbf{c} \in \mathbb{C}^{N_\phi \times 1}$.

We say that a source signal \mathbf{s} admits a sparse representation with respect to a particular dictionary Φ if the corresponding coefficient vector \mathbf{c} has only a few non-zero coefficients compared to its size. Hence, the concept of sparsity of source signals is generic and the sparsity of source signals is now measured using the coefficient vector \mathbf{c} instead of \mathbf{s} directly. The measures of sparsity introduced in Sec. 2.3, namely the family of ℓ^p ($p \geq 0$) norms still hold good on the coefficient vector \mathbf{c} .

With $\Phi = \mathbf{I}_{T \times T}$, the atoms are simply the Dirac functions $\varphi_n(t) = \delta(t - n)$, $1 \leq n \leq T$ and we return back to the notion of sparsity of signals in the time-domain. Van Hulle [36] considered this simple model of sparsity in time domain for separation of speech signals.

The dictionary Φ could also simply be a basis for the space \mathbb{R}^T , like the Discrete Fourier Transform (DFT) or Modified Discrete Cosine Transform (MDCT), in which case the dictionary is square. A basis spans the entire space and hence it is referred

2. THE SOURCE SEPARATION PROBLEM

to as a *complete dictionary*. It should be noted that in the case where the dictionary is a basis, the representation of a signal vector is unique. But, signals often might not be sparsely representable in a basis and hence we need to look for *redundant* or *overcomplete* dictionaries which allow sparse representation. This could be as simple as a concatenation of bases [37, 38, 39] to specialized construction of dictionaries like curvelets, ridgelets, noiselets, etc.

Most of the research effort in sparse signal processing is related to the study of sparse representations of general classes of signals in given a dictionary, and the design and analysis of algorithms for performing signal processing tasks exploiting sparsity. Another field of research in applied and computational harmonic analysis is to build or find dictionaries which allow sparse representations of a given class of signals. Recent research efforts by the community is focussed on finding techniques to learn dictionaries from training samples for specific class of signals and this is referred to as *dictionary learning* [35, 40, 41, 42].

Sparse representations have been quite extensively used in general audio and music signal processing, from coding to source separation [3]. We shall now turn our focus onto audio signals, their sparse representation and its relevance to source localisation and separation.

2.3.1.1 Sparsity of audio signals

Audio signals are known to be sparse in the Short-Time Fourier Transform (STFT) domain.

Let us recollect the example of time-frequency domain sparse signals. Fig. 2.1(a) shows the time domain plot of a flute signal. Note that the signal is active at almost all the times and hence the time domain sparsity assumption of Van Hulle is violated. Fig. 2.1(b) shows the magnitude of the STFT of the same signal with a frame size $F = 2048$ samples and *sine* window. We can clearly observe the sparse behaviour of the source. The frequency content of the source, although varying over time, is limited to only a few bins in each STFT frame. This property of sparseness of audio signals is central to a family of techniques developed for estimating the mixing parameters in linear-instantaneous and anechoic mixtures, and also the subsequent estimation of source signals.

2.3.1.2 Disjointness of the time-frequency supports

A closely related concept to sparsity is the disjointness of the supports of the sources in the time-frequency domain. Intuitively, the sources are said to be disjoint in the time-frequency domain if their time-frequency supports do not overlap with each other.

Let us illustrate the idea of time-frequency disjointness with an example. Fig. 2.2 shows the the STFT magnitude plot of two sources: Fig. 2.2(a) is a flute signal and Fig. 2.2(b) is a guitar signal. The dark regions in the plots indicate the presence of the source and it is visually noticeable that the sources are active at disjoint time-frequency locations when the frame index $\tau \leq 4000$, after which they overlap.

2.3 Sparse Component Analysis (SCA)




Figure 2.1: An example of time-frequency domain sparse signal

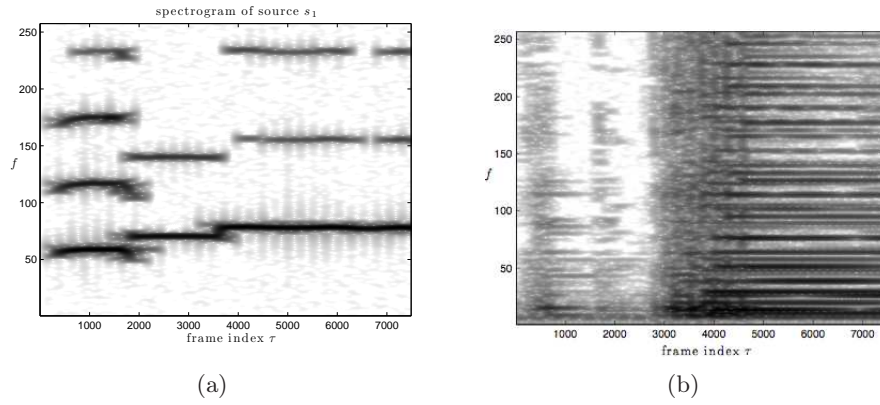


Figure 2.2: STFT of the two sources: (a) source s_1 is a flute sound and (b) source s_2 is a guitar sound.

Jourjine et al. [4] formalised the notion of time-frequency disjointness and introduced the concept of *w-disjoint orthogonality* of source signals. Two sources s_{j_1} and s_{j_2} are said to be *w-disjoint orthogonal*, if their corresponding STFTs obtained with the window function $w(t)$ have the following property:

$$\hat{s}_{j_1}(\tau, f)\hat{s}_{j_2}(\tau, f) = 0, \forall j_1 \neq j_2, \forall \tau, f. \quad (2.16)$$

Many a times, it is not possible for the sources to satisfy this *w-disjoint orthogonality* in the true sense. Yilmaz et al. [32] introduced the concept of approximate *w-disjoint orthogonality* and used it to achieve the separation of speech signals in anechoic settings. We shall now describe the general principle behind using the concepts of source sparsity and disjointness for mixing parameter estimation.

2. THE SOURCE SEPARATION PROBLEM

2.3.1.3 General principle of mixing parameter estimation

Let us begin our discussion with the simplest case of linear instantaneous mixtures in the absence of noise. The model that we are interested in is $\mathbf{X} = \mathbf{A}\mathbf{S}$. Let $\hat{x}_i(\tau, f)$ and $\hat{s}_j(\tau, f)$ be the STFT coefficients of the i^{th} mixture and j^{th} source respectively, at the time-frequency point (τ, f) obtained using (2.7) with a window function $w(t)$. Let a_{ij} be the mixing parameter between the i^{th} mixture and j^{th} source. Then we have the following relation:

$$\hat{x}_i(\tau, f) = \sum_j a_{ij} \cdot \hat{s}_j(\tau, f). \quad (2.17)$$

Let us assume that the sources are w -disjoint orthogonal and hence the supports are non-overlapping. Let Ω_j be the set of all time-frequency points where only the j^{th} source is active. Then, by (2.17) we have

$$\hat{x}_i(\tau, f) = a_{ij} \cdot \hat{s}_j(\tau, f), \quad \forall (\tau, f) \in \Omega_j. \quad (2.18)$$

It is evident from Eq. (2.18) that the mixtures at the time-frequency points Ω_j are nothing but scaled version of the mixing parameter a_{ij} , appropriately scaled by the factor $\hat{s}_j(\tau, f)$. If we collect all the M mixtures at a time-frequency point (τ, f) in a vector $\hat{\mathbf{x}}(\tau, f)$, then we can write Eq. (2.18) as

$$\hat{\mathbf{x}}(\tau, f) = \hat{s}_j(\tau, f) \cdot \mathbf{A}_j, \quad \forall (\tau, f) \in \Omega_j, \quad (2.19)$$

where $\mathbf{A}_j = [a_{1j}, \dots, a_{Mj}]^T$ is the j^{th} column of the matrix \mathbf{A} . In other words, we can say that the set of points $\{\hat{\mathbf{x}}(\tau, f) \in \mathbb{C}^M, (\tau, f) \in \Omega_j\}$ is aligned along the straight line passing through the origin and directed by the vector $\mathbf{A}_j \in \mathbb{C}^M$.

While it is hard to visualise the above statement in a general M dimensional setting, we can visualise the same easily when we have *stereo* mixtures, that is when $M = 2$, with the help of *scatter plots*. A scatter plot is basically a 2-dimensional plot of ordered pairs of real numbers. The STFT coefficients of the mixtures are complex numbers, and hence we plot either the real part $\Re \hat{\mathbf{x}}(\tau, f)$ or the imaginary part $\Im \hat{\mathbf{x}}(\tau, f)$ of the vector $\hat{\mathbf{x}}(\tau, f)$.

Fig. 2.3(a) shows the time-domain plots of three audio signals, and Fig. 2.3(b) shows the time domain plots of the stereo mixtures of the sources. Fig. 2.3(c) shows the STFT magnitude plots of the mixtures and finally Fig. 2.3(d) shows the scatter plot of the real part of the STFT of stereo mixtures. Notice that the points in the scatter plot tend to *cluster* around the vectors \mathbf{A}_j which are aligned towards the columns of the mixing matrix.

This geometrical structure of the scatter plots can be exploited to obtain the estimates $\hat{\mathbf{A}}_j$ of the columns of the mixing matrix. Clustering algorithms [43] can be used to cluster the data points and estimate the columns of the mixing matrix. These geometrical methods can be cleverly extended to the anechoic mixtures setting as well.

Fig. 2.4 depicts a typical mixing parameter estimation system that is primarily based on the source sparsity and disjointness assumption. The mixtures are transformed into

2.3 Sparse Component Analysis (SCA)

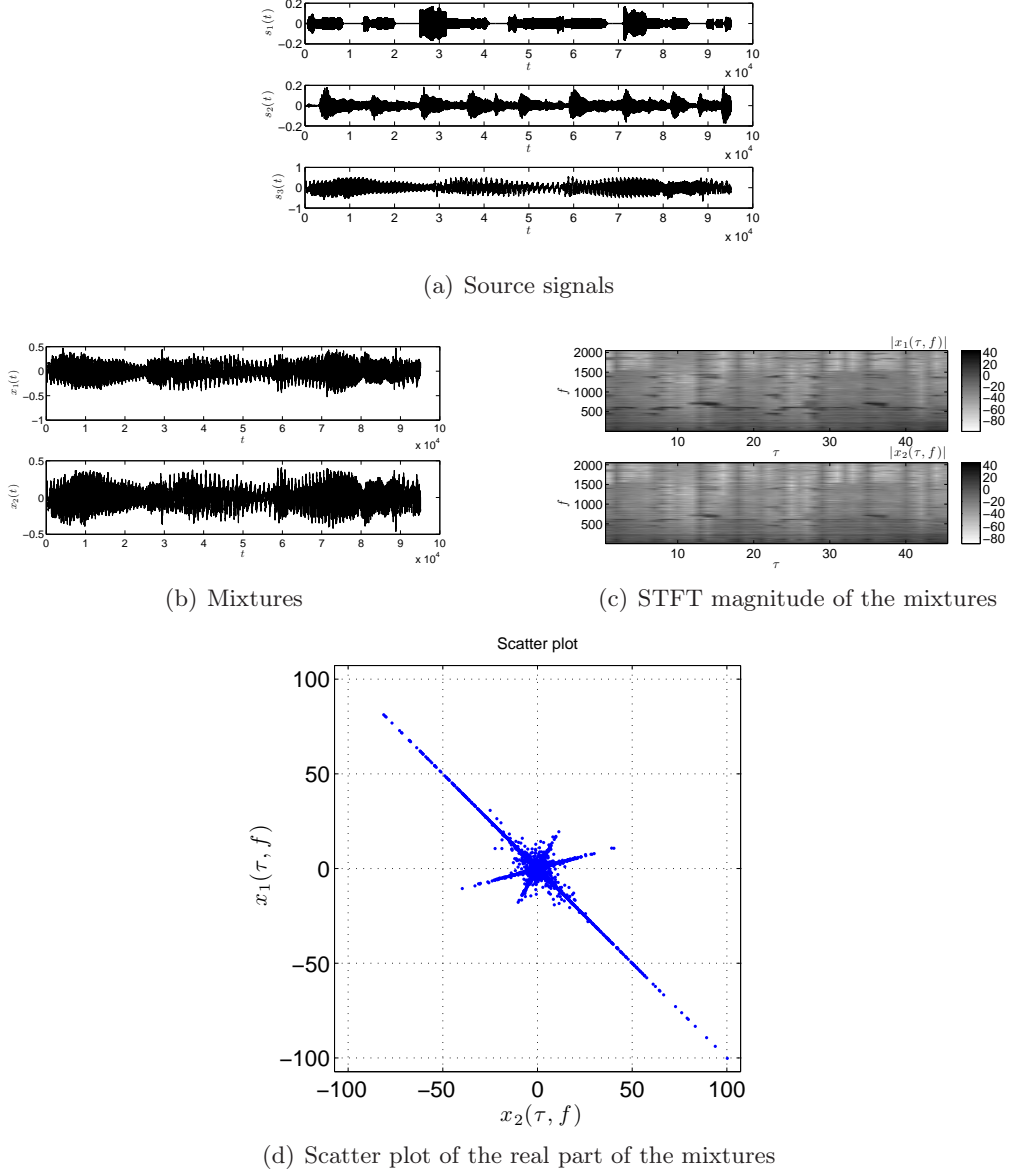


Figure 2.3: Geometric structure of scatter plot

a domain that admits sparse representation of the sources, the mixing parameters are then estimated using a clustering approach. Sparsity based methods are proven to be very effective for the undetermined mixtures setting. In the subsequent sections, let us review some of the key approaches based on the source sparsity hypothesis for linear instantaneous and anechoic mixtures setting.

2. THE SOURCE SEPARATION PROBLEM

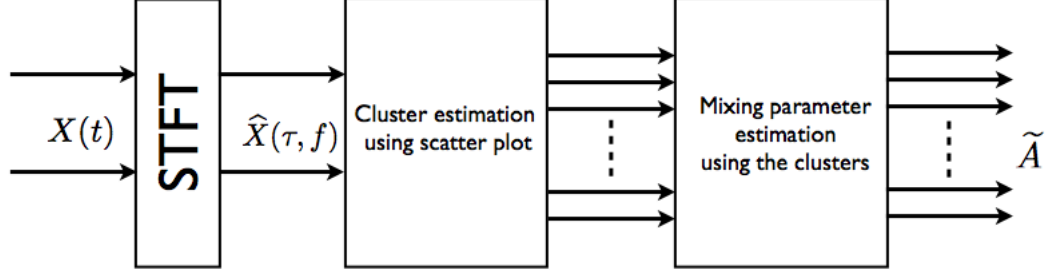


Figure 2.4: Workflow of a typical mixing parameter estimation system based on the scatter plot of the mixtures.

2.3.2 Linear instantaneous mixtures

In this section we shall review some approaches that are based on clustering of time-frequency points to estimate the mixing parameters in the linear instantaneous setting. Let us consider the linear instantaneous mixture model without noise: $\mathbf{X} = \mathbf{A}\mathbf{S}$. We wish to estimate the number of sources N , the mixing matrix \mathbf{A} from the mixtures \mathbf{X} . Since we neither know \mathbf{A} nor \mathbf{S} , the identification of \mathbf{A} is possible only up to a permutation and a scale factor [44]. However, barring the permutation, what we are really interested in is the relative amplitudes of the mixing parameters of a given source across all the channels. Hence, without loss of generality we can assume the mixing parameters across the channels to be normalized: $\sum_i a_{ij}^2 = 1$ and $a_{1j} \geq 0$ for $1 \leq j \leq N$.

In the stereo mixtures case ($M = 2$), each column \mathbf{A}_j of \mathbf{A} is a two dimensional vector, and they are assumed to be normalised having unit ℓ^2 norm. Hence, we can parameterize each column \mathbf{A}_j by an angle θ_j , that is, the columns can be written as:

$$\mathbf{A}_j = \begin{bmatrix} \cos \theta_j \\ \sin \theta_j \end{bmatrix} \in \mathbb{R}^2. \quad (2.20)$$

The parameter $\tan \theta_j$ is called the *intensity difference* and the parameter $\theta_j \in]-\pi/2, \pi/2]$ is called the *intensity parameter*. Hence, in the case of linear instantaneous stereo mixtures, all that is needed to characterise the mixing parameters are the intensity parameter θ_j for all $1 \leq j \leq N$.

2.3.2.1 Estimation of intensity parameter using global scatter plots

Suppose the sources are w -disjoint orthogonal, then from Eq. (2.19) it is clear that at each time-frequency point (τ, f) the following is valid

$$\forall(\tau, f) \in \Omega_j, \quad \begin{cases} \Re \hat{\mathbf{x}}(\tau, f) = \Re \hat{s}_j(\tau, f) \cdot \mathbf{A}_j \\ \Im \hat{\mathbf{x}}(\tau, f) = \Im \hat{s}_j(\tau, f) \cdot \mathbf{A}_j, \end{cases} \quad (2.21)$$

and hence we have the following relation

$$\forall(\tau, f) \in \Omega_j, \quad \frac{\Re x_2(\tau, f)}{\Re x_1(\tau, f)} = \frac{\Im x_2(\tau, f)}{\Im x_1(\tau, f)} = \frac{x_2(\tau, f)}{x_1(\tau, f)} = \tan \theta_j. \quad (2.22)$$

for some source index j .

Hence, at each time-frequency point (τ, f) , the ratio in Eq. (2.22) provides an estimate $\theta_{j(\tau, f)}$ of the intensity parameter at the time-frequency point (τ, f) , for some source j . Now, all that remains is to combine the information about these angles from all the time-frequency points and arrive at the global estimates of the intensity parameter $\tilde{\theta}_j$ for each source j .

Global scatter plot methods use all the time-frequency points simultaneously to arrive at the estimates. A naive approach is to calculate a histogram of the angles $\theta_{j(\tau, f)}$ and detect the peaks of the histogram to infer the estimates θ_j . The scatter plot contains a number of points of low magnitude, and hence these points might not well represent the column directions of the mixing matrix, making the histogram a poor choice for detecting the peaks.

A slightly sophisticated way to detect the angles is to use weighted histograms. A weighted histogram can be calculated in an efficient manner and may be smoothed with "potential functions", also called Parzen windows [5]. The weighting aims to obtain a histogram that depends on a small number of points with large amplitudes rather than a majority of points with negligible amplitudes, making the histogram more viable to detect peaks.

2.3.2.2 Estimation of intensity parameter using local scatter plots

Though the global scatter plots display geometric patterns which suggest the alignment of the points along the columns of the mixing matrix, the most traditional clustering approaches such as the K-means [43] might struggle to estimate the directions. While K-means algorithm is easy to implement, the performance of the algorithm depends on the proper initialisation and the pre-determination of the number of clusters. The usage of global scatter plots become much more difficult in the presence of noise, or when the w -disjoint orthogonality is not completely satisfied.

While the global scatter plot might fail to display any alignment of points, we can detect smaller regions of time-frequency points where the sources are disjoint and hence the local scatter plots display distinct directions. Abrard et al. [45], in their approach known as TIFROM, have proposed to use the local scatter plots to calculate the variance of the ratio defined in Eq. (2.22), in the local regions. In these regions, if there is only one dominant source then the variance will be close to zero and the mean of the ratio will be a good estimate of the intensity difference $\tan \tilde{\theta}_j$. Hence, the core step in TIFROM

2. THE SOURCE SEPARATION PROBLEM

is to identify time-frequency regions with small variance of the ratio, and then using the mean of the ratio in that region to obtain an estimate of the angle.

A variant of TIFROM, called TIFCORR was proposed by Deville et al. [46] which uses a different approach to obtain the regions in which only one source is dominant. However, the basic principle of estimating the angles from the local scatter plots remains the same. A similar approach consists in performing Principal Component Analysis (PCA) of the local scatter plots and selecting the regions whose main direction is the most dominant relative to others [47].

2.3.3 Anechoic mixtures

In this section, let us see how the scatter plot based methods can be extended to mixing parameter estimation in anechoic mixtures setting. The anechoic mixtures in the absence of noise can be modelled by

$$s_i(t) = \sum_{j=1}^N a_{ij} s_j(t - \delta_{ij}). \quad (2.23)$$

The only difference between the linear instantaneous model and anechoic model is the introduction of the delay δ_{ij} . The delay δ_{ij} represents the time interval between the origin of the sound by the j^{th} source and the observation of the same by the i^{th} microphone. Here again, we are not interested in the absolute delays, but only the relative delays between the channels and hence we can assume that $\delta_{1j} = 0$, $1 \leq j \leq N$.

On the same lines as the instantaneous mixtures, we can also assume that the coefficients a_{ij} are normalised: $\sum_i a_{ij}^2 = 1$ and $a_{1j} \geq 0$ for $1 \leq j \leq N$.

In the STFT domain, we have the relation (2.19) still valid even in the case of anechoic mixtures, but with the only difference that the columns \mathbf{A}_j of the matrix \mathbf{A} have the following structure: $\mathbf{A}_j = [a_{1j}, a_{2j}e^{-2i\pi\delta_{2j}f}, \dots, a_{Mj}e^{-2i\pi\delta_{Mj}f}]^T$.

In case of stereo mixtures ($M = 2$), each column \mathbf{A}_j now can be written as:

$$\mathbf{A}_j = \begin{bmatrix} \cos \theta_j \\ \sin \theta_j \cdot e^{-2i\pi\delta_j f} \end{bmatrix} \in \mathbb{C}^2. \quad (2.24)$$

The parameter $\tan \theta_j$ is called the *intensity difference* and the parameter $\theta_j \in] -\pi/2, \pi/2]$ is called the *intensity parameter*. The parameter $\delta_j \in \mathbb{R}$ is the delay associated with the j^{th} source between the channels. Hence, in the case of anechoic mixtures, all that is needed to characterise the mixing parameters are the intensity parameter and the delay pairs (θ_j, δ_j) for all $1 \leq j \leq N$.

If the sources satisfy the w -disjoint orthogonal hypothesis, then the following ratio

$$R_{21}(\tau, f) := \frac{\hat{x}_2(\tau, f)}{\hat{x}_1(\tau, f)}, \quad \hat{x}_1(\tau, f) \neq 0 \quad (2.25)$$

provides the necessary information about the parameters (θ_j, δ_j) . Specifically, the relationship will be:

$$R_{21}(\tau, f) = \tan \theta_j \cdot e^{-2i\pi\delta_j f}, \quad (2.26)$$

for some source j . DUET [4, 32] proposes to extract the intensity parameter and delay parameter using the following relationships:

$$\theta_j := \tan^{-1} |R_{21}(\tau, f)|, \quad (2.27)$$

$$\delta_j := -\frac{1}{2\pi f} \angle R_{21}(\tau, f), \quad (2.28)$$

where $\angle z \in]-\pi, \pi]$ is the phase of the complex number z . Hence, for each time frequency point (τ, f) , an estimate $\tilde{\theta}_j(\tau, f)$ of the intensity difference and an estimate $\tilde{\delta}_j(\tau, f)$ of the delay can be found for some unknown source j . Using these local estimates, a clustering approach can be used to estimate the global intensity difference and delay parameters $(\tilde{\theta}_j, \tilde{\delta}_j)$ for $1 \leq j \leq N$. DUET uses a two dimensional potential function (one for the intensity difference and another for the delay) and a K-means clustering algorithm is used to identify the peaks of the potential function.

In the next section, we will have a look at the convolutive mixtures setting and the challenges it poses for clustering based mixing parameter estimation technique.

2.3.4 Convolutive mixtures

Sparsity of sources has been explored to design systems to perform source localisation and separation in the convolutive settings as well. In this section, we would like to provide an overview of a few such approaches.

As a first approach, the convolutive mixtures can be transformed into complex instantaneous mixtures using the narrowband approximation assuming that the mixing filters are shorter compared to the STFT window size, as it was introduced in Sec. 2.2.4.2. Then for each frequency sub-band f , we can estimate the mixing parameters independently by clustering of the mixtures in each sub-band, as it is done for instantaneous mixtures that was described in the previous sections. However, with this scheme, we run into the problem of arbitrary scaling and permutations in each sub-band. That is, in each sub-band, the order of the filter estimates corresponding to the sources are arbitrary and hence the sub-bands will not be aligned in order to recover back the time-domain mixing filters. Hence, the permutation and scaling problems have to be solved before obtaining the time-domain filters.

Fig. 2.5 shows the workflow of a typical mixing filter estimation system based on narrowband approximation, which exploits source sparsity to estimate the mixing parameters in each sub-band. A typical system consists of the following steps:

1. Transform the mixtures into the time-frequency domain.
2. For each frequency sub-band f , obtain an estimate of the mixing filters by using the scatter plots of the mixtures. Let us denote the output of this stage by $\hat{\mathbf{A}}$. This new notation indicates that the sub-bands are arbitrarily permuted and scaled.

2. THE SOURCE SEPARATION PROBLEM

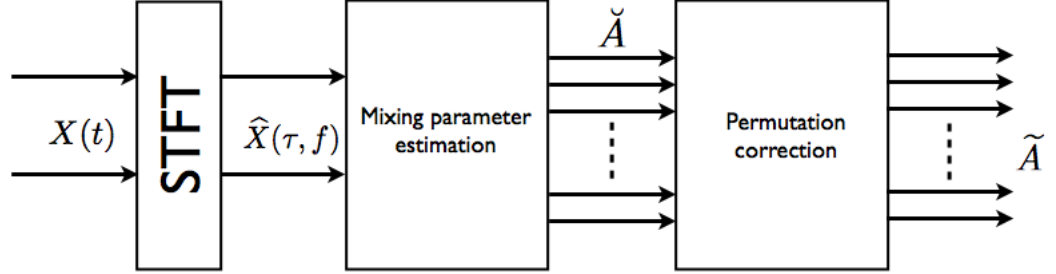


Figure 2.5: Workflow of a typical frequency domain based mixing parameter estimation for convolutive mixtures.

3. Solve the permutation and scaling problem to obtain the filter estimates $\tilde{\mathbf{A}}$.

The contribution of this thesis is to propose a framework to estimate the mixing filters from the mixtures by exploiting the time-frequency domain sparsity of the sources and the time-domain sparsity of the mixtures.

2.3.4.1 Brief review of existing methods for blind estimation of $\tilde{\mathbf{S}}$

Several methods for source separation in the convolutive setting, which exploit source sparsity, has been proposed based on the architecture presented in Fig. 2.5. However, as the goal of a source separation system is to estimate the sources, most of the methods which rely on this clustering approach directly estimate the sources by using a *time-frequency mask* obtained using the clusters, without going through the intermediate step of filter estimation. Time-frequency masks can just either be indicator functions over the entire time-frequency plane under consideration, which indicate the membership of time-frequency points to different sources, or they can be probabilistic masks allowing soft classification of points.

The central problems that are studied in the case of convolutive source separation based on narrowband approximation 2.8 and source sparsity are clustering and permutation correction.

Reju et al. [48] have recently proposed a deterministic method to cluster the time-frequency points of the mixtures based on the Hermitian angles between the mixture vectors. This approach is generic for mixtures with more than two components. Hermitian angles between the mixture vectors and a reference vector are analysed and they are clustered using k -means or a fuzzy c -means algorithm. Subsequently, the masks are obtained using the clusters, and the permutations are aligned using the mask functions rather than the time-frequency points themselves.

Sawada et al. [49, 50] have proposed a probabilistic approach for clustering of time-frequency points in the mixtures, and the subsequent permutation correction. The mixture vectors are modelled using a Gaussian mixture model, and an Expectation-Maximisation (EM) algorithm is then used to estimate the model parameters and the

classes to which the mixtures belong to. Mandel et al. [51] also have proposed a clustering scheme based on the probability distributions of the interaural phase and level differences in a stereo convolutive setting.

2.3.4.2 Estimation of $\tilde{\mathbf{S}}$ when \mathbf{A} is known

As an alternative to the narrowband formulation, the sparsity of sources can be exploited to estimate them directly in the time domain itself without resorting to the narrowband approximation, thus avoiding the problems of permutation and scaling. While this task is difficult in general, Kowalski et al. [2] have reported an interesting work regarding sparse source estimation using time-domain mixing equation and by assuming that the mixing filters are known (hence, it is not a blind setting). The authors refer to this approach to as *wideband formulation* because it directly uses the time-domain equations.

Let us describe the wideband formulation informally and for the sake of simplicity, let us use the following simplified, albeit abusive, notation. Let $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_N]$ be the matrix of sources, and let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$ be the matrix of mixtures. Let us use the shorthand notation $\mathbf{X} = \mathbf{A} \star \mathbf{S}$ to indicate the convolutive mixing process without noise, where \mathbf{A} is the mixing filter matrix.

Let the sources admit sparse representations in a dictionary $\Phi \in \mathbb{C}^T$ such that $\mathbf{S} = \Phi \mathbf{C}$, where each column of \mathbf{C} contains the vector of sparse representation coefficients for each source. Hence we have

$$\mathbf{X} = \mathbf{A} \star (\Phi \mathbf{C}). \quad (2.29)$$

Now, if the filter matrix \mathbf{A} is known to us, and if we assume that the coefficients \mathbf{C} are sparse, then we can use Eq. (2.30) to formulate an ℓ^1 minimisation problem in order to estimate the source coefficients. That is, one can solve the following optimisation problem in order to estimate the source coefficients $\tilde{\mathbf{C}}$:

$$\min_{\tilde{\mathbf{C}}} \lambda \|\tilde{\mathbf{C}}\|_1 + \frac{1}{2} \|\mathbf{X} - \mathbf{A} \star (\Phi \tilde{\mathbf{C}})\|_2^2. \quad (2.30)$$

Instead of plainly minimising the ℓ^1 norm of the sources, the authors in [2] propose to use specialised norms, called *mixed norms*, that promote certain specific structures on the sparse signals that one is looking for.

2.4 Summary

In this chapter, we provided an overview of two families approaches for solving the source separation problem, namely independent component analysis (ICA) and sparse component analysis (SCA).

ICA assumes that the sources which combine to form the mixtures are statistically independent and hence the objective of ICA is to separate them by maximising their independence. SCA is based on the hypothesis that the sources are sparse in the time-frequency domain. Source separation is generally performed by first estimating the

2. THE SOURCE SEPARATION PROBLEM

mixing parameters and subsequently estimating the sources by solving an optimisation problem that maximises the sparsity of the sources.

We described a general principle behind mixing parameter estimation based on the scatter plots of the mixtures in the stereo setting. This approach consists of a clustering step and a parameter estimation step and we described a few important clustering based parameter estimation techniques for linear-instantaneous and anechoic mixtures.

Parameter estimation techniques for convolutive cases based on the clustering approach suffer from permutation and scaling ambiguities, and these have to be corrected. We briefly described a few methods which rely on source sparsity for parameter estimation and have permutation correction stage in their workflow. Lastly, a source estimation method in a non-blind setting which exploits source sparsity is described.

In the next chapter we introduce the problem of finding solutions to underdetermined systems of linear equations and the role of sparsity in such tasks.

Chapter 3

Solutions of underdetermined linear systems

This chapter is dedicated to an introduction to the problem of finding sparse solutions to underdetermined systems of linear equations.

3.1 Introduction

Finding solutions to systems of linear equations is a central problem in numerical linear algebra, which in turn forms the basis for most of the engineering sciences. In a standard setting, we would like to solve a system of linear equations represented by the following matrix equation:

$$\mathbf{x} = \Phi \mathbf{a}, \quad (3.1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the known vector, $\Phi \in \mathbb{R}^{d \times D}$ is a known matrix and $\mathbf{a} \in \mathbb{R}^D$ is an unknown vector.

The vector \mathbf{x} can be interpreted as a vector of *linear measurements* or *observations* of the unknown vector \mathbf{a} through the matrix Φ .

If the dimension of \mathbf{x} is less than the dimension of \mathbf{A} , that is if $d < D$, then the system is referred to as *underdetermined*, for there are lesser number of measurements available than the dimension of the unknown vector. If $d = D$, then it is referred to as *determined* system and Φ is a square matrix. Lastly, if $d > D$ then it is called an overdetermined system of linear equations.

In the determined setting, if the matrix Φ has full rank d then it has an inverse Φ^{-1} and the unknown vector can be found by simply inverting the system: $\mathbf{a} = \Phi^{-1} \mathbf{x}$. If the matrix does not have an inverse, then the system is equivalent to an underdetermined system and finding an exact solution is not so straightforward when the system is underdetermined.

In the overdetermined case, if the matrix Φ has full rank d , then finding an exact solution for the system is generally difficult, and hence one looks for an approximate solution which minimises the least squares error between $\Phi \mathbf{a}$ and \mathbf{x} [52]. That is, one

3. SOLUTIONS OF UNDERDETERMINED LINEAR SYSTEMS

solves: $\min_{\mathbf{a}} \|\Phi \mathbf{a} - \mathbf{x}\|_2$. If the rank of the matrix is less than d , then the system can be transformed either into an underdetermined system.

When the system is underdetermined, then the system admits infinitely many solutions and hence it is not unique unless additional hypothesis is placed on the nature of the solution. As one possibility, we can look for a solution which has minimum ℓ^2 norm among all the admissible solutions, and this is called the *minimum norm* solution.

Alternatively, if we know that the solution that we are looking for is *sparse*, then we can incorporate this information to efficiently find the solution, which is possibly unique. This is the well known *sparse recovery problem*.

3.2 Sparse recovery problem

In a general sparse recovery problem, the challenge is to find a sparse solution to an underdetermined system of linear equations. As discussed in Sec. 2.3, the sparsity of a vector is measured using its ℓ^0 norm or the number of non-zero coefficients in the vector. Hence, the sparse recovery problem can be mathematically expressed as

$$\min_{\mathbf{a}} \|\mathbf{a}\|_0 \quad \text{subject to} \quad \mathbf{x} = \Phi \mathbf{a}. \quad (3.2)$$

To solve (3.2) directly, we must examine all possible combinations of non-zero coefficients of the vector \mathbf{a} . This is equivalent to the *subset selection* problem and it is intractable because the search space is exponentially large [53].

To mitigate this combinatorial complexity, there are two approaches that are considered.

3.2.1 Algorithmic approaches

Firstly, one can solve (3.2) by using heuristics such as *greedy approaches*. The idea behind a greedy approach is the following: the vector \mathbf{x} is successively approximated by choosing columns of the matrix Φ one at a time in a greedy fashion, and forming a linear combination of them. The approach consists of two steps: 1) finding which columns to choose and 2) finding the contribution of the chosen columns to the linear combination and obtaining a residual. The different ways in which the residual is computed gives rise to several flavours of the approach, collectively known as the *greedy pursuit* algorithms.

Examples of greedy pursuit algorithms include Matching Pursuit (MP) by Mallat and Zhang [54], Orthogonal Matching Pursuit (OMP) [55], Stagewise OMP (StOMP) [56], Regularized OMP (ROMP) [57, 58], etc. Free implementations of all these algorithms are available on the Internet. For example, see Matching Pursuit ToolKit (MPTK) by Krstulovic and Gribonval [59], Sparsify by Blumensath [60], etc.

The other approach that is commonly taken is to replace the ℓ^0 quasi-norm with the ℓ^1 norm to obtain the following optimisation problem

$$\min_{\mathbf{a}} \|\mathbf{a}\|_1 \quad \text{subject to} \quad \mathbf{x} = \Phi \mathbf{a}. \quad (\text{BP})$$

We noted in Sec. 2.3 that the ℓ^1 norm is the convex norm that is closest to the ℓ^0 quasi-norm and hence (BP) is a convex program. This substitution of the norms is called as *convex relaxation* and the mathematical program to find sparse solution to an underdetermined system is called *basis pursuit* [61].

The basis pursuit problem can be cast as a linear program and standard algorithms such as the interior point methods, homotopy methods, etc. can be used to solve it [62]. There are several standard numerical algorithm packages that are made freely available on the internet by the research community. Examples include ℓ^1 -magic by Candés and Romberg [63], Sparselab by Donoho [64], CVX by Boyd and Grant [65], etc. There are also algorithms such as the FOcal Underdetermined System Solver (FOCUSS) [66] which aim at recovering an approximate solution of (BP).

The basis pursuit (BP) demands the solution to satisfy an exact equality $\mathbf{x} = \Phi \mathbf{a}$. However, one might not be able to find an exact sparse solution always, say for example when there is additive noise on the measurements. Even in such cases, a modified version of (BP) can be used to recover the sparse solution by allowing some error margin on the constraint:

$$\min_{\mathbf{a}} \|\mathbf{a}\|_1 \quad \text{subject to} \quad \|\mathbf{x} - \Phi \mathbf{a}\|_2 \leq \epsilon. \quad (\text{BP-error})$$

It is also known that the solution to (BP-error) is exactly the same as the solution to the following unconstrained convex program, known as *basis pursuit denoising*, for an appropriate Lagrange multiplier λ

$$\min_{\mathbf{a}} \lambda \|\mathbf{a}\|_1 + \frac{1}{2} \|\mathbf{x} - \Phi \mathbf{a}\|_2^2. \quad (\text{BPDN})$$

The convex program (BP-error) can be cast as Second Order Cone Program (SOCP) and interior method points can be used to solve it [62]. The basis pursuit denoising problem (BPDN) is referred to as Least Absolute Shrinkage and Selection Operator (LASSO) [67] by the statistics community and a numerical method proposed in [68] solves (BPDN) for all possible values of λ at once. Several other methods to solve (BPDN) have been proposed in the literature such as Iterative Reweighted Least Squares (IRLS) for sparse recovery [69], etc.

Of late, a new family of numerical algorithms has been conceived to solve (BPDN) which are collectively called *iterative-shrinkage* algorithms. These algorithms are very simple in nature and can solve the optimisation problem very efficiently when the problem sizes are huge. A very elegant overview of these algorithms can be found in [70].

A framework called SMALLBox [71, 72] to evaluate various sparse recovery algorithms has been developed recently, which provides a common interface and interoperability among various sparse recovery toolboxes available on the Internet as open source implementations.

We can ponder about the ability of the greedy methods and the convex relaxation methods (BP), (BP-error) and (BPDN) to find a correct and unique solution to the problem (3.2) that we are originally interested in. The conditions under which the solution can be found and its uniqueness can be assessed are called *recovery conditions*.

3. SOLUTIONS OF UNDERDETERMINED LINEAR SYSTEMS

Lately, there has been an explosion of results on the recovery conditions for both algorithmic approaches and almost all the theoretical results are primarily characterised by the properties of the matrix Φ and the sparsity of the vector \mathbf{a} .

3.2.2 Recovery guarantees

One of the properties of the matrix Φ that governs the recovery abilities of (BP) is its *coherence*.

Coherence: If the columns φ_n , $1 \leq n \leq D$ of the matrix Φ are normalised to have unit ℓ^2 norm, i.e., $\|\varphi_n\|_2 = 1$, $1 \leq n \leq D$ then the coherence of Φ is defined as

$$\mu := \max_{m \neq n} |\langle \varphi_m, \varphi_n \rangle|. \quad (3.3)$$

In other words, the coherence of a matrix Φ is the maximum absolute off-diagonal element of its Gram matrix $\Phi^T \Phi$. Coherence tells us about the similarity between the columns of the matrix Φ . If the columns of Φ are orthogonal to each other, then its coherence is $\mu = 0$. Hence, if the value of μ of a matrix is smaller then the columns are nearly orthogonal to each other.

Another property that was introduced by Candès and Tao is the *Restricted Isometry Property* (RIP) of a given matrix Φ [73].

Restricted isometry property: For a matrix Φ of size d with ℓ^2 normalised columns, and for an integer $k \leq D$, let Φ_k denote a sub-matrix containing k columns of Φ . Then the matrix Φ is said to possess k -RIP with a constant δ_k if the following holds true

$$\forall \mathbf{c} \in \mathbb{R}^k \quad (1 - \delta_k) \|\mathbf{c}\|_2^2 \leq \|\Phi_k \mathbf{c}\|_2^2 \leq (1 + \delta_k) \|\mathbf{c}\|_2^2 \quad (3.4)$$

for all sub-matrices Φ_k of Φ . This property means that any sub-matrix of Φ containing k or less columns behaves like almost an orthonormal system.

These properties of the matrices are used to quantitatively establish the recovery performances of the algorithms. Results concerning the uniqueness and stability of the solution obtained by greedy pursuits and convergence of the algorithms in various problem settings are available in the literature [74, 75, 76, 77].

Results concerning the stability of solutions to (BP-error) and (BPDN), under random noise, based on the coherence of the dictionary can be found in [78].

3.3 Summary

This chapter provides an introduction to the problem of finding sparse solutions to underdetermined systems of linear equations, generally referred to as the sparse recovery problem. A brief overview of the two main approaches that have been developed to solve the sparse recovery problem has been given.

Greedy methods such as the matching pursuit and its variants find a sparse solution to a given system of linear equations in a step by step manner by choosing one coefficient of the solution at a time. Convex relaxation methods such as the ℓ^1 minimisation involves solving a convex mathematical program to obtain the solution.

3. SOLUTIONS OF UNDERDETERMINED LINEAR SYSTEMS

Chapter 4

Role of sparsity in channel estimation

Chapter 2 focussed on the general source separation problem and we introduced grand families of approaches that are found in literature. Chapter 3 introduced the standard sparse recovery problem and this chapter is focussed on the problem of channel estimation in both blind and non-blind settings, specifically using the convolutive observations.

4.1 Introduction

In the context of communications engineering, data transmitted over a channel is influenced by the channel characteristics. More so in wireless communications, because of the multiple reflection paths in the environment and the fading characteristics of the channels. When data is received at the receiver, due to the influence of the channel on the data, it will not be the same as the data that is originally transmitted at the source. So, one has to undo the influence of the channel in order to estimate or recover back the original source signal. In the communications engineering parlance, this is known as *channel equalisation*.

Communication channels (space, underwater, etc.) are often modelled by discrete finite impulse response filters and they are called *channel impulse responses*. Channel equalisation refers to finding a new set of filters called *equalisation filters*, which compensate the effect of channel impulse responses on the observations to provide an estimate of the source signals. In most cases, explicit knowledge of channels is not available and one has to estimate the equalisation filters by analysing the received signals, with or without the explicit knowledge of the source signals.

Estimating equalisation filters is a very well established area of research in communications engineering and most of the techniques that are used commonly are *non-blind* in nature. That is, the filters are estimated by periodically sending out known probe signals over the channels, and subsequently estimating the filters by analysing the received signals. These probe signals are called *training* or *pilot* signals.

4. ROLE OF SPARSITY IN CHANNEL ESTIMATION

In a typical multi-user communication system, where several users are sharing the transmission medium over the same time-frequency-code slot, the signals from various spatial origins interfere with the signal of interest (to a particular user) and hence it gives rise to *spatial mixtures*. This is called *co-channel interference* and this can be modelled as a blind source separation problem of linear-instantaneous mixtures. These transmission scenarios are known as instantaneous or static *Multi-Input-Multi-Output* (MIMO) channels.

The second kind of interference, called the *intersymbol interference* (ISI), arises when a transmitted signal gets corrupted by the time-delayed versions of itself. The time-delayed versions could arise due to multipath propagation of the signals, and what is received at the end is known as *temporal mixtures*. In a time-dispersive (MIMO) channel, which is often the case in multi-user wireless communication systems, the co-channel interference and the intersymbol interference have to be tackled simultaneously. This kind of equalisation is called *spatio-temporal equalisation*, which in the parlance of blind source separation, amounts to a convolutive source separation problem.

It should be possible to use the standard methods for convolutive mixtures to handle the channel equalisation methods. However, digital communication channels present particular features that can be capitalized to improve filter estimation and source recovery. The source signals in digital communications have a finite number of possible complex amplitudes, arising from the modulation techniques. Contrast functions that exploit this property make an excellent choice for source separation in communications context. Also, as mentioned earlier, pilot symbols are transmitted at regular intervals to aid the channel equalization.

Plan of the chapter

We start by briefly describing some of the classical approaches that are developed by the communications engineering community for channel equalisation, both blind and non-blind. We then turn our attention to the problem of estimating the channel impulse responses directly instead of the equalisation filters and describe a particular blind method for the same.

4.2 Methods for channel equalisation

A standard semi-blind method for channel equalisation in a single user setting is the *optimal Wiener filter* or the *minimum mean square error* (MMSE) equaliser. Consider a single source signal $s(t)$ of length T and a single user that receives the source signal that is transmitted through a channel whose impulse response is modelled by $a(t)$ of length L . Let the observed signal at the user be $x(t)$. By definition, the observed signal is related to the source signal by the following

$$x(t) = (s \star a)(t) + v(t), \quad (4.1)$$

4.3 Cross relation method for channel estimation

where $x(t)$ is of length $T + L - 1$, and $v(t)$ is the additive noise at the receiver, which is independent of the signal. Let $\mathbf{s}, \mathbf{a}, \mathbf{x}$ and \mathbf{v} denote the vectors corresponding to $s(t), a(t), x(t)$ and $v(t)$. Then, (4.1) can be written as

$$\mathbf{x} = \mathbf{s} \star \mathbf{a} + \mathbf{v}. \quad (4.2)$$

What we seek as an equaliser is another vector \mathbf{b} such that

$$\tilde{\mathbf{s}} = \mathbf{b} \star \mathbf{x}$$

is an accurate estimate of the source vector \mathbf{s} .

To achieve this, the supervised MMSE criterion aims at the minimisation of the cost function

$$ERROR_{MMSE}(\mathbf{b}) := \mathbb{E} \{ \|\tilde{\mathbf{s}} - \mathbf{s}\|_2^2 \} = \mathbb{E} \{ \|\mathbf{b} \star \mathbf{x} - \mathbf{s}\|_2^2 \}. \quad (4.3)$$

This is a non-blind method because it requires an explicit knowledge of the source term \mathbf{s} . Another widely used optimisation criterion is the standardized cumulant due to Donoho [79]. At fourth order, the standard cumulant is known as *kurtosis*, and is given by

$$ERROR_{KM}(\mathbf{b}) = \frac{\text{cum}_4\{\tilde{\mathbf{s}}\}}{\text{cum}_2^2\{\tilde{\mathbf{s}}\}} \quad (4.4)$$

where $\text{cum}_2\{\tilde{\mathbf{s}}\}$ is the variance of $\tilde{\mathbf{s}}$ and $\text{cum}_4\{\tilde{\mathbf{s}}\}$ is defined as

$$\text{cum}_4\{\tilde{\mathbf{s}}\} = \mathbb{E} \{ |\tilde{\mathbf{s}}|^4 \} - 2\mathbb{E} \{ |\tilde{\mathbf{s}}|^2 \}^2 - |\mathbb{E} \{ \tilde{\mathbf{s}}^2 \}|.$$

Cumulant based techniques do not require explicit knowledge of the original sources and hence they are blind.

Kurtosis maximisation has been widely used for MIMO channel equalisation [80] as well as source separation in instantaneous linear mixtures. Cumulant based method for coloured sources has been dealt with in [81].

Apart from the methods that are based on second order and higher order statistics, a host of other approaches are available to deal with the channel estimation problem [82, 83, 8, 84]. Our focus in this thesis is on channel estimation rather than channel equalisation and hence we shall not delve more on channel equalisation.

4.3 Cross relation method for channel estimation

We shall now focus on an interesting set of techniques for channel estimation in a blind *Single-Input-Multi-Output* scenario, based on the commutative and associative property of the convolution operation. Let us start with the simplest case of estimating filters when there is only one source s and two observations x_1 and x_2 . This is the single-input-two-output (SITO) case and we have:

$$x_i(t) = (a_i \star s)(t) + v_i(t), \quad i = 1, 2.$$

4. ROLE OF SPARSITY IN CHANNEL ESTIMATION

Let the length of s be T and the length of the filters be L , then the length of x_i will be $T + L - 1$. In the absence of noise, we have the following cross-relation (CR):

$$(x_2 \star a_1)(t) = (x_1 \star a_2)(t). \quad (4.5)$$

This is the key equation on which a number of blind channel estimation algorithms and methods have been proposed and we will exploit this relation in chapters 7 and 8 to propose a general framework for filter estimation. For convenience, let us associate the signal a_i to the column vector $\mathbf{a}_i = [a_i(t)]_{t=1}^L$ and likewise s to \mathbf{s} and x_i to \mathbf{x}_i .

The convolution $x_i \star a_j$ is associated to the multiplication between the Toeplitz matrix¹

$$\mathcal{T}[\mathbf{x}_i] = \begin{bmatrix} x_i(1) & 0 & \cdots & 0 \\ x_i(2) & x_i(1) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ x_i(L) & x_i(L-1) & \cdots & x_i(1) \\ x_i(L+1) & x_i(L) & \cdots & x_i(2) \\ \vdots & \vdots & \ddots & \vdots \\ x_i(T+L-1) & x_i(T+L-2) & \cdots & x_i(T) \end{bmatrix}, \quad (4.6)$$

and the vector \mathbf{a}_j . By using the shorthand $\mathcal{B}[\mathbf{x}_1, \mathbf{x}_2] = [\mathcal{T}[\mathbf{x}_2], -\mathcal{T}[\mathbf{x}_1]]$, we can write the CR (4.5) as

$$\mathcal{B}[\mathbf{x}_1, \mathbf{x}_2] \cdot \mathbf{a} = \mathbf{0}, \text{ where } \mathbf{a} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \end{bmatrix}. \quad (4.7)$$

When there are multiple channels instead of only two, we can write a relationship of the form (4.7) for every pair of channels, and then they can be stacked up to form a larger system of equations.

In contrast to the statistical approaches, the CR based approach relies only on solving a system of linear equations, which might lead to simpler implementations. Xu et al. [83] have proposed to minimise the squared energy of the term $\mathcal{B}[\mathbf{x}_1, \mathbf{x}_2] \cdot \mathbf{a}$ in order to obtain a solution. That is

$$\min_{\mathbf{a}} \|\mathcal{B}[\mathbf{x}_1, \mathbf{x}_2] \cdot \mathbf{a}\|_2^2. \quad (4.8)$$

They have further studied the identifiability properties of the problem and have derived the necessary and sufficient conditions for the solution to exist. The fundamental requirements of identifiability are 1) the polynomials corresponding to the filters \mathbf{a}_1 and \mathbf{a}_2 share no common roots and 2) the input signal $s(t)$ has sufficient "modes" or "diversity". The diversity of the input signal is nothing but the number of roots of a polynomial which characterises a rank deficient Hankel matrix (a square matrix with constant skew-diagonals) built using the input signal.

Similar methods which rely on the above cross relation include the component normalisation method by Avadeno et al. [85]. Huang and Benesty derive least mean square

¹Calligraphic letters will denote operators that map a vector to a matrix, e.g. $\mathcal{T}[\mathbf{x}_i]$.

and Newton algorithms based on the cross relation for blind channel identification in single-input-multi-output settings [86], and they further derive frequency domain algorithms, based on the same cross relation in [87].

Aïssa-El-Bey et al. [88] have further extended the formulation of the cross relation used in [8], and introduce a new identification method referred to as Minimum Cross-Relations (MCR) method which exploits the spatial diversity among the channel outputs. They also present a modified version of the MCR referred to as the "unbiased MCR" that leads to unbiased estimation of the channel parameters.

In all the representative approaches described till now in this section, the channel impulse responses were not assumed to be having any special structure. However, we are interested in the blind estimation of channels which are sparse in the time domain. Blind estimation of sparse channels is a recent trend in communication systems. In the next section, we shall embark upon the problem of sparse channel estimation and study some of these recent approaches that are based on the CR.

4.4 Sparse filter estimation

In the previous section, we described a few techniques that are generally used for estimating channel impulse responses in blind and semi-blind settings. In this section let us focus on the methods which take into account the sparse structure of the filters in order to aid its estimation. Before we get into sparse channel estimation approaches, we shall first describe the problem of recovering sparse vectors from its linear measurements in underdetermined settings. This is nothing but the problem of finding sparse solutions to underdetermined linear systems.

4.4.1 Sparse channel estimation: Single source setting

In the previous section, we described the general sparse recovery problem in the standard setting that is commonly addressed in the literature. In this section, let us re-examine the problem of channel impulse response estimation using the cross-relation approach that was already introduced in the beginning of this chapter, in the light of sparse recovery problem. The objective is to recast the channel estimation problem as a constrained ℓ^1 minimisation problem by considering the sparse structure of the channel impulse responses.

Let us recall that the impulse response of a communication channel can be interpreted as the multiple paths through which the signals get transmitted, and the magnitude of the impulse response peaks are the different attenuation and delays that are associated with the paths. In a variety of settings such as underwater communication, geological acoustics, etc., the number of reflection paths that the source signals take are very few compared to the delays associated and hence the channel impulse responses are sparse in the time-domain.

In the single source setting, we introduced the notion of cross relation (CR) between the channels in Eq (4.5), and the corresponding matrix form in Eq. (4.7). Now, with

4. ROLE OF SPARSITY IN CHANNEL ESTIMATION

the additional assumption that the vector $\mathbf{a} = [\mathbf{a}_1 \ \mathbf{a}_2]^T$ is sparse and denoting $\mathbf{B} := \mathcal{B}[\mathbf{x}_1, \mathbf{x}_2]$, the channel estimation problem can be cast as the following ℓ^1 minimisation problem:

$$\text{minimize } \|\mathbf{a}\|_1 \quad \text{subject to } \mathbf{B} \cdot \mathbf{a} = \mathbf{0}. \quad (4.9)$$

Solving this convex problem without further assumptions on the nature of \mathbf{a} will not be fruitful because the algorithm would give us the trivial solution $\mathbf{a} = \mathbf{0}$, which satisfies the constraint and also has the least possible ℓ^1 norm. Hence, we need additional constraint that prevents this. Also, it is sometimes unreasonable to expect the filters to satisfy an exact equality $\mathbf{B} \cdot \mathbf{a} = \mathbf{0}$ because of truncated source signals, noise at the observations, etc. Hence, it is more appropriate to allow a margin of error and solve a (BP-error) type of problem in order to estimate the filters.

Aïssa-El Bey *et al.* [89] have proposed to use the constraint $\|\mathbf{a}\|_2 = 1$ for avoiding the trivial solution and they formulate the channel estimation problem in the following way:

$$\text{minimize } \|\mathbf{a}\|_1 \quad \text{subject to } \|\mathbf{B} \cdot \mathbf{a}\|_2 \leq \epsilon \quad \text{and} \quad \|\mathbf{a}\|_2 = 1. \quad (4.10)$$

However, the normalisation $\|\mathbf{a}\|_2 = 1$ makes the problem non-convex and therefore, as an alternative, we use the constraint $\mathbf{a}_1(t_0) = 1$, where t_0 is an arbitrarily chosen time index which results in the following convex problem:

$$\text{minimize } \|\mathbf{a}\|_1 \quad \text{subject to } \|\mathbf{B} \cdot \mathbf{a}\|_2 \leq \epsilon \quad \text{and} \quad \mathbf{a}_1(t_0) = 1. \quad (4.11)$$

As mentioned in the previous section, any standard algorithmic package such as the CVX [65] can be used to solve (4.11). It is this template of the convex problem that is central to the framework that we eventually develop in this thesis.

In the following section, let us consider the problem of estimating channel impulse responses in the presence of multiple sources and how the CR based method can be tailored to approach the problem.

4.4.2 Sparse channel estimation: Multiple source setting

In a setting where there are multiple sources $s_j(t)$, $1 \leq j \leq N$ instead of just one, the CR formulation (4.7) is not valid anymore for any particular source $s_j(t)$. This is because the sources interfere with each other and what is observed at the receiver will be a sum of filtered versions of different source signals. This is exactly the model of source mixtures (1.1) that was introduced in Sec. 1.2. However, it is possible to use the CR based approach to identify channel impulse responses with some additional assumption on the sources.

Aïssa-El Bey *et al.* [90] have extended the CR in SITO based approach to a setting with N sources. The authors assume that there exists time segments in general where only one source is active alone and all the other sources are inactive. They further assume that at least one time segment is available for each of the sources and therefore

corresponding to these time segments, there are segments in the observations where the contribution is only from one of the sources. If these segments are identified from the observations, then the multiple source setting is now reduced to SITO setting for each segment, and a CR can be formulated locally at such segments.

Let us illustrate this using a simple example of two sources. Fig. 4.1(a) shows the time-domain plot of two sources s_1 and s_2 . Notice that at time segment I_1 , only source s_1 is active and at I_2 only source s_2 is active. Fig. 4.1(b) shows the plot of their mixtures, x_1 and x_2 , obtained by convolving the sources with the filters $a_{ij}, i = 1, 2$. These mixtures do not satisfy the CR (4.7) over the entire time frame because of time segments where both the sources are active simultaneously. However, corresponding to the time intervals I_1 and I_2 of the sources, there exists time intervals \tilde{I}_1 and \tilde{I}_2 in the observations, where only one source is contributing.

If the time intervals \tilde{I}_1 and \tilde{I}_2 can be identified, then we can extract the mixtures at these time segments to obtain $y_i^{(j)} = \{x_i(t)\}_{t \in \tilde{I}_j}$, which depend on only one source j .

The vectors $\mathbf{y}_i^{(j)}$ corresponding to $y_i^{(j)}$ now satisfy the CR $\mathcal{B}[\mathbf{y}_1^{(j)}, \mathbf{y}_2^{(j)}] \cdot \mathbf{a}^{(j)} \approx \mathbf{0}$. Note that the relationship is not an equality due to the boundary effect of convolution. Therefore, this relationship can be used to estimate the filters for source j by solving the optimisation problem (4.11) with $\mathbf{B} = \mathcal{B}[\mathbf{y}_1^{(j)}, \mathbf{y}_2^{(j)}]$.

The authors of [90] have developed a multiple channel estimation framework based on the assumption of time-domain disjointness of sources. The framework consists of the following steps:


- Blindly estimate the number of sources and the intervals $\tilde{I}_j, 1 \leq j \leq N$ based on the rank of the covariance matrices of the mixtures at various instances.
- Extract the mixtures at each of these intervals and solve the channel estimation problem using the CR.
- It could so happen that multiple time intervals for the same source are identified and hence the channel impulse responses are identified multiple times for the same source. Hence, a clustering step follows the filter estimation step to cluster the impulse responses belonging to the same source.

The key factor which enables the use of CR based method to estimate channels in multiple source setting is the time-domain disjointness of the sources. The framework that is developed in this thesis is a generalisation of the work presented in [90], and it is based on the time-frequency domain disjointness of the sources, which has been dealt with in chapter 2.


4.5 Summary

In this chapter, we introduced the problem of estimating channel impulse responses from the observed signals in the context of communication systems for channel equalisation purposes. We first described some of the legacy approaches that are used in

4. ROLE OF SPARSITY IN CHANNEL ESTIMATION



(a) Sources with intervals where only one source is active.



(b) Mixtures from the sources.

Figure 4.1: An example of time-frequency domain sparse signal

the communication systems for channel equalisation. We then looked at the channel estimation problem and described a blind approach based on the cross-relation that exists between the channels in the SITO setting.

Later, we described how the sparse information about the filters can be incorporated in the channel estimation problem to formulate a convex program for filter estimation in the single source setting. Then we described an existing work about the extension of SITO based sparse filter estimation approach in the multiple source setting, which primarily relies on the time-domain disjointness of the sources.

Part II

Contributions

Proposition of the thesis

In Sec. 1.3.1, we saw how the explicit knowledge of the source and the sparsity of the filter can be used to estimate a filter, with just one observation. Instead of a single observation, suppose we have two observations x_1 and x_2 , through the filters a_1 and a_2 . This is the single-input-two-output (SITO) setting, and due to the commutativity and associativity property of the convolution operator the following holds

$$(x_2 \star a_1)(t) = (x_1 \star a_2)(t) = (a_2 \star a_1 \star s)(t), \quad \forall t. \quad (4.12)$$

This implies the cross-relation (CR) [8]:

$$(x_2 \star a_1 - x_1 \star a_2)(t) = 0, \quad \forall t. \quad (4.13)$$

The CR provides us with a constraint that does not require the knowledge of the sources explicitly. This has been extensively used in communications engineering to estimate the filters blindly and further if the filters are sparse, then an optimisation problem can be formulated and solved to estimate a_1 and a_2 simultaneously. This is a blind estimation scheme, and it is still in the single source setting. The question now we can ask is: how can we exploit this CR to estimate mixing filters in the multiple source setting?

It is at this juncture that the time-frequency domain sparsity of the sources plays a role. We have mentioned in Sec. 1.3.2 that a commonly used hypothesis is that the sources are sparse and disjoint in the time-frequency domain. With this hypothesis, we can find time-frequency points in the mixtures where only one of the sources is active. Hence, even though the CR in Eq (4.12) might not be satisfied in the time-domain for any source, because of the presence of other sources, a suitable time-frequency domain version of the CR will be satisfied at those time-frequency points where only one source is active. Then these time-frequency domain CR for each such point can be used to formulate an optimisation problem for filter estimation.

This thesis is focussed on the development of a framework which relies on the time-frequency domain sparsity and disjointness of the sources and the time-frequency domain version of the CR. Fig. 4.2 shows a symbolic diagram of our contribution in the thesis.

We propose a framework which exploits the source and filter sparsities for the task of estimating multiple sparse filters from stereo convolutive mixtures. The main tasks that are involved in such a framework are: 1) Identifying the time-frequency points in the mixtures where only one source is active, and clustering them according to the sources and 2) estimating the filters by exploiting the time-frequency domain CR in each cluster.

Such a system would have a workflow as depicted in the Fig. 4.3. The steps involved can be summarized as:

1. Transformation of the mixtures into the time-frequency domain where the sparse structure of the sources are made explicit;

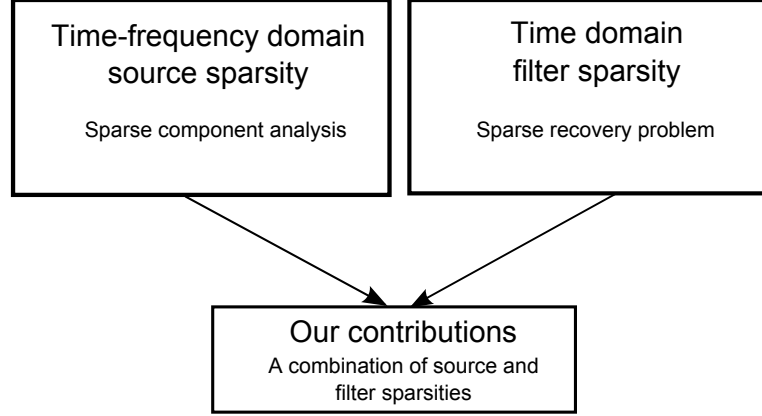


Figure 4.2: Our contribution in the thesis.

2. Identification of the time-frequency points where only one source is active;
3. Clustering of the points and formulation of the filter estimation problem;
4. Solution of the filter estimation problem.




Figure 4.3: Workflow of the filter estimation system based on the sparsity of sources and mixing filters.

Before we go on with the development of a framework for multiple sparse filter estimation, we first study the utility of sparsity of mixing filters to solve the problem of permutation ambiguity in blind source localisation with convolutive filters.

Chapter 5

Convolutional source localisation: Could we exploit filter sparsity?

One of the classical approaches for convolutional blind source localisation and separation is to transform a single, time domain convolutional mixtures problem into multiple, time-frequency domain instantaneous mixture problems using the narrowband approximation. This transformation enables us to use the techniques developed for separating linear instantaneous mixtures, such as independent component analysis (ICA) and sparse component analysis (SCA), to perform localisation and separation in each frequency sub-band. However, as pointed out in Sec. 2.3.4, this scheme gives rise to the arbitrary permutation and scaling in each sub-band of the estimated objects (filters or sources) and they have to be corrected in order to perform localisation and separation.

As it was mentioned in chapter 2, the permutation and scaling problems are inherent to the mathematical formulation of the source separation itself. Hence, whenever we intend to solve the convolutional problem using the narrowband approximation, these problems are bound to arise irrespective of which technique we use to tackle the problem. One possibility to avoid these problems is to handle the problem in the time-domain itself, like the time-domain convolutional ICA. However, these time-domain techniques also suffer from the permutation and scaling ambiguities [29] and also they are computationally very intensive and hence it is advantageous to work in the time-frequency domain.

Another advantage of working with the time-frequency domain and narrowband approximation is that we can exploit the sparsity of the sources in the time-frequency domain. As it was discussed in chapter 2, the time-frequency domain sparsity of the sources can be very elegantly exploited to design techniques for source localisation and separation. And with this time-frequency domain approach, it is imperative that we have to solve the permutation and scaling ambiguities.

In this chapter, we only consider the permutation ambiguity of filter coefficients. We suppose that the filter estimates in each frequency sub-band are obtained by using one of the techniques that we have described in chapter 2, and the contribution of this chapter is to assess the potential of sparsity to solve the permutation ambiguity by exploiting the

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

time-domain sparsity of the mixing filters. A combinatorial ℓ^1 minimisation algorithm is also presented to correct the permutation ambiguities in the absence of scaling.

Plan of the chapter

We begin with a recollection of the narrowband approximation of the convolutive mixture model and state the permutation and scaling problems. Then, a brief survey of approaches that are found in the literature to solve these problems is presented. Subsequently, we show how the sparsity of mixing filters can be used to solve the permutation problem by presenting some experimental results which bring out the connections between the sparsity of the filters and permutations. Motivated by these experimental results, we then present combinatorial algorithm to solve the permutation problem. The algorithm is experimentally evaluated and theoretical connections between filter permutations and filter sparsity is presented at the end.

5.1 Permutation and scaling problems: State of the art

Let us consider N_T short time Fourier transform (STFT) frames of the mixtures \mathbf{x}_i , $1 \leq i \leq M$, with a framesize F . If the STFT window size F is sufficiently larger than the filter length L , then the mixtures can be approximated by:

$$\hat{x}_i(\tau, f) \approx \sum_{j=1}^N \hat{a}_{ij}(f) \hat{s}_j(\tau, f), \quad 1 \leq i \leq M, \quad 1 \leq \tau \leq N_T, \quad 0 \leq f < F, \quad (5.1)$$

where $\hat{x}_i(\tau, f)$ and $\hat{s}_j(\tau, f)$ represent the time-frequency representation of the mixtures and sources respectively at the time-frequency index (τ, f) and $\hat{a}_{ij}(f)$ represents the Discrete Fourier Transform (DFT) of the filters at frequency f . This formulation transforms a single convolutive source separation problem into F independent complex instantaneous mixtures problem in the time-frequency domain.

Let us denote the filter coefficients $a_{ij}(t)$ at a given time index t by a matrix $\mathbf{A}(t)$ whose size is $M \times N$. That is, the $(ij)^{th}$ element of $\mathbf{A}(t)$ is $a_{ij}(t)$. Similarly, let us denote the DFT coefficients of the filters at a frequency f by a matrix $\hat{\mathbf{A}}(f)$ of size $M \times N$.

Using the narrowband approximation, we can estimate the filter matrices $\hat{\mathbf{A}}(f)$ independently for each frequency bin f using one of the techniques based on ICA, or NMF or SCA. Irrespective of the approach used, the filter matrices can be only estimated up to a permutation and scaling, which might differ for each frequency sub-band f . Hence, the filter estimates are of the form

$$\check{\mathbf{A}}(f) = \hat{\mathbf{A}}(f) \mathbf{\Lambda}(f) \mathbf{P}(f), \quad 0 \leq f \leq F/2, \quad (5.2)$$

where $\mathbf{\Lambda}(f)$ is an $N \times N$ diagonal matrix containing arbitrary complex scaling factors and $\mathbf{P}(f)$ is an $N \times N$ permutation matrix. Because of the conjugate symmetric structure of the DFT, the conjugate frequencies f and $F - L$ are permuted in a similar way and they

are scaled by conjugates. Hence, it is sufficient to have the frequency indices running only up to $F/2$.

Due to these permutation and scaling ambiguities, the frequency domain filters cannot be right away transformed back into the time domain. Hence, the filter estimation stage of any source separation system that relies on the narrowband approximation has to correct these ambiguities before the filters are used in the source estimation stage.

We recall that Fig. 2.5 shows a typical architecture of a filter estimation block of a convolutive source separation system that is based on narrowband approximation. The estimates of the filter sub-bands $\check{\mathbf{A}}(f)$ are obtained for each frequency f , which has a form of (5.2). The permutation and scaling correction stage has to resolve the ambiguities for each sub-band to obtain the corrected estimates $\tilde{\mathbf{A}}(f)$, $0 \leq f \leq F/2$. However, there still could be a global permutation and a global scaling, which is independent of the frequencies. This global ambiguity does not pose any further problems for the source estimation stage, and hence it is often considered to be not an issue. Concisely, the goal of the permutation and scaling correction stage of the filter estimation process can be summarised as:

Goal of the permutation and scaling correction stage: *Starting with the filter estimates $\check{\mathbf{A}}(f)$, obtain a scaling and permutation corrected version $\tilde{\mathbf{A}}(f)$ up to a global scaling $\mathbf{\Lambda}_g$ and global permutation \mathbf{P}_g with respect to the true filter matrices $\hat{\mathbf{A}}(f)$, such that*

$$\tilde{\mathbf{A}}(f) = \hat{\mathbf{A}}(f)\mathbf{\Lambda}_g\mathbf{P}_g, \quad \forall f. \quad (5.3)$$

However, in this chapter we will assume that the filter estimates do not possess any scaling ambiguity and we deal only with filter permutations. Hence, $\mathbf{\Lambda}(f) = \mathbf{I}_{N \times N}$, $\forall f$ and $\tilde{\mathbf{A}}(f) = \hat{\mathbf{A}}(f)\mathbf{P}(f)$, $\forall f$. Though in reality the filter estimates suffer from scaling ambiguities also, we will consider only the permutation ambiguity in this chapter to show that the time-domain sparsity can indeed help solve the permutation problem.

Due to the fact that a global permutation of the filters is immaterial to the source separation step, we can arbitrarily fix a permutation, whose permutation matrix is \mathbf{P}_g , as a reference permutation for the purpose of aligning the sub-bands. Clearly, in this case we only need to align the sub-bands which have undergone permutations other than the permutation corresponding to the matrix \mathbf{P}_g . Let us denote the set of frequency sub-bands, whose permutations are other than the one corresponding to \mathbf{P}_g , by $\Omega(\mathbf{P}_g)$. That is:

$$f \in \Omega(\mathbf{P}_g) \iff \mathbf{P}(f) \neq \mathbf{P}_g. \quad (5.4)$$

With this definition of the set $\Omega(\mathbf{P}_g)$, we can now state that the goal of the permutation correction stage is to make the set $\Omega(\mathbf{P}_g) = \emptyset$, for some choice of \mathbf{P}_g .

5.1.1 Notations

For further discussions related to sub-band permutations, let us set up some additional notations. In Sec. 1.2.3 we introduced the notion of sparse filters, and in Sec. 4.4

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

we introduced the relevance of using ℓ^1 norm of vectors to measure its sparsity. In a standard sparse recovery problem, we are interested in recovering a single sparse vector and hence the ℓ^1 norm of that single vector is sufficient for consideration. However, in the context of source localisation with N sources and M channels we have $M \times N$ filter vectors \mathbf{a}_{ij} , $1 \leq i \leq M$, $1 \leq j \leq N$, and we are interested in the ℓ^1 norm the entire filter matrix.

Let \mathbf{A}_{TD} be a shorthand notation for representing all the MN time-domain filter vectors \mathbf{a}_{ij} . By a straightforward extension of the definition of the ℓ^1 norm for a single vector, we can define the ℓ^1 norm of \mathbf{A} as the sum of the individual ℓ^1 norms of each filter \mathbf{a}_{ij} . That is:

$$\|\mathbf{A}_{\text{TD}}\|_1 := \sum_{i=1}^M \sum_{j=1}^N \|\mathbf{a}_{ij}\|_1. \quad (5.5)$$

It should be noted that the sparsity of the filter matrix is measured in the time domain, whereas the permutations occur in the frequency domain. In order to study the interplay between sub-band permutations and filter sparsity, we have to often switch between the frequency-domain and time-domain representations and hence it is important to distinguish the notations. Table 5.1 list the notations the we shall use and their associated meanings.

Notation	Meaning
$\mathbf{A}(t)$	An $M \times N$ matrix which contains all the filter coefficients at time t
$\hat{\mathbf{A}}(f)$	An $M \times N$ matrix which contains all the filter coefficients at frequency f
\mathbf{A}_{TD}	Shorthand notation representing all the MN time-domain filter vectors \mathbf{a}_{ij}
$\hat{\mathbf{A}}_{\text{fd}}$	Shorthand notation representing all the MN frequency-domain filter vectors $\hat{\mathbf{a}}_{ij}$
$\check{\mathbf{A}}(f)$	An $M \times N$ filter matrix of the form $\check{\mathbf{A}}(f) = \hat{\mathbf{A}}(f)\mathbf{P}(f)$
$\check{\mathbf{A}}_{\text{fd}}$	Shorthand notation for the frequency domain matrices $\check{\mathbf{A}}(f)$, $\forall f$
$\check{\mathbf{A}}_{\text{TD}}$	Shorthand notation for the time-domain filters obtained by the inverse DFT of $\check{\mathbf{A}}_{\text{fd}}$
$\tilde{\mathbf{A}}_{\text{fd}}$	Shorthand notation for the permutation corrected frequency domain matrices $\tilde{\mathbf{A}}(f)$, $\forall f$
$\tilde{\mathbf{A}}_{\text{TD}}$	Shorthand notation for the time-domain filters obtained by the inverse DFT of $\tilde{\mathbf{A}}_{\text{fd}}$

Table 5.1: List of notations.

Before we go ahead with the description of our contributions, let us first look at some of the existing work in the literature which are aimed at tackling the permutation problem.

5.1.2 Existing approaches for permutation correction

In the two stage source separation architecture which relies on narrowband approximation for filter estimation, the permutations can be corrected either on the sub-bands of the estimated filter coefficients, or on the sub-bands of the estimated source signals, or jointly on both. Irrespective of whether the permutation correction is done on the filter sub-bands or on the estimated source signals, the underlying idea is to exploit some sort of *consistency* across the sub-bands of the underlying object (filters or sources). As mentioned in the beginning of this chapter, our contribution is to show how the permutations of the estimated filter coefficients can be corrected using filter sparsity as a consistency measure. Therefore we discuss only the methods that aim to correct the permutations on the estimated filter coefficients.

Several approaches that are found in the literature use different methods to establish the consistency of filter coefficients across the sub-bands. For example, constraints on the length of the filters, information about the location of the sources, etc.

Exploiting spectral consistency

One simple way of measuring the consistency of filter coefficients is to look at the continuity of the filter values across the sub-bands in the frequency domain. A naive approach could be to pick those permutations that minimise the Euclidean distance between neighbouring frequencies [91]. That is, the set of permutation matrices $\{\mathbf{P}(f)\}_{f=1}^F$ are chosen such that

$$\min_{\{\mathbf{P}(f)\}_{f=1}^F} \sum_{f=1}^F |\check{\mathbf{A}}(f)\mathbf{P}(f) - \check{\mathbf{A}}(f-1)|. \quad (5.6)$$

Continuity in frequency domain can also be interpreted as the smoothness of the filter coefficients in the frequency domain, and smooth frequency domain representation implies a limited temporal support of the filters. A simple way to ensure this is to zero-pad the time domain filters (essentially, increasing the frequency resolution by zero padding) before frequency transformation [92].

Instead of using merely the differences between adjacent frequency sub-bands, Asano et al. have suggested to use the cosine of the angle between the filter coefficients of different frequencies for measuring the smoothness across the sub-bands [93]. They assume that due to the smooth nature of the filters in the time domain the columns of the mixing filter matrix for adjacent frequencies have minimal angle. Hence, the permutations can be corrected by maximising the cosine of the angle between adjacent frequencies. Given two adjacent frequencies f_1 and f_2 , the cosine of the angle between the columns of $\check{\mathbf{A}}(f_1)$ and $\check{\mathbf{A}}(f_2)$ is defined by

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

$$\cos \alpha_{(f_1, f_2)}^{(j)} = \frac{\check{\mathbf{A}}_j(f_1)^H \check{\mathbf{A}}_j(f_2)}{\|\check{\mathbf{A}}_j(f_1)\|_2 \|\check{\mathbf{A}}_j(f_2)\|_2}, \quad (5.7)$$

where $\check{\mathbf{A}}_j(f)$ indicates the j^{th} column of the matrix $\check{\mathbf{A}}(f)$. Hence, the permutation across adjacent frequencies can be corrected by maximising the following cost function

$$C = \sum_{f=1}^F \sum_{j=1}^N \cos \alpha_{(f, f-1)}^{(j)}. \quad (5.8)$$

Exploiting spatial information

Continuity across the frequencies can also be assessed in terms of the estimated spatial locations of the sources. The mixing filters are essentially the impulse responses between the source locations and the microphone locations and hence the location information such as the Direction Of Arrival (DOA) can be used to correct the permutations. Methods which use this information are known as *beamforming* approaches. However, most beamforming approaches work on the unmixing filters rather than the mixing filters [94, 95, 96, 97].

If the sources are at a sufficient distance from a microphone array, then the sound waves arrive as plane waves to the microphones. Let the microphone array have a separation distance of d , and the plane waves from a given source be incident on the microphone array at an angle θ , called the DOA. Then the microphones sense the sound waves with a delay δ which is given by

$$\delta = \frac{d}{c} \sin \theta \quad (5.9)$$

where c is the velocity of the sound waves.

The separation filters can be interpreted as *null-steering beamformers* which respond only to a particular source and reject all other sources. A beamforming filter which responds constructively to a specific DOA and reject other interfering DOA has to account for the delay in arrival of sound waves, and these delays correspond to the phase characteristics of the filter. Hence, DOA information can be used for correcting permutations of the separation filters. DOA information is generally estimated from the observed signals and estimated separation filter matrices at different frequencies.

Other methods

Several other methods based on various assumptions have also been proposed. Methods utilising the visual clues of the sources [98], maximum likelihood approach [99], split spectral difference [100], etc. have been proposed. Table 5.2 is a partial reproduction of table 4 from [101] which contains a listing of various consistency measures that are used for correcting filter permutations, along with the references.

In what follows, we propose to use the time-domain sparsity of the mixing filters as a consistency measure to correct the permutations.

5.2 Assessing the impact of permutations on ℓ^1 norm of filters

Metric used	Reference
Smooth spectrum	[92], [28]
Source locations	[102]
Directivity pattern	[99], [94], [95]
Location vectors	[93]
Direction of arrival	[96], [97], [103]
Adjacent matrix distance	[91]
Invariances	[104]
Split spectrum	[100]
Vision	[98]

Table 5.2: List of metrics used to solve the permutation problem in the frequency domain and their corresponding references.

5.2 Assessing the impact of permutations on ℓ^1 norm of filters

In this section, we shall demonstrate that the ℓ^1 norm of the filters is effectively correlated to the permutations, by assuming that the mixing filters are sparse in the time-domain.

In order to establish the connection between permutations and the ℓ^1 norm of the filter matrix, we present some experimental results that shows empirically that if the filters are sparse, then only when all the permutations are aligned correctly the ℓ^1 norm of the filters reaches the minimum. We illustrate this claim with two types of experiments:

Experiment Type 1: The purpose of this experiment is to demonstrate that the ℓ^1 norm of the filter matrix has its minimum only when all the permutations are correctly aligned.

Experiment Type 2: The purpose of this experiment is to demonstrate the sensitivity of the ℓ^1 norm of the filter matrix.

We shall reiterate that we are treating the permutation problem alone in the absence of scaling. That is $\mathbf{\Lambda}(f) = \mathbf{I}_{N \times N}$, $0 \leq f \leq F/2$, where $\mathbf{I}_{N \times N}$ is an identity matrix of size $N \times N$.

For both types of the experiments, we shall use the following filter generation method.

Sparse filter generation for both experiments

For each time-domain filter $a_{ij}(t)$ of length L and sparsity k , let Γ_{ij} denote the support of filters such that $|\Gamma_{ij}| = k$. That is

$$\Gamma_{ij} := \{t \mid a_{ij}(t) \neq 0\}. \quad (5.10)$$

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

In all our experiments, the non-zero coefficients of each filter $a_{ij}(t)$ are generated as i.i.d. Gaussian with mean zero and variance one and the support sets Γ_{ij} of size k are selected uniformly at random from the set $\{1, 2, \dots, L\}$. That is

$$a_{ij}(t) = \begin{cases} \sim \mathcal{N}(0, 1) & \text{if } t \in \Gamma_{ij}, \\ 0 & \text{otherwise.} \end{cases} \quad (5.11)$$

These filters are transformed into the frequency domain by taking their discrete Fourier transform (DFT) of length $F = L$ and subsequently the matrices of filter coefficients $\hat{\mathbf{A}}(f)$, $0 \leq f < L$ are obtained. These are our original filters and let us denote their ℓ^1 norm by $\|\mathbf{A}_{\text{td}}\|_1$.

Protocol for experiment type 1:

The set of frequencies at which the filter matrices undergo permutation is denoted by Ω and in this experiment, the set $\Omega \subseteq \{0, 1, \dots, L/2\}$ is chosen by selecting randomly the frequency indices. For each $f \in \Omega$, a permutation matrix $\mathbf{P}(f)$ is randomly generated and the permuted matrices $\check{\mathbf{A}}(f) = \hat{\mathbf{A}}(f)\mathbf{P}(f)$ are obtained. For every f , the same permutation is applied on the corresponding mirror sub-band as well. That is

$$\check{\mathbf{A}}(f) = \begin{cases} \hat{\mathbf{A}}(f)\mathbf{P}(f) & \text{if } f \in \Omega \text{ or } (L - f) \in \Omega, \\ \hat{\mathbf{A}}(f) & \text{otherwise.} \end{cases} \quad (5.12)$$

With these permuted matrices $\check{\mathbf{A}}(f)$, $0 \leq f < L$, the corresponding time domain filters $\check{\mathbf{A}}_{\text{td}}$ are obtained by the inverse DFT, and their ℓ^1 norm $\|\check{\mathbf{A}}_{\text{td}}\|_1$ is computed using (5.5).

Results

We conducted experiments with filter matrices having the number of channels $M = 3$, number of sources $N = 4$, filter length $L = 1024$ for different values of sparsity k of individual filters. Fig. 5.1 shows the plot of maximum, minimum and mean of the relative difference in ℓ^1 norms $(\|\check{\mathbf{A}}_{\text{td}}\|_1 - \|\mathbf{A}_{\text{td}}\|_1)/\|\mathbf{A}_{\text{td}}\|_1$ versus $|\Omega|$ over 50 draws for different values of k .

We can notice from the figures that in each draw of the experiment, the quantity $\|\check{\mathbf{A}}_{\text{td}}\|_1 - \|\mathbf{A}_{\text{td}}\|_1$ is never negative, which means that the ℓ^1 norm tends to increase when the sub-bands are not aligned. Also, the trend is the same for larger values of $|\Omega|$.

We conducted experiments for different configurations of the problem sizes and the observations were consistent with the ones presented here.

Conclusion

From the experimental results presented in this section, we conclude that when the filter sub-bands are permuted randomly the sparse structure of the filters is disturbed and the ℓ^1 norm of the filters increases compared to the true ℓ^1 norm. Therefore, the

5.2 Assessing the impact of permutations on ℓ^1 norm of filters

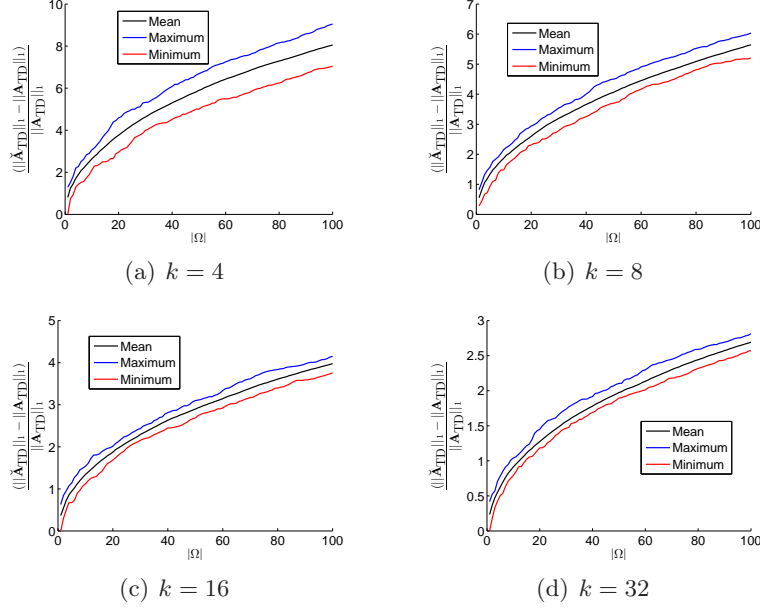


Figure 5.1: Plots showing the increase of ℓ^1 norm of filter matrices due to permutations.

sparsity of the filters can be effectively used as a consistency measure to correct the permutations.

Now that we have established that random permutations in the sub-bands increase the ℓ^1 norm, we shall see how sensitive the ℓ^1 norm of the filters is to the permutations in the next experiment.

Protocol for experiment type 2:

The previous experiment empirically established that random permutations of the sub-bands does increase the ℓ^1 norm of the filters. The following experiment shows that even a *single* permutation between two sources in only one sub-band ($|\Omega| = 1$) can increase the ℓ^1 norm of the filters.

The frequency f at which we wish to permute the sources is chosen randomly from the set $\{0, 1, \dots, L/2\}$. At a single frequency f and its conjugate, we swap *only* two sources j_1 and j_2 that are pre-determined to obtain the permuted matrices $\tilde{\mathbf{A}}(f)$. Let us call this as $j_1 - j_2$ source-pair permutation. The corresponding time domain filters $\tilde{\mathbf{A}}(t)$ are obtained by the inverse DFT, and their ℓ^1 norm $\|\tilde{\mathbf{A}}_{\text{td}}\|_1$ is computed using (5.5).

Results

We conducted experiments with filter matrices having the number of channels $M = 3$, number of sources $N = 4$, filter length $L = 1024$ for different values of sparsity k of

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

individual filters. Fig. 5.2 shows the box plot of the relative difference in ℓ^1 norms $(\|\hat{\mathbf{A}}_{\text{td}}\|_1 - \|\mathbf{A}_{\text{td}}\|_1)/\|\mathbf{A}_{\text{td}}\|_1$ versus k over 50 draws for different source-pair permutation.

For a given sparsity level and a trial number, and at a random frequency f , we can have $\binom{N}{2}$ different source pair permutations. Fig. 5.2(a) shows the variation in ℓ^1 norm when the sources 2 and 3 are permuted at a single randomly chosen frequency index using Matlab's `boxplot()` function. Similarly, Figs. 5.2(b), 5.2(c) and 5.2(d) correspond to 1 – 4, 4 – 2 and 3 – 1 source pair permutations.

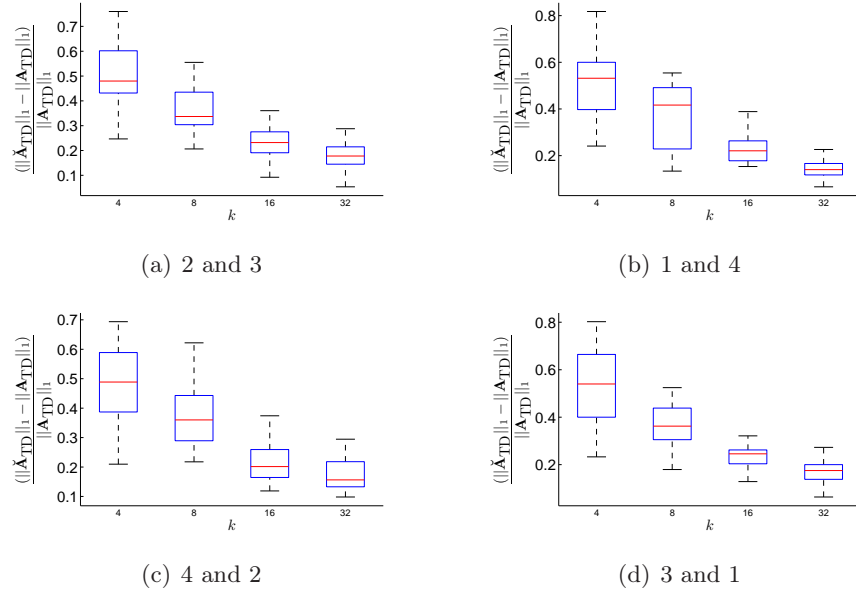


Figure 5.2: Box plots showing the sensitivity of ℓ^1 norm of filter matrices due to a single permutation of two sources.

On each box, the central horizontal line in red colour is the median, the edges of the box in blue colour are the 25th and 75th percentiles. The extreme data points that are not considered outliers are marked by the whiskers in black colour.

Here again, we can notice from the figures that in each draw of the experiment, the quantity $\|\hat{\mathbf{A}}_{\text{td}}\|_1 - \|\mathbf{A}_{\text{td}}\|_1$ is never negative, which means that the ℓ^1 norm increases even when one sub-band is permuted.

Conclusion

From the above experiments, we have gathered an empirical evidence that when the mixing filters are sparse, the sub-band permutations perturb the sparse structure of the filters and they increase the ℓ^1 norm from their true ℓ^1 norm, in the absence of any scaling of sub-bands. Hence, we can use the time-domain sparsity of the filters, quantified using the ℓ^1 norm, as a consistency measure to correct the permutations of the filter sub-bands. Inspired by this, we now propose a combinatorial ℓ^1 minimisation

algorithm to correct the sub-band permutations in the absence of scaling.

5.3 Combinatorial ℓ^1 minimisation for solving permutation ambiguity

The objective of a permutation correction algorithm based on the time-domain sparsity of the filters can be summarised as:

Given the frequency domain estimates of the filter matrix $\check{\mathbf{A}}(f)$, $0 \leq f \leq L$, find a set of permutations $\mathbf{P}(f)$, $0 \leq f < L$ such that the time-domain filters $\tilde{\mathbf{A}}_{\text{td}}$, which are obtained as the Inverse Discrete Fourier Transform (IDFT) of the matrices $\tilde{\mathbf{A}}(f) = \check{\mathbf{A}}(f)\mathbf{P}(f)$, $0 \leq f < L$, has the least ℓ^1 norm.

If there are N sources, then at each sub-band there are $N!$ possible permutations for the sources, and hence there are totally $(N!)^{F/2+1}$ possible permutations (due to the symmetry of the sub-bands) of the filter sub-bands. A straightforward approach to align the permutations is to explore all the possible permutations, compute the ℓ^1 norm for each possible permutation, and choose the set of permutations which yields the minimum ℓ^1 norm. However, this is a very expensive affair. Even for a small problem size like $N = 2$ and $L = 32$, the number of permutations to be explored is 2^{17} , which is already a big number.

We now propose our algorithm to minimise the ℓ^1 norm of the filter matrix. The algorithm which operates iteratively on each frequency at a time and chooses a permutation which has the least ℓ^1 norm among possible permutations. The proposed algorithm is based on the *greedy* principle, where in at each step of the algorithm a permutation is chosen to locally minimise the ℓ^1 norm.

5.3.1 Algorithm

Let denote the set of all the possible source permutations by \mathcal{P} ($|\mathcal{P}| = N!$). Starting with the first sub-band till the last, at each sub-band the sources are permuted using every possible $\mathbf{P} \in \mathcal{P}$, keeping the other sub-bands fixed. For each of these explored source permutations, the ℓ^1 norm of the filter matrix is computed and that permutation which minimises the ℓ^1 norm is retained and declared as the optimal permutation for that particular sub-band. This ensures that ℓ^1 -norm of the filter matrix is lowered to the local minimum at each step by aligning one particular sub-band. At the end of the first sweep through all the sub-bands, the norm of the filter matrix would be less than what was started with. However, as the sub-bands are locally examined, the resulting norm may not be the global minimum. Hence, the entire process of sweeping through the sub-bands is repeated until the difference in the filter norms between two successive sweeps is less than a certain threshold ϵ . The pseudocode of the algorithm is presented in Algorithm 1.

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

Algorithm 1: Algorithm to solve the permutation indeterminacy by minimizing the ℓ^1 -norm of the time domain filter matrix.

Input: $\check{\mathbf{A}}(f)$, $0 \leq f < L$ and ϵ : The estimated sub-band coefficients and a threshold

Output: $\tilde{\mathbf{A}}(f)$, $0 \leq f < L$: The sub-band coefficient matrix after solving for the permutations

```

1 (1) Initialize ;
2    $\tilde{\mathbf{A}}(f) \leftarrow \check{\mathbf{A}}(f)$ ,  $0 \leq f < L$ ;
3 (2) Update all the sub-bands;
4   foreach  $f = 1 : L$  do
5      $old\tilde{\mathbf{A}}(f) \leftarrow \tilde{\mathbf{A}}(f)$ ;
6     foreach  $\mathbf{P} \in \mathcal{P}$  do
7        $\tilde{\mathbf{A}}(f) \leftarrow \check{\mathbf{A}}(f)\mathbf{P}$ ;
8        $\tilde{\mathbf{A}}_{td} \leftarrow IDFT(\tilde{\mathbf{A}}(f'), 0 \leq f' < L)$ ;
9        $val(\mathbf{P}) \leftarrow \|\tilde{\mathbf{A}}_{td}\|_1$ ;
10     $\mathbf{P}(f) \leftarrow \arg \min_{\mathbf{P} \in \mathcal{P}} val(\mathbf{P})$ ;
11     $\tilde{\mathbf{A}}(f) \leftarrow \check{\mathbf{A}}(f)\mathbf{P}(f)$ ;
12 (3) Test if the algorithm should stop;
13   if  $\|\tilde{\mathbf{A}}_{td}\|_1 \geq \{\|old\tilde{\mathbf{A}}_{td}\|_1 - \epsilon\}$  then Output  $\tilde{\mathbf{A}}(f)$ ,  $0 \leq f < L$ ;
14   else Go to step (2)
```

5.3.2 Complexity of the algorithm

To evaluate the ℓ^1 norm of the filters, the frequency domain filter coefficients have to be transformed back into the time domain by $IDFT$ and for each filter, the cost of IDFT through a Fast Fourier Transform (FFT) is $L \log L$. There are MN filters totally and hence the cost of ℓ^1 norm evaluation for a given configuration of sub-bands is $MNL \log L$.

The brute force approach described in Sec. 5.3 to solve the ℓ^1 -minimisation problem would then need $(N!)^L MNL \log L$ operations. In the case of Algorithm 1, each sweep needs to explore $N!L$ permutations and hence, the complexity of each sweep is $N!MNL^2 \log L$. This is still expensive because the computational cost grows in factorial with the number of sources and in square with the filter length, but it is tractable

for small problem sizes and very efficient compared to the brute force approach.

In the next section, we shall demonstrate the working of the algorithm with experimental results.

5.3.3 Experimental results

In this section we shall present the experimental results that evaluate the performance of the proposed algorithm. Firstly, we will consider an ideal situation where in the frequency domain estimates of the filters are noiseless and perfect, but only the sub-bands are permuted, then we will consider a scenario where the estimates are assumed to be imperfect and this is modelled by an additive noise term to the filter estimates in the frequency domain.

5.3.3.1 Recovery in noiseless condition

In this set of experiments, we assume that the sub-band coefficients are estimated without noise and without scaling ambiguity. The filter generation protocol and the way in which the permuted matrices $\check{\mathbf{A}}(f)$, $0 \leq f < L$ are obtained for a given Ω is the same as described for the first experiment in Sec. 5.2.

50 filter matrices with $N = 3$, $M = 2$ and $L = 1024$ for different values of sparsity k are created, transformed and sub-bands are randomly permuted with various sizes of $|\Omega|$ and fed as the input to the proposed algorithm.

The output of the algorithm is transformed back to the time domain to compute the reconstruction error. The following definition of output Signal-to-Noise Ratio (SNR) is used for performance evaluation.

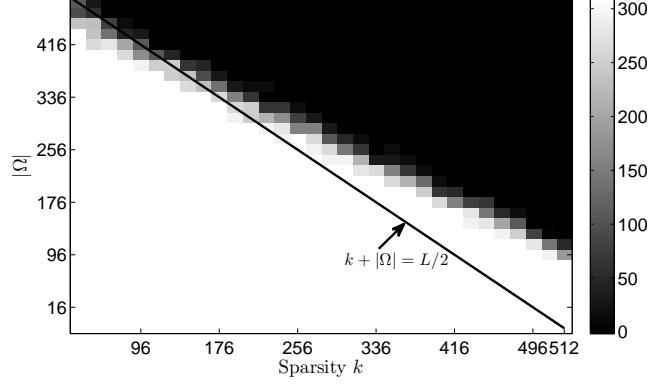
$$SNR_{out} = 20 \log_{10} \left(\frac{\|\mathbf{A}_{td}\|_2}{\|\mathbf{A}_{td} - \tilde{\mathbf{A}}_{td}\|_2} \right). \quad (5.13)$$

The recovered filters \mathbf{A}_{td} suffer from a global permutation ambiguity and hence it has to be taken into account while computing the output SNR. The output SNR is computed with all the possible ($N!$) global permutations of \mathbf{A}_{td} , and the permutation which yields minimum error is retained and then the corresponding output SNR is taken into account for analysis purposes.


Fig 5.3 shows the performance of the algorithm in the noiseless condition in terms of average output SNR of the recovered filters and the number of sweeps the algorithm needs for convergence. Fig. 5.3(a) shows the phase transition diagram of the filter recovery with respect to the filter sparsity k and the number of sub-bands permuted $|\Omega|$. The white region represents a very high SNR of the recovered filters, indicating success and the black region corresponds to very low (almost zero) output SNR which indicates recovery failure. The transition is quite sharp as in there is not much of gray region. When the sparsity k and the number of sub-bands permuted $|\Omega|$ are simultaneously high, the recovery fails.

The thick dark diagonal line on the figure is the straightline given by $k + |\Omega| = L/2$. As a rule of thumb we can say that the recovery is successful if $k + |\Omega| \leq L/2$. With

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?



(a) Average output SNR versus $|\Omega|$ for different values of sparsity k .



(b) Average number of sweeps versus $|\Omega|$ for different values of sparsity k .

Figure 5.3: Performance of Algorithm 1 in terms of average output SNR and number of sweeps.

$|\Omega| > L/4$, the recovery fails when $k + |\Omega| > L/2$. For $|\Omega| < L/16$, the recovery is successful even when $k = L/2$. In chapter 6 we present a theoretical analysis of the effect of permutations on filter sparsity which throws light on the recoverability conditions.

Fig. 5.3(b) shows the average number of sweeps that is needed for the algorithm to converge with the convergence threshold ϵ set to 10^{-4} . We can observe that for smaller values of k and $|\Omega|$, the algorithm takes at most two sweeps to converge. As expected, the number of sweeps increases for larger values of k and $|\Omega|$.

Conclusion

These experimental results in noiseless condition support our claim that we can correct the permutation ambiguity of the filter matrices by minimising the ℓ^1 norm of the filters. However, the filter estimation is not always perfect and hence the robustness of the algorithm is very important. In the next section, let us study how the algorithm performs in the presence of noise.

5.3.3.2 Robustness to noise

The estimation of $\check{\mathbf{A}}(f)$ by an actual BSS algorithm invariably involves some level of noise (as well as scaling, which we do not deal with here). Hence, the permutation solving algorithm needs to be robust to certain level of noise. In this section, we present experiments results that establish the effectiveness of sparsity criterion and the proposed algorithm for solving the permutation problem when the estimates are not perfect.

The following model was used to introduce noise to the filter sub-bands:

$$\check{\mathbf{A}}(f) = \hat{\mathbf{A}}(f)\mathbf{P}(f) + \hat{\mathbf{V}}(f), \quad (5.14)$$

where $\hat{\mathbf{V}}(f)$ is a $M \times N$ matrix having entries that are i.i.d. complex Gaussian with mean zero and variance σ^2 .

The filter generation protocol and the way in which the permuted matrices $\check{\mathbf{A}}(f)$, $0 \leq f < L$ are obtained for a given Ω is same as described for the first experiment in Sec. 5.2. The only change being the addition of noise in each sub-band. The extent of noise added is measured with the following definition of oracle input SNR:

$$SNR_{in}^{oracle} = 20 \log_{10} \left(\frac{\|\mathbf{A}_{td}\|_2}{\|\mathbf{V}_{td}\|_2} \right). \quad (5.15)$$

where \mathbf{V}_{td} is the noise in the time-domain obtained by taking the inverse DFT of $\hat{\mathbf{V}}(f)$, $0 \leq f < L$. The input SNR defined by

$$SNR_{in} = 20 \log_{10} \left(\frac{\|\check{\mathbf{A}}_{td}\|_2}{\|\hat{\mathbf{A}}_{td} - \mathbf{A}_{td}\|_2} \right). \quad (5.16)$$

is actually much larger than the oracle input SNR. That is $SNR_{in} \gg SNR_{in}^{oracle}$.

Note that the output SNR is measured in the time domain, and hence it becomes easier if the input SNR is also measured in the time domain even though the noise added is in the frequency domain. For a specified input SNR and a specific realisation of the filter matrix \mathbf{A}_{td} the noise power is computed using Eq. (5.15). Then the time-domain noise signal \mathbf{V}_{td} is generated with i.i.d. Gaussian distribution with zero mean and unit variance and it is scaled to match the required noise power. This noise signal is then transformed into the frequency domain to obtain $\hat{\mathbf{V}}(f)$, $0 \leq f < L$ and added to the frequency domain filter matrices.

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

Results

50 filter matrices with $N = 3$, $M = 2$ and $L = 1024$ for different values of sparsity k are created, transformed and sub-bands are randomly permuted with various sizes of $|\Omega|$. Noise is then added to the frequency domain filters such that the input SNR is made to vary between -30 dB and 30 dB in steps of 10 dB. These frequency domain filter matrices are fed as input to the proposed algorithm.

The output of the algorithm is transformed back into the time domain and the reconstruction error is computed using Eq. 5.13 after taking care of the global permutation.

Fig 5.4 shows the performance of the algorithm in the noisy settings. Fig. 5.4(a) (resp. 5.4(c) and 5.4(e)) shows the plot of average output SNR (Eq. (5.13)) versus input SNR for $|\Omega| = 8$ (resp. $|\Omega| = 40$ and $|\Omega| = 64$) for various filter sparsity values k .

We observe that $SNR_{out} = SNR_{in}^{oracle}$ within 0.01 dB when $-30 \text{ dB} \leq SNR_{in}^{oracle} \leq 30 \text{ dB}$. This means that the output SNR closely follows the input SNR, and the algorithm can recover the filters to an output SNR level as good as the input SNR.

Fig. 5.4(b) (resp. 5.4(d) and 5.4(f)) shows the average number of sweeps that is needed for the algorithm to converge against the input SNR level for $|\Omega| = 8$ (resp. $|\Omega| = 40$ and $|\Omega| = 64$) with the convergence threshold ϵ set to 10^{-4} . We can observe that for low SNR levels, the algorithm takes more number of sweeps to converge compared to high levels of SNR.

5.4 Summary

In this chapter, we dealt with the problem of permutation and scaling ambiguities that arises in the context of frequency domain convolutive blind source separation based on the narrowband approximation approach. We presented a brief survey of approaches found in the literature to solve these problems explicitly. Then, we assessed the use of the time-domain sparsity of the mixing filters as a consistency measure to correct the permutations in the absence of scaling. Our experiments showed that the permutations can be corrected by minimising the ℓ^1 norm of the filters.

Motivated by the empirical evidence of the connections between ℓ^1 norm and filter permutations, we proposed a combinatorial ℓ^1 norm minimisation algorithm for permutation correction. We experimentally evaluated the algorithm under noiseless and noisy conditions, and demonstrated the ability of the algorithm to correct the permutation ambiguity effectively.



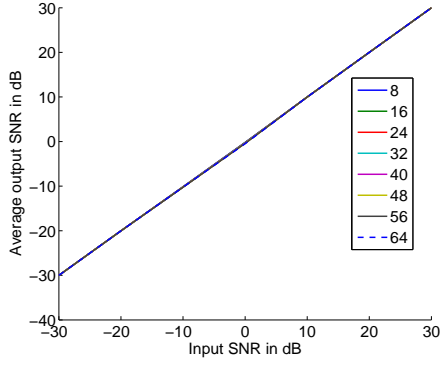
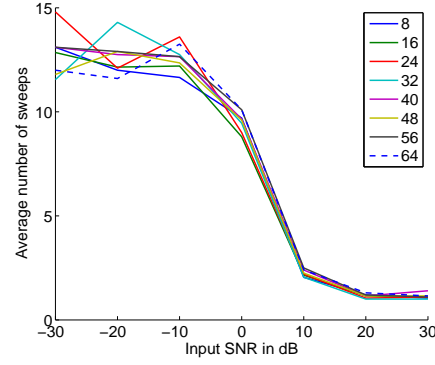


(a) Output Vs input SNR for $|\Omega| = 8$.(b) Average number of sweeps versus input SNR for $|\Omega| = 8$.(c) Output Vs input SNR for $|\Omega| = 40$.(d) Average number of sweeps versus input SNR for $|\Omega| = 40$.(e) Output Vs input SNR for $|\Omega| = 64$.(f) Average number of sweeps versus input SNR for $|\Omega| = 64$.

Figure 5.4: Performance of Algorithm 1 in terms of average output SNR and number of sweeps under noisy settings.

5. CONVOLUTIVE SOURCE LOCALISATION: COULD WE EXPLOIT FILTER SPARSITY?

Chapter 6

Permutations and Sparsity: Theory

Chapter 5 was devoted to the empirical demonstration of the connections between the time-domain sparsity of the filters and sub-band permutations and the possibility of using of sparsity of filters as a consistency criterion for solving the permutation problem in the absence of scaling. This chapter presents some theoretical connections between sparsity and permutations.

The experiments presented in Sec. 5.2 show that for a set of filters with a given sparsity level k , the ℓ^1 norm of the filters defined in Eq. (5.5) increases only if the size of the set of frequencies where the permutations occur $|\Omega|$ is smaller than a certain number. In this chapter, our objective is to understand analytically how the filter sparsity and number of permuted frequencies are related to ℓ^1 norms of the filters before and after permutations.

6.1 Notations

To simplify the analysis let us consider the single channel setting $i = 1$ with N sources. Let $\mathbf{a}_j \in \mathbb{C}^L$, $1 \leq j \leq N$ be the true filter vectors with sparsity less than or equal to k with support sets Γ_j . Let $\hat{\mathbf{a}}_j = \mathbf{F}^* \mathbf{a}_j$, $1 \leq j \leq N$, be the true frequency-domain filter vectors \mathbf{a}_j , where $\mathbf{F}^* \in \mathbb{C}^{L \times L}$, is the forward Fourier matrix.

Let $\check{\mathbf{a}}_j$, $1 \leq j \leq N$ be the frequency-domain filters obtained after permuting the frequency-domain true filters at frequencies given by the set Ω and $\tilde{\mathbf{a}}_j = \mathbf{F} \check{\mathbf{a}}_j$ be the corresponding time-domain filters where $\mathbf{F} \in \mathbb{C}^{L \times L}$, is the inverse Fourier matrix.

We are interested in knowing for what values of k and $|\Omega|$ we can expect the ℓ^1 norm $\sum_{j=1}^N \|\tilde{\mathbf{a}}_j\|_1$ to be greater than $\sum_{j=1}^N \|\mathbf{a}_j\|_1$.

6.2 Main results

The first main result is related to the setting where the supports of the filters are pairwise disjoint. In such a case the ℓ^p , $0 \leq p \leq 1$ norm of the filters obtained after permutations is bounded below by the ℓ^p , $0 \leq p \leq 1$ norm of the true filters, irrespective of the size of the permuted frequency set.

6. PERMUTATIONS AND SPARSITY: THEORY

Theorem 6.1. *Let $\mathbf{a}_j \in \mathbb{R}^L$, $1 \leq j \leq N$ with support sets Γ_j . Assume $\Gamma_i \cap \Gamma_j = \emptyset$ for $i \neq j$, then for $0 \leq p \leq 1$, we have $\sum_{j=1}^N \|\tilde{\mathbf{a}}_j\|_p^p \geq \sum_{j=1}^N \|\mathbf{a}_j\|_p^p$.*

If the value of sparsity k is very small compared to the length of the filters L , and if there are only a few filters (i.e. small N), then the chance of having filters with disjoint supports is high and hence Theorem 6.1 is relevant in such situations.

The second main result is related to ℓ^0 norm of the vectors in a two filter setting.

Theorem 6.2. *Let $\mathbf{a}_j \in \mathbb{R}^L$, $1 \leq j \leq 2$ with $\|\mathbf{a}\|_0 \leq k$ with support sets Γ_j and let Ω be the set of frequencies where the permutations occur.*

(a) *If $|\Omega| < L/2k$ then $\sum_{j=1}^2 \|\tilde{\mathbf{a}}_j\|_0 > \sum_{j=1}^2 \|\mathbf{a}_j\|_0$.*

(b) *If $|\Omega| = L/2k$ and $\Gamma_1 \cap \Gamma_2 = \emptyset$, then $\sum_{j=1}^2 \|\tilde{\mathbf{a}}_j\|_0 \geq \sum_{j=1}^2 \|\mathbf{a}_j\|_0$, with equality iff either i) $\tilde{\mathbf{a}}_j = \mathbf{a}_j$, $\forall j$, or ii) $\tilde{\mathbf{a}}_1 = \mathbf{a}_2$ and $\tilde{\mathbf{a}}_2 = \mathbf{a}_1$.*

Part (a) of Theorem 6.2 gives a condition for the sum of ℓ^0 norms of the new vectors to be strictly greater than the sum of the ℓ^0 norms of the true vectors, irrespective of the supports of the filters. Part (b) tells that if the supports are disjoint and when the sum of ℓ^0 norms of the new vectors is equal to the sum of the ℓ^0 norms of the true vectors, then it implies a global permutation of the filters.

Experimental results in Sec. 5.3.3.1 showed that under noiseless conditions sparse filters can be successfully recovered when $k + |\Omega| \leq L/2$, but theorem 6.2 states that recovery is possible when $k \cdot |\Omega| \leq L/2$. This theoretical result is pessimistic compared to the experimental observations.

6.3 Proofs

In this section, we present the proofs of the stated theorems.

6.3.1 Proof of Theorem 6.1

Consider the vector $\sum_{j=1}^N \check{\mathbf{a}}_j$. This vector has the same discrete Fourier transform (DFT) as the vector $\sum_{j=1}^N \hat{\mathbf{a}}_j$. This is due to the fact that the vectors $\check{\mathbf{a}}_j$ and $\hat{\mathbf{a}}_j$ are related through permutations at certain frequency indices. Hence,

$$\sum_{j=1}^N \check{\mathbf{a}}_j = \sum_{j=1}^N \hat{\mathbf{a}}_j$$

Therefore, the corresponding time-domain vectors are also the same. That is

$$\sum_{j=1}^N \tilde{\mathbf{a}}_j = \sum_{j=1}^N \mathbf{F} \check{\mathbf{a}}_j = \mathbf{F} \left(\sum_{j=1}^N \check{\mathbf{a}}_j \right) = \mathbf{F} \left(\sum_{j=1}^N \hat{\mathbf{a}}_j \right) = \sum_{j=1}^N \mathbf{F} \hat{\mathbf{a}}_j = \sum_{j=1}^N \mathbf{a}_j$$

Hence the p^{th} power of their ℓ^p norms are equal:

$$\begin{aligned} \left\| \sum_{j=1}^N \tilde{\mathbf{a}}_j \right\|_p^p &= \left\| \sum_{j=1}^N \mathbf{a}_j \right\|_p^p \\ &= \sum_{j=1}^N \|\mathbf{a}_j\|_p^p \quad (\text{Due to disjoint supports}) \end{aligned} \quad (6.1)$$

By generalised Triangle inequality, we also have

$$\sum_{j=1}^N \|\tilde{\mathbf{a}}_j\|_p^p \geq \left\| \sum_{j=1}^N \tilde{\mathbf{a}}_j \right\|_p^p, \quad 0 \leq p \leq 1 \quad (6.2)$$

Therefore, by Eqs. (6.1) and (6.2) we have

$$\sum_{j=1}^N \|\tilde{\mathbf{a}}_j\|_p^p \geq \sum_{j=1}^N \|\mathbf{a}_j\|_p^p.$$

6.3.2 Proof of Theorem 6.2

To prove Theorem 6.2 we make use of the following lemma from Elad and Bruckstein (Theorem 1, [37]), which characterises the ℓ^0 norms of a vector in Dirac and Fourier representations.

Lemma 6.1 (Elad and Bruckstein). *Let $\Delta \in \mathbb{C}^L$ and $\mathbf{F} \in \mathbb{C}^{L \times L}$ be the inverse Fourier matrix. Then, a) we have the uncertainty principle $\|\Delta\|_0 \|\mathbf{F}\Delta\|_0 \geq L$ and b) equality implies the non-zero values of Δ have constant magnitude:*

$$\exists c, \Gamma \quad |\Delta_l| = \begin{cases} 0 & \text{if } l \in \Gamma, \\ c & \text{if } l \notin \Gamma. \end{cases}$$

In the case of two vectors, a permutation of sub-bands is nothing but just the swapping of the coefficients and hence it is easier to characterise the permutations. The difference between the vectors $\hat{\mathbf{a}}_j$ and $\check{\mathbf{a}}_j$, $j = 1, 2$, are only due to the coefficients swapped at the locations indexed by Ω .

Let $\mathbb{1}_\Omega$ be the characteristic function of Ω :

$$\mathbb{1}_\Omega = \begin{cases} 1 & \text{at } \Omega, \\ 0 & \text{elsewhere;} \end{cases}$$

and therefore the vector that captures the difference between the vectors $\hat{\mathbf{a}}_1$ and $\hat{\mathbf{a}}_2$ at frequency indices given by Ω is

$$\Delta = \text{diag}(\mathbb{1}_\Omega) \cdot (\hat{\mathbf{a}}_2 - \hat{\mathbf{a}}_1) \quad (6.3)$$

6. PERMUTATIONS AND SPARSITY: THEORY

where $\text{diag}(\mathbf{1}_\Omega)$ is a diagonal matrix whose diagonal is the characteristic function $\mathbf{1}_\Omega$. Now, clearly we have

$$\begin{aligned}\check{\mathbf{a}}_1 &= \widehat{\mathbf{a}}_1 + \Delta, \\ \check{\mathbf{a}}_2 &= \widehat{\mathbf{a}}_2 - \Delta.\end{aligned}\tag{6.4}$$

and

$$\begin{aligned}\widetilde{\mathbf{a}}_1 &= \mathbf{a}_1 + \mathbf{F}\Delta, \\ \widetilde{\mathbf{a}}_2 &= \mathbf{a}_2 - \mathbf{F}\Delta.\end{aligned}\tag{6.5}$$

Let $\mathcal{C}_0^{(2)}(\Omega) = \sum_{j=1}^2 \{\|\widetilde{\mathbf{a}}_j\|_0 - \|\mathbf{a}_j\|_0\}$. Let Γ_1 and Γ_2 be the support sets of the vectors \mathbf{a}_1 and \mathbf{a}_2 respectively, and Γ_1^c and Γ_2^c be the complements of the support sets. That is $\Gamma_j^c = \{l : \mathbf{a}_j(l) = 0\}$.

Proof of (a)

By recalling the definition of Δ in Eq. (6.3), and by using Eq. (6.5) we have

$$\begin{aligned}\mathcal{C}_0^{(2)}(\Omega) &= \|\mathbf{a}_1 + \mathbf{F}\Delta\|_0 + \|\mathbf{a}_2 - \mathbf{F}\Delta\|_0 - \|\mathbf{a}_1\|_0 - \|\mathbf{a}_2\|_0 \\ &= \|(\mathbf{a}_1 + \mathbf{F}\Delta)|_{\Gamma_1}\|_0 + \|(\mathbf{F}\Delta)|_{\Gamma_1^c}\|_0 + \|(\mathbf{a}_2 - \mathbf{F}\Delta)|_{\Gamma_2}\|_0 + \|(\mathbf{F}\Delta)|_{\Gamma_2^c}\|_0 \\ &\quad - \|\mathbf{a}_1\|_0 - \|\mathbf{a}_2\|_0 \\ &\geq \|(\mathbf{F}\Delta)_{\Gamma_1^c}\|_0 + \|(\mathbf{F}\Delta)_{\Gamma_2^c}\|_0 - \|(\mathbf{F}\Delta)_{\Gamma_1}\|_0 - \|(\mathbf{F}\Delta)_{\Gamma_2}\|_0 \\ &\quad \text{(By Triangle inequality)} \\ &= 2\|\mathbf{F}\Delta\|_0 - 2\|(\mathbf{F}\Delta)_{\Gamma_1}\|_0 - 2\|(\mathbf{F}\Delta)_{\Gamma_2}\|_0 \\ &\geq 2\frac{L}{|\Omega|} - 2|\Gamma_1| - 2|\Gamma_2| \quad \text{(Using Lemma 6.1)} \\ &\geq 2\frac{L}{|\Omega|} - 4k.\end{aligned}$$

Therefore if $|\Omega| < L/2k$, then $\sum_{j=1}^2 \|\widetilde{\mathbf{a}}_j\|_0 > \sum_{j=1}^2 \|\mathbf{a}_j\|_0$.

Proof of (b)

The trivial case is $\Delta = 0$.

$$\Delta = 0 \implies \mathbf{F}\Delta = 0 \implies \check{\mathbf{a}}_1 = \mathbf{a}_1, \check{\mathbf{a}}_2 = \mathbf{a}_2$$

When $\Delta \neq 0$ then by Lemma 6.1 we have

$$\|\Delta\|_0 \|\mathbf{F}\Delta\|_0 \geq L$$

and by assumption

$$\begin{aligned} |\Omega| \cdot k = \frac{L}{2} &\implies \|\Delta\|_0 = |\Omega| = \frac{L}{2k} \\ &\implies \|\mathbf{F}\Delta\|_0 \geq 2k \end{aligned} \quad (6.6)$$

Since $\mathcal{C}_0^{(2)}(\Omega) = 0$, $\|\tilde{\mathbf{a}}_1\|_0 + \|\tilde{\mathbf{a}}_2\|_0 = \|\mathbf{a}_1\|_0 + \|\mathbf{a}_2\|_0 = 2k$, i.e.

$$\begin{aligned} 2k = \|\mathbf{a}_1\|_0 + \|\mathbf{a}_2\|_0 &= \|(\mathbf{a}_1 + \mathbf{F}\Delta)|_{\Gamma_1}\|_0 + \|(\mathbf{F}\Delta)|_{\Gamma_1^c}\|_0 \\ &\quad + \|(\mathbf{a}_2 - \mathbf{F}\Delta)|_{\Gamma_2}\|_0 + \|(\mathbf{F}\Delta)|_{\Gamma_2^c}\|_0 \end{aligned} \quad (6.7)$$

Moreover, using Eq. (6.6)

$$\begin{aligned} \|(\mathbf{F}\Delta)|_{\Gamma_1^c}\|_0 &= \|(\mathbf{F}\Delta)\|_0 - \|(\mathbf{F}\Delta)|_{\Gamma_1}\|_0 \geq 2k - k = k \\ \|(\mathbf{F}\Delta)|_{\Gamma_2^c}\|_0 &= \|(\mathbf{F}\Delta)\|_0 - \|(\mathbf{F}\Delta)|_{\Gamma_2}\|_0 \geq 2k - k = k \end{aligned}$$

With the above conditions, Eq. (6.7) can be satisfied only when

$$\|(\mathbf{a}_1 + \mathbf{F}\Delta)|_{\Gamma_1}\|_0 = 0 \quad (6.8)$$

$$\|(\mathbf{a}_2 - \mathbf{F}\Delta)|_{\Gamma_2}\|_0 = 0 \quad (6.9)$$

and

$$\begin{aligned} \|(\mathbf{F}\Delta)|_{\Gamma_1^c}\|_0 &= k \\ \|(\mathbf{F}\Delta)|_{\Gamma_2^c}\|_0 &= k \end{aligned}$$

Eqs. (6.8) and (6.9) imply

$$\begin{aligned} \mathbf{F}\Delta|_{\Gamma_1} &= -\mathbf{a}_1|_{\Gamma_1} \\ \mathbf{F}\Delta|_{\Gamma_2} &= \mathbf{a}_2|_{\Gamma_2} \end{aligned}$$

We then have

$$\|(\mathbf{F}\Delta)\|_0 = \|(\mathbf{F}\Delta)|_{\Gamma_1}\|_0 + \|(\mathbf{F}\Delta)|_{\Gamma_1^c}\|_0 = \|-\mathbf{a}_1\|_0 + k = k + k = 2k$$

Since $\Gamma_1 \cap \Gamma_2 = \emptyset$, we have

$$\|(\mathbf{F}\Delta)|_{\Gamma_1} + (\mathbf{F}\Delta)|_{\Gamma_2}\|_0 = \|(\mathbf{F}\Delta)|_{\Gamma_1}\|_0 + \|(\mathbf{F}\Delta)|_{\Gamma_2}\|_0 = \|\mathbf{a}_1\|_0 + \|\mathbf{a}_2\|_0 = 2k = \|(\mathbf{F}\Delta)\|_0$$

So,

$$\mathbf{F}\Delta = \mathbf{F}\Delta|_{\Gamma_1} + \mathbf{F}\Delta|_{\Gamma_2} = -\mathbf{a}_1 + \mathbf{a}_2$$

with

$$\Gamma_1 \cup \Gamma_2 = \text{support}(\mathbf{F}\Delta)$$

Hence,

$$\begin{aligned} \tilde{\mathbf{a}}_1 &= \mathbf{a}_1 + \mathbf{F}\Delta = \mathbf{a}_2, \\ \tilde{\mathbf{a}}_2 &= \mathbf{a}_2 - \mathbf{F}\Delta = \mathbf{a}_1. \end{aligned}$$

6.4 Summary

This chapter focussed on the theoretical connections between filter permutations and sparsity. The first result is regarding filters with disjoint supports and second result is about a special case of two filters and ℓ^0 norm. In comparison with the empirical results from chapter 5, we notice that the theoretical results are pessimistic.

Chapter 7

A convex optimisation framework for single source localisation in sparse convolutive setting

In the blind source separation problem, the objects of primary interest are the sources and the mixing filters. Several properties of these objects (independence, non-negativity, short filters, etc.) are exploited, either individually or jointly, to perform source separation. As it was discussed in chapter 2, sparsity of sources in the time-frequency domain is a property that has been widely used for designing source separation systems. This category of approaches, collectively called by the name *sparse component analysis*, forms an important tool to approach the source separation problem on the one hand.

On the other hand, sparse channel (filter) estimation in the single source, semi-blind setting is a widely studied problem in the context of communication systems. As we discussed in chapter 4, the problem of sparse channel estimation in a semi-blind setting can be formulated as a convex optimisation problem by exploiting the cross-relation that exists across the channels. The motivation for convex optimisation approach for filter estimation stems from the more generic field of *sparse recovery* problems which was discussed in chapter 3.

In this chapter, we shall present a generalised form of the cross-relation between the channels in the time-frequency domain, and formulate an ℓ^1 minimisation problem for filter estimation in the single source setting. We also present experimental results regarding the recovery performance of the proposed approach.

Plan of the chapter

In the following section, we will introduce the time-frequency domain cross-relation and show how to exploit the cross-relation for sparse filter estimation in the single source setting. Subsequently, we will evaluate the filter estimation method based on the time-frequency domain CR for different kinds of data models.

7. A CONVEX OPTIMISATION FRAMEWORK FOR SINGLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

7.1 Time-frequency domain cross-relation

Let us recall the time-domain CR, with a single source signal $s(t)$ of length T and two observations $x_i(t)$, $i = 1, 2$ of the source after being filtered by filters $a_i(t)$ of length L , in the absence of noise. That is

$$x_i(t) = (s \star a_i)(t), \quad i = 1, 2. \quad (7.1)$$

The length of each mixture¹ $x_i(t)$, $i = 1, 2$ is $T + L - 1$. We have

$$(x_2 \star a_1)(t) = (x_1 \star a_2)(t), \quad \forall t.$$

and therefore the following time-domain CR holds:

$$(x_2 \star a_1)(t) - (x_1 \star a_2)(t) = 0, \quad \forall t. \quad (\text{CR-TIME})$$

We shall now see how the time-domain CR ([CR-TIME](#)) can be transformed into the time-frequency domain. As a first approach, one can transform the mixtures $x_i(t)$ into the time-frequency domain and then write down the CR equality. Let us call this as the *narrowband* approximation of the CR.

7.1.1 Narrowband approximation

Let us now consider the time-frequency representation of the mixtures using the Short-Time Fourier Transform (STFT). The definition of the STFT for source signals is given by Eq. (2.7), and the STFT of the mixtures is defined likewise. Let us consider N_T frames of STFT of the mixtures, obtained with a window length of F . Let $\widehat{s}(\tau, f)$ and $\widehat{x}_i(\tau, f)$ denote the STFT coefficients of the source and i^{th} mixture at frame index τ and frequency index f . Then by using the narrowband approximation [26], we have

$$\widehat{x}_i(\tau, f) \approx \widehat{a}_i(f) \cdot \widehat{s}(\tau, f), \quad i = 1, 2, \quad (7.2)$$

where $\widehat{a}_i(f)$ is the DFT coefficient of the filter $a_i(t)$ at frequency index f . Note that the filters do not vary with time and hence the DFT coefficients, of length F , are not dependent on τ .

7.1.2 Time-frequency domain cross-relation

We can write the narrowband time-frequency domain CR (equivalent to the time-domain CR ([CR-TIME](#))) as

$$\widehat{a}_2(f) \cdot \widehat{x}_1(\tau, f) - \widehat{a}_1(f) \cdot \widehat{x}_2(\tau, f) \approx 0, \quad \forall(\tau, f). \quad (\text{CR-NB})$$

Alternatively we have:

¹When there is only one source, $x_i(t)$, $i = 1, 2$ are technically not *mixtures*. However, we will abuse the terminology and call them mixtures.

$$[\hat{x}_2(\tau, f) \quad -\hat{x}_1(\tau, f)] \begin{bmatrix} \hat{a}_1(f) \\ \hat{a}_2(f) \end{bmatrix} \approx 0. \quad (7.3)$$

Let $\hat{\mathbf{x}}_i(\tau) = [\hat{x}_i(\tau, f)]_f$ be the vector corresponding to the STFT coefficients of i^{th} mixture at frame τ at all the frequencies. Also, let $\hat{\mathbf{a}}_i = [\hat{a}_i(f)]_f$ be the DFT vector corresponding to the time domain filter vector \mathbf{a}_i . Then, we have

$$\hat{\mathbf{a}}_i = \mathbf{F}^* \cdot \begin{bmatrix} \mathbf{a}_i \\ \mathbf{0}_{(F-L)} \end{bmatrix}, \quad i = 1, 2, \quad (7.4)$$

where \mathbf{F}^* is the forward Fourier matrix of size $F \times F$ and $\mathbf{0}_{(F-L)}$ is a zero vector of length $(F - L)$.

Using (CR-NB) and (7.4) the time-frequency domain CR can be written in the matrix form as

$$\underbrace{\begin{bmatrix} \text{diag}(\hat{\mathbf{x}}_2(1)) & -\text{diag}(\hat{\mathbf{x}}_1(1)) \\ \text{diag}(\hat{\mathbf{x}}_2(2)) & -\text{diag}(\hat{\mathbf{x}}_1(2)) \\ \vdots & \vdots \\ \text{diag}(\hat{\mathbf{x}}_2(N_T)) & -\text{diag}(\hat{\mathbf{x}}_1(N_T)) \end{bmatrix}}_{\mathbf{B}_{\text{NB}}} \begin{bmatrix} \mathbf{F}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{F}^* \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{0}_{(F-L)} \\ \mathbf{a}_2 \\ \mathbf{0}_{(F-L)} \end{bmatrix} \approx \mathbf{0}, \quad (7.5)$$

where $\text{diag}(\hat{\mathbf{x}}_i(\tau))$ is an operator which maps the vector $\hat{\mathbf{x}}_i(\tau)$ into a diagonal matrix, whose diagonal entries are the elements of the vector $\hat{\mathbf{x}}_i(\tau)$.

For brevity, let's denote the product of the first two matrices in (7.5) by \mathbf{B}_{NB} , and let

$$\mathbf{a} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{0}_{(F-L)} \\ \mathbf{a}_2 \\ \mathbf{0}_{(F-L)} \end{bmatrix}. \quad (7.6)$$

Hence we have

$$\mathbf{B}_{\text{NB}} \cdot \mathbf{a} \approx \mathbf{0}. \quad (7.7)$$

The matrix \mathbf{B}_{NB} is of size $(N_T \cdot F) \times 2F$ and each row of the matrix corresponds to a time-frequency point of the mixture. Hence, the rows of this matrix can be indexed by the ordered pair (τ, f) , $1 \leq \tau \leq N_T$, $0 \leq f \leq F - 1$.

7.1.3 Filter estimation by convex optimisation

If we have a single source, then from Eq. (7.2) it is clear that the mixture coefficient $\hat{x}_i(\tau, f)$ indirectly gives us some information about $\hat{a}_i(f)$, at every frequency f . Thus, the mixture vector in the time-frequency domain $\hat{\mathbf{x}}_i(\tau)$ can be regarded as a linear

7. A CONVEX OPTIMISATION FRAMEWORK FOR SINGLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

measurement of the frequency domain filter coefficient vector \mathbf{a}_i at frame index τ . At each (τ, f) , what we actually observe is $\hat{a}_i(f)$ which is modulated by $\hat{s}(\tau, f)$.

The aim of blind filter estimation is to recover the sparse filters \mathbf{a}_i , $i = 1, 2$ from the mixtures $\mathbf{x}_i(\tau)$, $1 \leq \tau \leq N_T$. We can regard this problem as the task of recovering sparse filters from their indirect linear observations at various frequencies.

Recovering sparse vectors from their frequency information is a well studied problem in the area of compressed sensing. If we are allowed to observe the coefficients $\hat{a}_i(f)$, $i = 1, 2$ directly for all the frequencies f , then recovering the filter is a simple straightforward task of inverting the discrete Fourier transform, even when the filters are not sparse. However, if we have an incomplete set of observations (direct observations on only on a partial set of frequencies Ω), then results from compressed sensing literature [105] show that it is still possible to recover the filter by solving an ℓ^1 minimisation problem, as long as it is sparse enough and the size of the set of observed frequencies Ω are large enough, though incomplete.

Notice that the vector \mathbf{a} is a time domain vector and when the filters \mathbf{a}_i are sparse in the time domain, in the spirit of the sparse recovery problems we propose to solve the following ℓ^1 minimisation problem to estimate the filters.

$$\text{minimize } \|\mathbf{a}\|_1 \quad \text{subject to} \quad \|\mathbf{B} \cdot \mathbf{a}\|_2 \leq \epsilon \quad \text{and} \quad \mathbf{a}_1(t_0) = 1, \quad (7.8)$$

where $\mathbf{B} := \mathbf{B}_{\text{NB}}$. The normalisation $\mathbf{a}_1(t_0) = 1$ is forced to avoid the all-zero trivial solution. It should be noted that this constraint also makes the problem convex and efficient algorithms can be used to solve the problem. In all our experiments, we used the CVX [65] package to solve the ℓ^1 minimisation problem.

For a given sparsity k of a vector, the ℓ^1 minimisation problem can successfully recover the vector provided Ω is big enough. Theoretical and empirical limits of recovery have been established in literature for sparse recovery with incomplete frequency observations [105]. However, in our current setting, we don't observe the filter coefficients directly at various frequencies and we don't use the same formulation of the ℓ^1 minimisation problem as that of direct observations setting. Hence the existing recovery results for direct observations cannot be used to assess the recovery performance in our setting.

Notice the similarity of (7.8) with the ℓ^1 minimisation problem (4.11) corresponding to the time-domain CR. The form of the problem is same but for the way in which the matrix \mathbf{B} is constructed.

7.2 Filter recovery in single source setting: Experiments

In this section, we shall evaluate the filter recovery performance of the narrowband CR approach in the single source setting. The objective of the experiments are twofold:

1. To assess the empirical recovery performance of the narrowband CR based approach and

7.2 Filter recovery in single source setting: Experiments

2. To experimentally demonstrate the possibility of sparse recovery with incomplete information.

By the term *recovery performance* we mean that we are interested in knowing the range of values of sparsity k for which the filters are recovered with a certain output performance level. Hence, the problem complexity is driven by the sparsity level k and the number of observations $|\Omega|$. We shall consider these two factors to be given by the problem setting.

The parameter that influences the performance of the presented approach is the error term ϵ . In the single source case, the only factor that affects the accuracy of the CR is the extent of narrowband approximation, which is governed by the length of the DFT F . For large sizes of F compared to the filter length L , the CR is accurate enough to choose and fix a small value for ϵ .

7.2.1 Data models and performance measure

Before we go on to describe the experimental evaluation of the narrowband CR approach in the single source setting, let us first describe the different data models that we will be using throughout this chapter and Chapter 8. The data generation protocol is generic and we shall reference this section at appropriate places in the rest of the thesis.

7.2.1.1 Sparse filter generation

Each filter \mathbf{a}_i , $i = 1, 2$ of length L is generated to have $k/2$ non-zero coefficients, for various values of k . That is, $\|\mathbf{a}_i\|_0 = k/2$, $i = 1, 2$. The $k/2$ support indices on each channel are chosen uniformly at random in the set $(\frac{L}{4}, \frac{3L}{4})$ (See explanation below). The filter coefficients are generated i.i.d. Gaussian with zero mean, unit variance and sorted to have decreasing magnitudes along the time axis within the support. Every filter $a_1(t)$ is then normalised and shifted to have $a_1(L/2) = 1$. Hence, the vector \mathbf{a} defined in the Eq. (7.6) has k non-zero coefficients.

When $k = 2$, it refers to the case where there is only one peak in each filter and hence it corresponds to anechoic filters. Also, given a solution $\tilde{\mathbf{a}}$ to the convex problem (7.8), arbitrarily shifted and scaled versions of $\tilde{\mathbf{a}}$ are also valid solutions due to the properties of DFT and the formulation of the optimisation problem. The solvers which solve the convex program (7.8) arbitrarily fixes the shift of the solution depending on its non-zero values, and the scaling is fixed by the constraint imposed on the solution. Hence, in order to avoid a rollover of the support of the estimated filters, we restrict the support of the input filters to the set $(\frac{L}{4}, \frac{3L}{4})$.

7.2.1.2 Source generation

We generate the sources according to one of the three models described below.

7. A CONVEX OPTIMISATION FRAMEWORK FOR SINGLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

Source Type 1: Gaussian

The length of the sources is set to $T = 3L$ and to ensure the source activity across all frequencies, they are generated as i.i.d. Gaussian with zero mean and unit variance.

Source Type 2: Sinusoid

For each source j , N_T independent time frames $s_j^\tau(t)$, $1 \leq \tau \leq N_T$ are generated. Each frame $s_j^\tau(t)$ of length $T = 3L$ is a sum of N_F sinusoids:

$$s_j^\tau(t) = \sum_{w=1}^{N_F} A_{jw}^\tau \sin(2\pi f_{jw}^\tau t + \phi_{jw}^\tau), \quad (7.9)$$

where the frequencies f_{jw}^τ are chosen uniformly at random in $[0, 1/2]$. The amplitudes A_{jw}^τ are generated i.i.d. Gaussian with zero mean, unit variance, and the phases ϕ_{jw}^τ are chosen uniformly at random in $[0, 2\pi]$.

The number of sinusoids N_F can be considered as the source sparsity per frame. It should be noted that even though N_F remains fixed across the frames, the frequency content of the sources is varying to provide spectral diversity to the sources, and hence to the observations.

Source Type 3: Gabor

Each source signal of length T is generated as a sum of sinusoids of random durations at random locations with Gaussian envelopes. We choose the name Gabor for this type of source as a mnemonic for referencing in the rest of this thesis.

Given the length of a source T , we first generate a set of window lengths $\{W_{jn}\}_{n=1}^\Lambda$ randomly according to an exponential distribution with parameter λ , such that $\sum_n W_{jn} \leq T$. Then, we set the starting indices $\{t_{jn}\}_{n=1}^\Lambda$ of the windows as

$$\begin{aligned} t_{j1} &= 1, \\ t_{jn} &= t_{j(n-1)} + W_{j(n-1)}. \end{aligned}$$

We generate N_F such sets of window lengths and their corresponding starting indices, for each source j . Hence, we add one more index m , $1 \leq m \leq N_F$, which indexes the set of window lengths, and the number of windows Λ now depends on m . Hence, we have for each source j :

$$\begin{aligned} \left\{ W_{jn}^m \right\}_{n=1, m=1}^{\Lambda_m, N_F} &: N_F \text{ sets of window lengths, each set having } \Lambda_m \text{ windows.} \\ \left\{ t_{jn}^m \right\}_{n=1, m=1}^{\Lambda_m, N_F} &: N_F \text{ sets of starting indices, each set having } \Lambda_m \text{ indices.} \end{aligned}$$

With these parameters, we generate the sources as

7.2 Filter recovery in single source setting: Experiments

$$s_j(t) = \sum_{m=1}^{N_F} \sum_{n=1}^{\Lambda_m} \sin(2\pi f_{jn}^m t + \phi_{jn}^m) \cdot w_{W_{jn}^m}(t - t_{jn}^m) \quad (7.10)$$

where the frequencies f_{jn}^m are chosen uniformly at random in $[0, 1/2]$, the phases ϕ_{jn}^m are chosen uniformly at random in $[0, 2\pi]$ and $w_W(t)$ is a discrete Gaussian window of length W centered at $(W - 1)/2$, defined by

$$w_W(t) = e^{-\frac{1}{2} \left(\frac{t - (W-1)/2}{\sigma(W-1)/2} \right)^2}. \quad (7.11)$$

We use the default value of $\sigma = 0.4$ as set by the Matlab function `gausswin()`. The sources generated using this model will have at most N_F active frequencies at any given time index.

7.2.1.3 Performance measure

In the case of narrowband CR, the unknown vector \mathbf{a} (7.6) is of length $2F$, but each filter that we are looking for is of length $L < F$. Due to the nature of the convex problem at hand, the solution obtained by (7.8) has a shift ambiguity, and is also scaled appropriately to satisfy the normalisation constraint. Eq. (7.12) is the definition of SNR in decibel (dB), which accounts for the scaling and shift ambiguity. The definition of output SNR, which is generic for N sources (which will be used further in our discussion), is used as the recovery performance measure.

$$\text{SNR}_{out} = 10 \log_{10} \left(\frac{\sum_{j=1}^N \|\mathbf{a}^{(j)}\|_2^2}{\sum_{j=1}^N \min_{t', \mu} \sum_t (|a^{(j)}(t) - \mu \cdot \tilde{a}^{(j)}(t - t')|^2)} \right), \quad (7.12)$$

where $\mathbf{a}^{(j)}$ represents the filter vector corresponding to the j^{th} source and $a^{(j)}(t)$ is its coefficient at index t .

It should be highlighted that by observing Eq. (7.6), we can infer that the vector $\tilde{\mathbf{a}}$ that we are looking for has only $2L$ unknowns and is k -sparse, in spite of its actual length being $2F$. Once the estimated vectors are obtained, the output Signal-to-Noise Ratio (SNR) is computed using Eq. (7.12).

7.2.2 Recovery with full set of observations

The first set of experiments concerns the limit of recovery when the observations (mixtures) are available at all frequencies. This is an ideal situation: the recovery performance in this setting puts a limit on the recovery capability of the ℓ^1 minimisation formulation.

7. A CONVEX OPTIMISATION FRAMEWORK FOR SINGLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

Experimental protocol

The sparse filters that are needed for this experiment are generated according to the procedure described in Sec. 7.2.1.1 and the source model used is Gaussian, described in Sec. 7.2.1.2. In all the experiments the filter length is set to $L = 256$. For each value of k between 2 and 60, in steps of 2, 50 sets of the filters and sources are generated.

In each experimental realisation, each mixture \mathbf{x}_i , $i = 1, 2$ of length $T + L - 1$, obtained according to Eq. (7.1) is treated as a single frame and the corresponding frequency domain representation (DFT) $\hat{\mathbf{x}}_i$ is computed. Hence, the total number of frames is $N_T = 1$ and the window length is $F = T + L - 1$. Then, the matrix \mathbf{B}_{NB} is built according to Eq. (7.5) and the ℓ^1 minimisation problem (7.8) is solved to obtain the filter estimate $\tilde{\mathbf{a}}$. In all the experiments, the value of the parameter ϵ is set to 10^{-3} .

Results




Figure 7.1: Average output SNR versus sparsity k

Fig. 7.1 shows the plot of average SNR_{out} as a function of sparsity k , when the filters are recovered using the full set of frequencies. The output SNR drops below 20dB when $k > 30$. Therefore, with the narrowband CR approach, we can recover pairs of filters of length $L = 256$ having up to $k/2 = 15$ non-zero coefficients each (Note that by convention each filter \mathbf{a}_i , $i = 1, 2$ has $k/2$ non-zero coefficients) with more than 20dB output SNR, when the full set of frequencies are available.

7.2.3 Recovery with limited set of observations

In the previous section, we provided the recovery results in a case where the source is active in all the frequencies and hence the filters can be observed at all possible frequencies. However, this might not be the case always and hence we might not be able to observe the filters at all the frequencies. We shall now describe this with an example.

Fig. 7.2(a) shows the magnitude of the STFT $\{|\hat{s}(\tau, f)|\}_{(\tau, f)}$ of a source signal in the deciBel (dB) scale. The blue region represents the source inactivity and we can notice that the source is inactive for most of the time-frequency regions.

Fig 7.2(b) shows the magnitude of the STFT $\{|\hat{x}_i(\tau, f)|\}_{(\tau, f)}$ of a mixture obtained by convolving the source with a filter \mathbf{a}_i in the dB scale. The right hand side of the figure symbolically shows the frequencies along which the observation are available. The shaded boxes indicate the presence of observations and the white boxes indicate the *missing* frequencies.

We now present experimental results that show the recovery performance of the time-frequency domain CR based method varies with respect to $|\Omega|$, the size of the set of frequencies over which the observations are available.

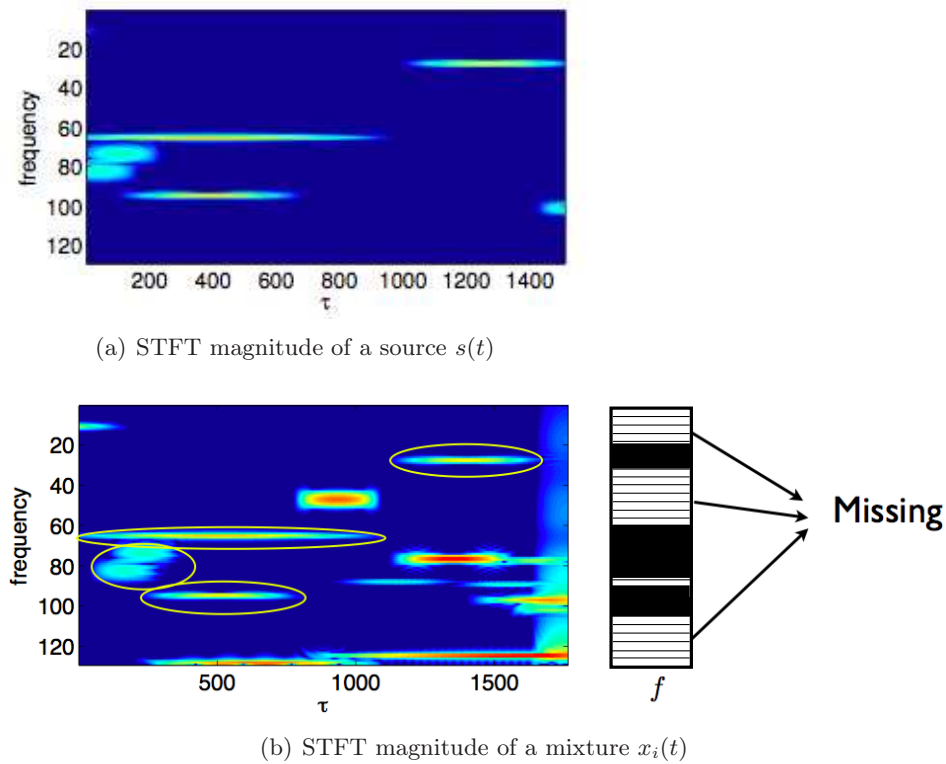


Figure 7.2: Illustration of a mixture with limited activity of the source in the time-frequency domain.

7. A CONVEX OPTIMISATION FRAMEWORK FOR SINGLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

Experimental protocol

The sparse filters that are needed for this experiment are generated according to the procedure described in Sec. 7.2.1.1 and the source model used is Gaussian, described in Sec. 7.2.1.2. In all the experiments the filter length is set to $L = 256$. For each value of k between 2 and 60, in steps of 2, 50 sets of the filters and sources are generated.

In each experimental realisation the mixtures \mathbf{x}_i , $i = 1, 2$ of length $T + L - 1$, obtained according to Eq. (7.1), are treated as a single frame and the corresponding frequency domain representation (DFT) $\hat{\mathbf{x}}_i$ is found. Hence, the total number of frames is $N_T = 1$ and the window length is $F = T + L - 1$. Then, the matrix \mathbf{B}_{NB} is built according to Eq. (7.5). Recall from Sec. 7.1 that the rows of this matrix can be indexed by the ordered pair (τ, f) . Since we have only one frame of the mixtures, the matrix \mathbf{B}_{NB} is of size $F \times 2F$ and the row indices are $\{(1, f)\}_{0 \leq f \leq F-1}$.

In order to recover the filters using observations over only a subset of frequencies of a given size, we randomly generate the set $\omega \subset \{0, 1, \dots, F-1\}$. And, given a set of frequencies ω , let $\Omega = \{(1, f) : f \in \omega\}$ be the set of time-frequency indices and $\mathbf{B}_{\text{NB}}^\Omega$ be the restriction of matrix \mathbf{B}_{NB} to the rows indexed by the set Ω . The estimate $\tilde{\mathbf{a}}$ of the filters are obtained by solving the convex program (7.8) using $\mathbf{B} = \mathbf{B}_{\text{NB}}^\Omega$. In all the experiments, the value of the parameter ϵ is set to 10^{-3} .

Results

Fig. 7.3 shows the transition diagram of the average recovery performance in SNR as a function of the size of the set of observed frequencies Ω relative to the number of unknowns $2L$ in percentage and the sparsity of the filters.

For sparsity levels $k \leq 26$, the output SNR is at least 20dB when the number of observations is greater than 30% of $2L$. For sparsity levels $k > 40$, the output SNR is close to zero even when the relative size of observations is 90%, and this is consistent with the plot obtained for full set of observations on Fig. 7.1. As the size of the set Ω decreases, the recovery performance also gradually drops. Only for $k = 2$, which corresponds to anechoic setting, the performance is very good even when the relative size of Ω is as low as 10%.

7.3 Summary

This chapter focussed on the development of a time-frequency domain CR based on the narrowband approximation of the mixtures and the usage of this time-frequency domain CR in filter estimation. We empirically demonstrated the ability of an ℓ^1 minimisation approach to recover sparse filters from their indirect observations.

In the experiments we presented Sec. 7.2.3 concerning the recoverability of filters using limited observations, the experimental setup was such that the observations were discarded deliberately in spite of its availability. This was actually to empirically demonstrate the possibility of sparse filter recovery even with limited information. As the




Figure 7.3: Transition diagram for sparse filter recovery as a function of $|\Omega|/2L$ and sparsity k .

experiments suggested, we can recover filters of sparsity $k \leq 30$ with more than 20dB output SNR even with small set of observations.

Further, the sources can actually be sparse in the time-frequency domain and hence we might not be able to observe the filters at all the frequencies. But, thanks to the sparsity of filters, it is still possible for us to recover them with limited information.

In the next chapter, we shall consider the multiple source setting and describe how the sparsity of sources can help us to formulate multiple ℓ^1 minimisation problems of the form (7.8) for each pair of filters corresponding the sources. These ℓ^1 minimisation problems can in turn be solved to estimate the filters.

7. A CONVEX OPTIMISATION FRAMEWORK FOR SINGLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

Chapter 8

A convex optimisation framework for multiple source localisation in sparse convolutive setting

In this chapter we present an approach to generalise the time-frequency domain CR based method for filter estimation presented in Chapter 7 to the multiple source setting by exploiting the time-frequency domain sparsity of the sources.

We shall recall that sparse component analysis for blind source separation task in linear-instantaneous and anechoic settings largely utilises the sparsity property of the sources, sparse channel estimation using sparse recovery techniques rely on the sparsity of the channels.

This chapter proposes an approach to combine these two different notions of sparsity and presents a unified framework for multiple filter estimation in the multiple source setting. The proposed framework consists of two stages: 1) a time-frequency points clustering stage where the source sparsity is exploited and 2) an ℓ^1 minimisation formulation for filter estimation where the filter sparsity is exploited.

Plan of the chapter

In the first part of the chapter we describe the basic idea about how to exploit source sparsity in the time-frequency domain to formulate a CR for each of the sources involved and then present a two-stage framework for filter estimation. We then present experimental results which evaluate the filter estimation stage of the framework by assuming that the time-frequency clusters are available as side information.

We then develop a time-frequency domain cross-relation that is not based on the narrowband approximation of the mixtures. We call this a *wideband* cross-relation and we describe how to use the wideband CR for multiple filter estimation. We present the experimental results which evaluate the filter estimation stage by assuming the availability of clusters by some side information, as it is done for the narrowband CR experiments.

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

In the last part of this chapter, we shall consider a special setting of only two sources where only one of them is mixed through a convolutive filter and the remaining source is mixed in an anechoic fashion. In such a setting, we show that it is possible to blindly obtain the clusters and they can be subsequently used for filter estimation. We will present the filter recovery results for this particular setting.

8.1 General idea

In chapter 7, the time-frequency domain cross-relation (CR-NB) played the key role in the formulation of the ℓ^1 minimisation problem (7.8) for filter estimation.

When we consider a general case of ($N \geq 2$) sources, the mixtures x_i in the time-frequency domain using the narrowband approximation will be of the form

$$\hat{x}_i(\tau, f) \approx \sum_{j=1}^N \hat{a}_{ij}(\tau, f) \cdot \hat{s}_j(\tau, f), \quad i = 1, 2, \quad (8.1)$$

at each time frequency point (τ, f) . Therefore the narrowband CR (CR-NB) is not valid anymore due to the contribution of multiple sources.

In such cases, the time-frequency domain sparsity and disjointness of the sources enables enable us to formulate the time-frequency domain CR locally at points where there is at most one source is active. These points where the CR is satisfied can then be grouped together according to the sources and subsequently used to estimate the filters by formulating and solving a convex problem of the form (7.8) for each source.

Let us recall the *approximate w-disjoint orthogonality* [4, 32] hypothesis of the sources introduced in the Sec. 2.3.1.1. Two sources $s_1(t)$ and $s_2(t)$ are said be approximately *w-disjoint orthogonal* when

$$\hat{s}_1(\tau, f) \hat{s}_2(\tau, f) \approx 0, \quad \forall \tau, f. \quad (8.2)$$

Approximate *w-disjoint orthogonality* of sources implies that at each time-frequency point, there is *at most only one dominant source*. Let Ω_j , $1 \leq j \leq N$ be the set of time-frequency points where only source j is dominant, then by using the narrowband approximation of the mixing equations we have

$$\hat{x}_i(\tau, f) \approx \hat{a}_{ij}(f) \cdot \hat{s}_j(\tau, f), \quad \forall (\tau, f) \in \Omega_j, i = 1, 2. \quad (8.3)$$

Hence, when a time-frequency point belongs to a set Ω_j , then it satisfies the narrowband CR (CR-NB) for source j . Even though the time-domain CR breaks down due to the presence of multiple sources, the narrowband CR is still satisfied at certain time-frequency points due to the time-frequency domain disjointness of the sources.

If the sets of time-frequency points Ω_j , $1 \leq j \leq N$ are known to us, then we can build the matrices $\mathbf{B} = \mathbf{B}_{\text{NB}}^{\Omega_j}$ for each source j as the restrictions of the matrix \mathbf{B}_{NB} defined in (7.5), to the rows indexed by the sets of time-frequency points Ω_j . Then these matrices can be used in the ℓ^1 minimisation problem (7.8) in order to estimate

the filters $\tilde{\mathbf{a}}^{(j)}$, $1 \leq j \leq N$. The filter vector $\tilde{\mathbf{a}}^{(j)}$ has the same structure as in (7.6), but with an extra index to the sources.

Fig. 8.1 illustrates the idea pictorially. On the first column of the figure are the STFT magnitude plots of two sources which are sparse in the time-frequency domain in dB scale. The yellow ellipses show the time-frequency regions where only source 2 is active and source 1 is inactive. The second column shows the STFT magnitude of the mixtures. The yellow ellipses in the mixtures correspond to the second source, and hence the corresponding time-frequency indices belong to the set Ω_2 . With this set, we can form the matrix $\mathbf{B}_{\text{NB}}^{\Omega_j}$ and use it to estimate the filters corresponding to source number 2.




Figure 8.1: Illustration of the time-frequency domain sparsity and disjointness of sources.

Let us henceforth refer to the sets of time-frequency points Ω_j , $1 \leq j \leq N$ as *clusters* and the process of identifying them as *clustering*.

8.2 A two stage framework

Sec. 8.1 described the way in which we can exploit the narrowband approximation and time-frequency domain sparsity of sources to formulate the ℓ^1 minimisation problem for filter estimation in the multiple source setting. A key step in this approach is the identification of time-frequency points where the cross-relations are satisfied. The sparse filter estimation process can now be summarised in the following steps.

1. Compute the time-frequency representations $\hat{\mathbf{x}}_i$, $i = 1, 2$.
2. For each source j ,

P1 { (a) Identify the set Ω_j .

P2 { (b) Build the matrix $\mathbf{B} := \mathbf{B}_{\text{NB}}^{\Omega_j}$.

(c) Solve (7.8) with ϵ to obtain the estimated filter $\tilde{\mathbf{a}}^{(j)}$.

8: A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

It is clear that the framework consists of two distinct stages: 1) a time-frequency points clustering stage for each source and 2) a filter estimation stage by forming the cross-relation and solving a convex optimisation problem. Fig. 8.2 shows the block diagram of the filter estimation framework that we have proposed. The block **P1** performs the clustering step by exploiting the time-frequency sparsity of the sources. Once the clusters are found, then **P2** formulates and solves an ℓ^1 minimisation problem of the form in (7.8) for each source to estimate the filters.

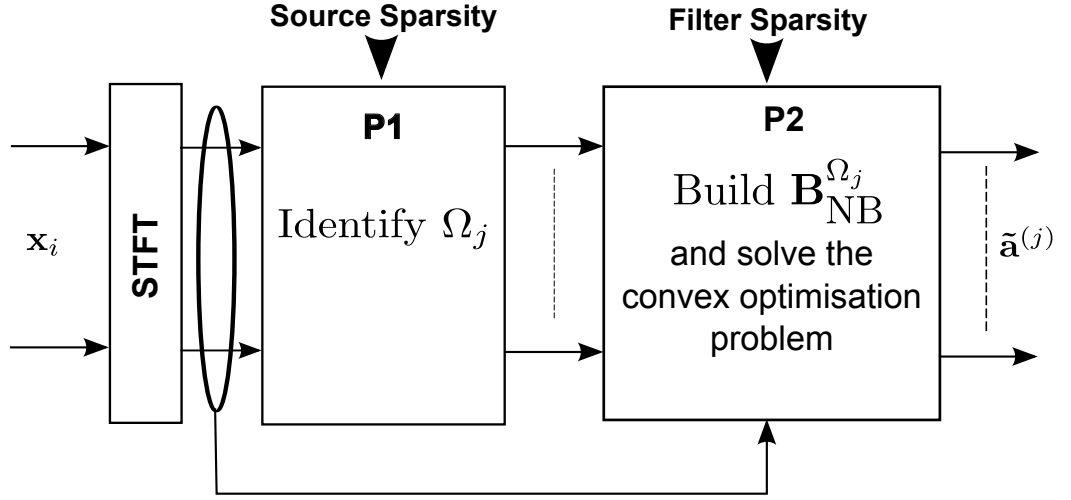


Figure 8.2: Block diagram of the proposed framework.

Solving both **P1** and **P2** simultaneously from the mixtures in a blind setting is a very hard problem. Also, the performance of the filter estimation stage **P2** depends on the performance of the clustering stage **P1**. Hence we shall first experimentally study the performance of the filter estimation step by performing the clustering step using side information about either the sources or filters.

Unlike the single source case where the accuracy of the CR is affected only by the narrowband approximation, in the multiple source case the accuracy of the CR for a given source depends on the interference from the other sources as well. This is because when we assume only one source to be dominant at a given time-frequency point, the other sources are likely to be nearly zero but not exactly zero. Hence, the error parameter ϵ has to be carefully chosen in order to accommodate for both the sources of error: a) due to narrowband approximation and b) due to interference from other sources. Choosing ϵ is a difficult task and in all our experiments it was fixed to a value which gave best results after experiments with different values.

8.3 Multiple filter recovery using oracle clustering

In this section, we shall study the performance of the filter estimation stage in a non-blind setting when the clustering is performed using side information from either the

8.3 Multiple filter recovery using oracle clustering

sources or the filters. For the experiments that we present in this section, we will be using source models Type 2 (sums of sinusoids) and Type 3 (sources with Gabor atoms) described in Sec. 7.2.1.2.

As in the single source setting, we are interested in knowing the range of values of sparsity k for which the filters are recovered with a certain output performance level. The problem complexity is driven by the sparsity level k , the number of sources N and the size of the sets Ω_j for each source j , which in turn depends on the length of the sources.

The performance of the filter estimation stage depends on two factors: a) the accuracy of the CR at the time-frequency points where only one source is dominant and b) the clustering step where the time-frequency points which satisfy the CR for each source are identified and clustered. The accuracy of the CR in turn depends on the time-frequency domain sparsity of the sources and hence it is useful to have a control on the sparsity of the sources. Hence, for the experiments we present in this section we use the sinusoid and Gabor synthetic source models, where the number of active frequencies of the sources can be controlled while generation.

In all the experiments, we fix the length of the filters L and the sources T , and study the recovery performance for different filter sparsity and source sparsity levels.

The clustering step in the experiments is performed using some side information about either the sources or the filters and the quality of clustering step is driven by a threshold parameter. We shall describe the clustering step in detail at appropriate places.

8.3.1 Experiments with sinusoid sources

Data generation

The sparse filters that are needed for this experiment is generated according to the procedure described in Sec. 7.2.1.1 and the sources the source model used is sinusoid, described in Sec. 7.2.1.2. In all the experiments, the length of the filters was set to $L = 256$.

How many number of source frames N_T ?

For a given filter sparsity k , the recovery performance depends on the number of observations available, which is related to the spectral diversity of the sources in each frame and that in turn depends on source sparsity N_F . Hence, for a fixed source sparsity N_F , different values of filter sparsity k demands different number of source frames for a fixed performance guarantee.

For all our experiments, we first determine the number of frames needed for a filter recovery performance guarantee of 20dB in the single source setting with missing frequencies as described in the Sec. 7.2.3. This value depends on k and N_F and hence let us denote this by $\#(k, N_F)$. Fig. 8.3 shows the plot $\#(k, N_F)$ for $N_F = 30$.

For the multiple source setting, for a given k and N_F , we generate

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING




Figure 8.3: Number of source frames required for the output SNR to reach 20dB in the single source setting $N = 1$.

$$N_T(k, N_F) = 2 \cdot (\#(k, 30) \cdot 30) / N_T \quad (8.4)$$

frames per source.

Experimental protocol

For every combination of filter sparsity k and source sparsity N_F , 20 sets of sources and filters are generated according to the protocol described previously. In each trial, the mixtures x_i are obtained and the vectors $\hat{\mathbf{x}}_i$ are formed. For each source j , **P1** and **P2** are solved as described below:

P1: Obtaining Ω_j using side information:

The set Ω_j is constructed using as side information the frequencies $\{f_{jw}^\tau\}_{w=1}^{N_F}$ that are used to generate the source in (7.9). Let $\hat{\mathbf{s}}_j^\tau$ be the time-frequency domain vector of length $F = T + L - 1$ obtained by appropriate zero-padding and transformation of the frame τ of the j^{th} source $s_j^\tau(t)$.

For a fixed threshold ν , we construct the set Ω_j as follows: a time-frequency index (τ, f) belongs to Ω_j iff the magnitude of the coefficient of j^{th} source $\hat{\mathbf{s}}_j^\tau(f)$ dominates

8.3 Multiple filter recovery using oracle clustering

all other sources $\widehat{\mathbf{s}}_{j'}^\tau(f)$, $j' \neq j$ in magnitude at (τ, f) by at least ν dB. That is

$$(\tau, f) \in \Omega_j \iff f \in \{f_{jw}^\tau\}_{w=1}^{N_F} \text{ and } 20 \cdot \log_{10} \left(\frac{|\widehat{\mathbf{s}}_j^\tau(f)|}{|\widehat{\mathbf{s}}_{j'}^\tau(f)|} \right) \geq \nu, \forall j' \neq j. \quad (8.5)$$

P2: Filter estimation by ℓ^1 minimisation:

For each source j , the matrix $\mathbf{B}_{\text{NB}}^{\Omega_j}$ is built using the set Ω_j as a restriction of the matrix \mathbf{B}_{NB} to the rows indexed by the set Ω_j . Then, the resulting convex optimisation problem (7.8) is solved with $\epsilon = 10^{-3}$ to obtain the filter estimates $\widetilde{\mathbf{a}}^{(j)}$.

The performance measure defined in Eq. (7.2.1.3) is used to assess the recovery performance. In the following section, we present the results.

Results

Figures 8.4(a) and 8.4(b) show the average output SNR for various filter sparsity k and source sparsity N_F for 2 and 3 sources respectively, with the parameter $\nu = 10\text{dB}$. In both cases, the anechoic filters ($k = 2$) are recovered with very high SNR, and the output SNR is at least 10dB when $N_F \leq 3$ and $k \leq 10$. For a given sparsity k , the output SNR drops as the number of sinusoids per frames N_F increases. We experimented with higher values of N_F and we found that the performance continues to degrade.

This is possibly because the sources tend to interfere more as N_F increases, thereby violating the CR badly. Indeed, even though we generated sums of sinusoids, their Fourier transform has peaks at the associated frequencies that can have a large main lobe and secondary lobes, leading to interference. This could be compensated by setting a higher threshold ν to obtain the set Ω_j , at the price of a smaller number of “visible” frequencies per time frame, which in turn could be compensated by increasing the number of observed time frames N_T . However, as the sinusoid source model under consideration is extremely simple, we shall not delve further into the experiments for different values of the threshold ν .


8.3.2 Experiments with Gabor sources

In the previous section, the idea of having independent frames with varying frequency content as sources was to simplify the “clustering” step using the source frequencies. In this section, we shall consider another synthetic source model described in Sec. 7.2.1.2. The clustering stage is performed using a heuristic based on the true filters.


Data generation

The sparse filters that are needed for this experiment is generated according to the procedure described in Sec. 7.2.1.1 and the sources the source model used is Gabor, described in Sec. 7.2.1.2. In all the experiments, the length of the filters was set to $L = 256$.

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING



(a) Recovery result for number of sources $N = 2$.



(b) Recovery result for number of sources $N = 3$.

Figure 8.4: Recovery performance of the filter estimation stage based on narrowband CR approach for Type 2 sources when the clusters Ω_j are obtained using side information.

Experimental protocol

For every combination of filter sparsity k and source sparsity N_F , 20 sets of filters and sources are generated according to the protocol described previously. In each trial, the mixtures x_i are obtained and the STFT vectors of the mixtures $\hat{\mathbf{x}}_i(\tau) = [\hat{x}_i(\tau, f)]_f$ are obtained with a window length of F using a Blackman-Harris window [25]. For each source j , **P1** and **P2** are solved as described below:

P1: Obtaining Ω_j using side information:

Unlike the previous experiment, we shall solve the clustering step by using the knowledge about the true filters themselves. For a give source j , the true filter $\mathbf{a}^{(j)}$ indeed satisfies the cross relation at those time-frequency points where only source j is dominant. Hence, we can identify the time-frequency points where source j is active by looking at the points where the CR vanishes. However, if the DFT coefficients of the filters at a frequency f , $a_1^{(j)}$ and $a_2^{(j)}$, are simultaneously close to zero, then the CR

$$\hat{x}_2(\tau, f)\hat{a}_1^{(j)}(f) - \hat{x}_1(\tau, f)\hat{a}_2^{(j)}(f)$$

is satisfied at all the time frames τ for frequency f , irrespective of the magnitude of the sources.

One heuristic to measure the strength of the filters at a given frequency is to look at the term $\hat{x}_2(\tau, f)\hat{a}_1^{(j)}(f) + \hat{x}_1(\tau, f)\hat{a}_2^{(j)}(f)$. We would like to classify a time-frequency point as belonging to source j when the CR is satisfied and the term $\hat{x}_2(\tau, f)\hat{a}_1^{(j)}(f) + \hat{x}_1(\tau, f)\hat{a}_2^{(j)}(f)$ does not vanish.

Denoting

$$\begin{aligned} \text{Num}^{(j)} &= \hat{x}_2(\tau, f)\hat{a}_1^{(j)}(f) + \hat{x}_1(\tau, f)\hat{a}_2^{(j)}(f), \\ \text{Den}^{(j)} &= \hat{x}_2(\tau, f)\hat{a}_1^{(j)}(f) - \hat{x}_1(\tau, f)\hat{a}_2^{(j)}(f), \end{aligned}$$

we consider the following time-frequency points clustering method with side information, given a threshold parameter ν :

$$(\tau, f) \in \Omega_j \iff 20 \cdot \log_{10} \left(\frac{|\text{Num}^{(j)}|}{|\text{Den}^{(j)}|} \right) \geq \nu. \quad (8.6)$$

P2: Filter estimation by ℓ^1 minimisation:

For each source j , the matrix $\mathbf{B}_{\text{NB}}^{\Omega_j}$ is built using the set Ω_j as a restriction of the matrix \mathbf{B}_{NB} to the rows indexed by the set Ω_j . Then, the resulting convex optimisation problem (7.8) is solved with $\epsilon = 6 \times 10^{-4}$ to obtain the filter estimates $\tilde{\mathbf{a}}^{(j)}$.

The performance measure defined in Eq. (7.2.1.3) is used to assess the recovery performance. In the following section, we present the results.

Results

There are two important parameters that drive the performance of the proposed approach: a) The size of the STFT window chosen for the transforming the mixtures into the time-frequency domain and b) the value of the clustering threshold parameter ν . Recovery experiments for three sources case $N = 3$ are performed for input data with various filter and source sparsity levels for various settings of the STFT window sizes and the clustering parameter ν .

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

Effect of STFT window size

Three Gabor sources of length $T = 30000$ samples are generated along with corresponding mixing filters with $\lambda = 1/512$ and $1 \leq N_F \leq 4$. Fig. 8.5 shows the performance of the NB CR based filter recovery approach for different STFT window sizes for four different values of N_F with clustering threshold $\nu = 30$ dB.

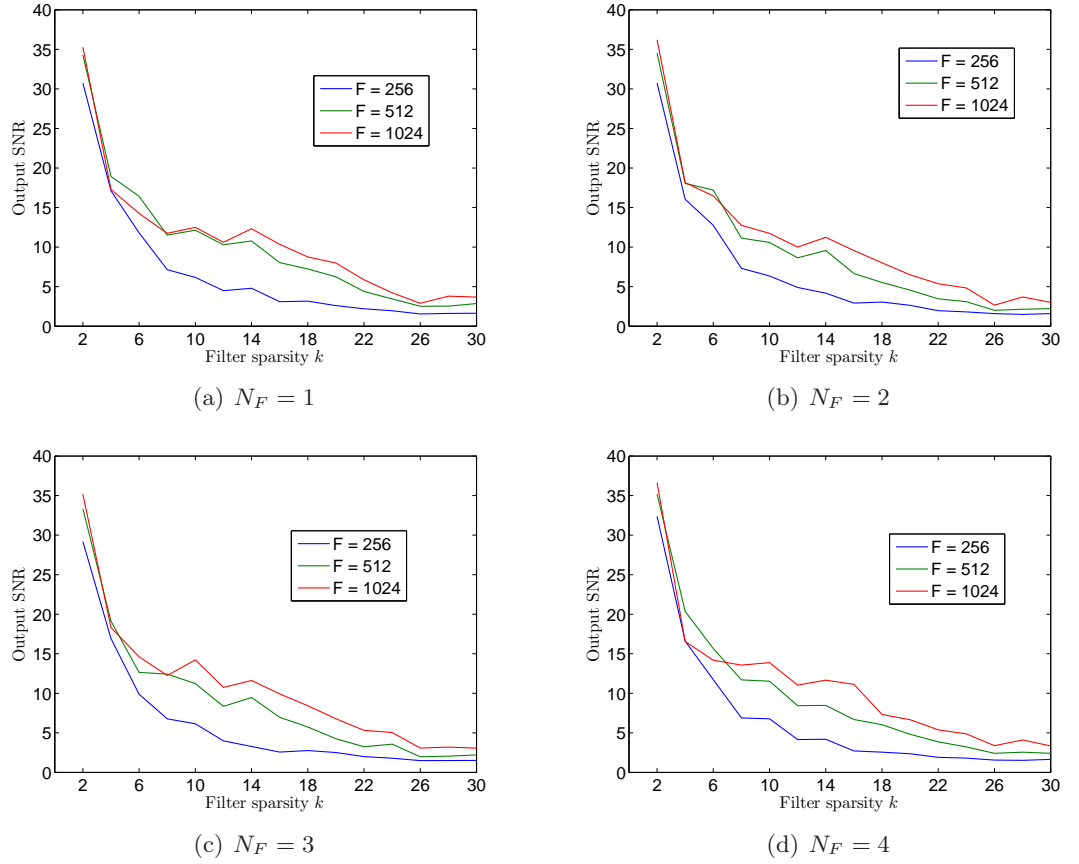


Figure 8.5: Effect of STFT window size on the performance of narrowband CR approach.

In all the four figures we notice that for sparsity $k \leq 6$ the performance is almost the same irrespective of the STFT window size, whereas for $k > 6$, the performance improves when the window size increases. This is due to the reason that the narrowband approximation is more accurate for longer window sizes compared to the filter length and hence the CR is satisfied more accurately.

Also, we can note that the performance is comparable for different source sparsity levels N_F . Even with $N_F = 4$, the sources are sparse enough to have time-frequency points with only one active source.

Effect of clustering threshold ν

Fig. 8.6 shows the recovery performance with three different values of clustering threshold ν for four different values of N_F with $F = 1024$.

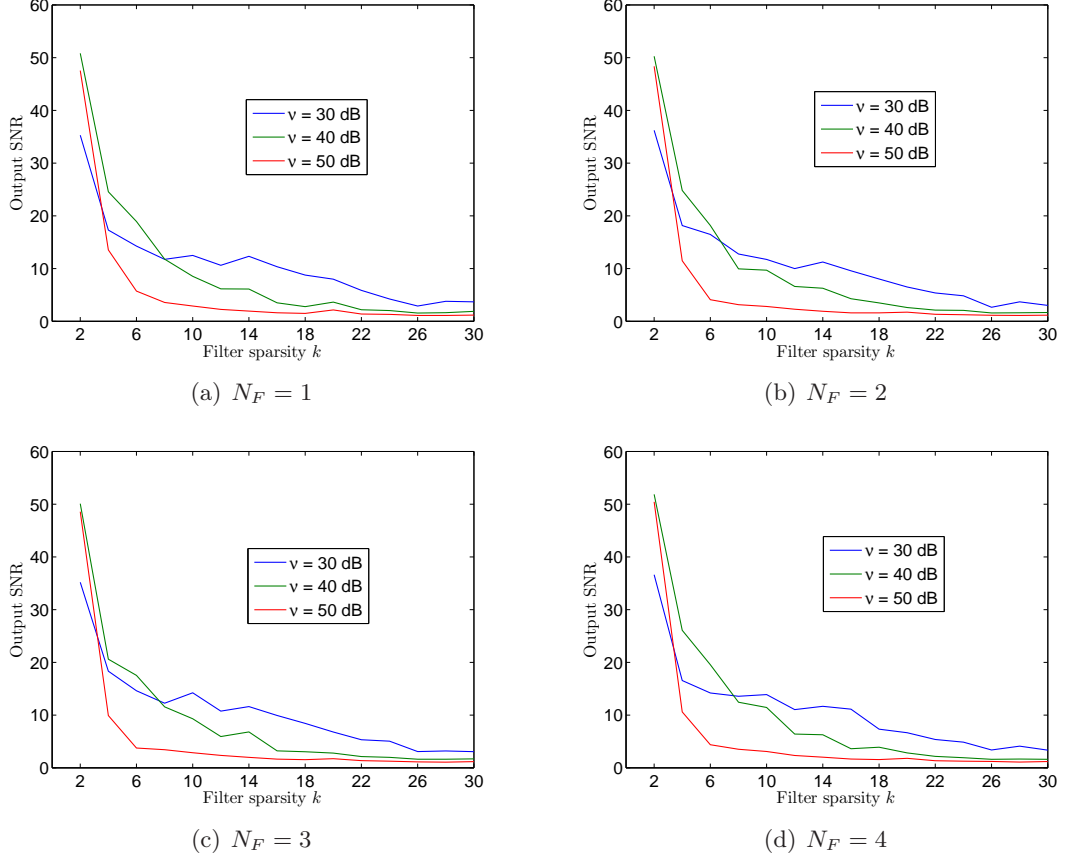


Figure 8.6: Effect of clustering threshold ν on narrowband CR approach.

When the threshold ν is increased, it means that the time-frequency points are selected with a stronger requirement to satisfy the CR, resulting in lesser number of good quality points. Hence we notice in the figures that the performance for $\nu = 40$ dB is better compared to $\nu = 40$ when $k \leq 6$, whereas it falls down rapidly for bigger values of k . For $\nu = 50$ dB, the performance is comparable to $\nu = 40$ dB for the anechoic case ($k = 2$) but it falls off very rapidly for larger k . This is due to the reason that we have fixed the length of the sources and having a higher threshold effectively reduces the number of points available per source for reconstruction.

The narrowband CR based filter recovery method performs well when the size of the STFT analysis window is larger compared to the filter length. The clustering threshold also plays a role in the performance and it has to be chosen carefully depending upon the sparsity of the filters to be recovered and the length of the sources (which decides

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

the number of points available for recovery).

8.4 From narrowband to wideband cross-relation

To build the narrowband CR ([CR-NB](#)) we first obtained the STFT of the mixtures and DFT of the filters and then formulated the CR. The STFT of the mixtures at a given time-frequency point (t, f) is only approximated by the product of the STFT of the source and the DFT of the filter at the same time-frequency point, because the narrowband CR is intrinsically approximate.

As an alternative, we can consider directly the time-domain CR ([CR-TIME](#)), and then examine the STFT of the same. This formulation does not depend upon the narrowband approximation and hence it is called the *wideband* CR.

8.4.1 Time-frequency wideband formulation

Let $x_i(t)$, $i = 1, 2$ be the observations of a source $s(t)$ convolved through the filters $a_i(t)$, $i = 1, 2$ correspondingly in the absence of noise. If the time-domain CR ([CR-TIME](#)) is satisfied across the observations, then we consider the time-frequency domain representation of the sequence $(x_2 \star a_1 - x_1 \star a_2)(t)$. The sequence $(x_2 \star a_1 - x_1 \star a_2)(t) = 0$, $\forall t$ and hence we have

$$\begin{aligned}
 0 &= [x_2 \star a_1 - x_1 \star a_2](\tau, f) \\
 &= \sum_t [x_2 \star a_1 - x_1 \star a_2](t) \cdot w(\tau + t) e^{-2\pi i f t} \\
 &= \sum_{lt} [a_1(l)x_2(t-l) - a_2(l)x_1(t-l)] \cdot w(\tau + t) e^{-2\pi i f t} \\
 &= \sum_l \left(a_1(l) \left[\sum_t x_2(t-l) \cdot w(\tau + t) e^{-2\pi i f t} \right] - a_2(l) \left[\sum_t x_1(t-l) \cdot w(\tau + t) e^{-2\pi i f t} \right] \right).
 \end{aligned}$$

Consider only the term $\sum_l a_1(l) \left[\sum_t x_2(t-l) \cdot w(\tau + t) e^{-2\pi i f t} \right]$ and by making a change of variable $t' = t - l$, we have

$$\sum_l a_1(l) \left[\sum_t x_2(t-l) \cdot w(\tau + t) e^{-2\pi i f t} \right] = \sum_l a_1(l) \sum_{t'} x_2(t') \cdot w(\tau + l + t') e^{-2\pi i f t'} e^{-2\pi i f l}. \tag{8.7}$$

We now notice that the term $\sum_{t'} x_2(t') \cdot w(\tau + l + t') e^{-2\pi i f t'}$ is nothing but the windowed discrete Fourier transform of the mixture x_2 , evaluated at the time-frequency index $(\tau + l, f)$. This is nothing but the projection of the mixture x_2 onto the STFT atom centered at $(\tau + l, f)$.

8.4 From narrowband to wideband cross-relation

Let us denote the STFT atom centered at (τ, f) by $g_{(\tau, f)}$. Then we have

$$\sum_l a_1(l) \sum_{t'} x_2(t') \cdot w(\tau + l + t') e^{-2\pi i f t'} e^{-2\pi i f l} = \sum_l a_1(l) e^{-2\pi i f l} \langle x_2, g_{(\tau+l, f)} \rangle. \quad (8.8)$$

Further, the product $e^{-2\pi i f l} \langle x_2, g_{(\tau+l, f)} \rangle$ is nothing but the windowed DFT of the mixture x_2 shifted by l . Let us denote the DFT of x_2 shifted by l at (τ, f) by $\hat{x}_{2,l}(\tau, f)$. Therefore, by Eqs. (8.7) and (8.8), we have

$$\sum_l a_1(l) \left[\sum_t x_2(t-l) \cdot w(\tau+t) e^{-2\pi i f t} \right] = \sum_l a_1(l) \hat{x}_{2,l}(\tau, f) \quad (8.9)$$

and similarly we have

$$\sum_l a_2(l) \left[\sum_t x_1(t-l) \cdot w(\tau+t) e^{-2\pi i f t} \right] = \sum_l a_2(l) \hat{x}_{1,l}(\tau, f). \quad (8.10)$$

Therefore the CR at a time-frequency point (τ, f) can be written as

$$\sum_l a_1(l) \hat{x}_{2,l}(\tau, f) - \sum_l a_2(l) \hat{x}_{1,l}(\tau, f) = 0. \quad (8.11)$$

Let $\hat{\mathbf{x}}_{j,l}^\tau = \{\hat{x}_{j,l}(\tau, f)\}_f$ be the vector of STFT coefficients of the sequence x_j shifted by l samples at frame index τ . Considering N_T such STFT frames of the mixtures, we can now build the matrix \mathbf{B}_{WB}

$$\mathbf{B}_{\text{WB}} = \begin{bmatrix} \mathbf{B}_{\text{WB},1} \\ \mathbf{B}_{\text{WB},2} \\ \vdots \\ \mathbf{B}_{\text{WB},N_T} \end{bmatrix} \quad (8.12)$$

with

$$\mathbf{B}_{\text{WB},\tau} = [\mathcal{T}[\{\hat{\mathbf{x}}_{2,l}^\tau\}_l] \quad -\mathcal{T}[\{\hat{\mathbf{x}}_{1,l}^\tau\}_l]] \quad (8.13)$$

where

$$\mathcal{T}[\{\hat{\mathbf{x}}_{j,l}^\tau\}_l] = \begin{bmatrix} \hat{x}_{j,0}(\tau, 0) & \hat{x}_{j,1}(\tau, 0) & \dots & \hat{x}_{j,L-1}(\tau, 0) \\ \hat{x}_{j,0}(\tau, 1) & \hat{x}_{j,1}(\tau, 1) & \dots & \hat{x}_{j,L-1}(\tau, 1) \\ \vdots & \vdots & \dots & \vdots \\ \hat{x}_{j,0}(\tau, F-1) & \hat{x}_{j,1}(\tau, F-1) & \dots & \hat{x}_{j,L-1}(\tau, F-1) \end{bmatrix}. \quad (8.14)$$

Let us recall that $\hat{x}_{j,l}(\tau, f)$ is the STFT coefficient at (τ, f) of x_j shifted by l samples.

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

In the single source setting, the matrix \mathbf{B}_{WB} is of size $(F \cdot N_T) \times (2L)$ and it satisfies the CR $\mathbf{B}_{\text{WB}} \cdot \mathbf{a} = \mathbf{0}$, with

$$\mathbf{a} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \end{bmatrix}. \quad (8.15)$$

Therefore, we can estimate the filters by solving the ℓ^1 minimisation problem (7.8) with \mathbf{B}_{WB} .

Modification of the two-stage framework for wideband CR approach

In the case of multiple source setting, we again rely on the time-frequency domain sparsity of the sources to formulate an ℓ^1 minimisation problem to estimate each pair of filters. In case of the narrowband CR based approach for multiple filter estimation, to assert a time-frequency point of the mixture at (τ, f) to belong to a particular source, it is sufficient that all the other sources at (τ, f) has negligible energy so that the narrowband CR is satisfied.

In the case of wideband CR, the relation in Eq. (8.11) depends on L STFT points of the mixtures at (τ, f) obtained using L one-sample shifts of the mixtures. This means that if the CR has to be satisfied for a particular source j at a time-frequency point (τ, f) , then a sufficient condition is that all the other sources ($j' \neq j$) should have their STFT to be equal zero at (τ, f) with $2L$ one-sample shifts.

That is, the CR at (τ, f) is satisfied for source j if

$$\hat{s}_{j',l}(\tau, f) = 0, 0 \leq l \leq 2L - 1, j' \neq j. \quad (8.16)$$

Therefore, the wideband CR is more demanding than the narrowband CR.

For each source j , if we can identify sets of time-frequency points Ω_j where the other sources satisfy Eq. (8.16), then we can obtain the matrices $\mathbf{B}_{\text{WB}}^{\Omega_j}$ as a restriction of the matrix \mathbf{B}_{WB} to the rows indexed by Ω_j . We then can solve the ℓ^1 minimisation problem (7.8) with the matrices $\mathbf{B}_{\text{WB}}^{\Omega_j}$ to estimate the filters.

The two-stage framework based on the wideband CR can now be written as:

1. Compute the STFT of the mixtures $x_j, j = 1, 2$ with shifts $0 \leq l \leq L - 1$: $\hat{x}_{j,l}(\tau, f), 1 \leq \tau \leq N_T, 0 \leq f < F$.
2. For each source j ,

P1 { (a) Identify the set Ω_j .

P2 { (b) Build the matrix $\mathbf{B} := \mathbf{B}_{\text{WB}}^{\Omega_j}$.

(c) Solve (7.8) with ϵ to obtain the estimated filter $\tilde{\mathbf{a}}^{(j)}$.

8.4.2 Experiments with Gabor sources

In this section, we shall present the experimental results for filter estimation using the wideband CR based approach. We shall only consider Gabor sources in this section. The input data that was used for the experiments reported in Sec. 8.3.2 is used for the current set of experiments as well.

For every combination of filter sparsity k and source sparsity N_F , 20 sets of filters and sources are generated. In each trial, the mixtures x_i are obtained and the STFT vectors of the mixtures $\hat{\mathbf{x}}_i(\tau) = [\hat{x}_i(\tau, f)]_f$ are obtained with a window length of F using a Blackman-Harris window. For each source j , **P1** and **P2** are solved as described below:

P1: Obtaining Ω_j using side information:

The clustering of the time-frequency points is done by using the true filters as the side information. The strategy is similar to the one reported in Sec. 8.3.2 for the narrowband CR method: to make use of the CR with the true filters for identifying the time-frequency points where only one source is dominant. For a given source j , the idea is to locate the points where the CR is satisfied. We define the following two terms

$$\begin{aligned} \text{Num}^{(j)} &= \sum_l \hat{x}_{2,l}(\tau, f) a_1^{(j)}(l) + \sum_l \hat{x}_{1,l}(\tau, f) a_2^{(j)}(l), \\ \text{Den}^{(j)} &= \sum_l \hat{x}_{2,l}(\tau, f) a_1^{(j)}(l) - \sum_l \hat{x}_{1,l}(\tau, f) a_2^{(j)}(l), \end{aligned}$$

and the time-frequency points clusters, given a threshold parameter ν , are obtained by evaluating:

$$(\tau, f) \in \Omega_j \iff 20 \cdot \log_{10} \left(\frac{|\text{Num}^{(j)}|}{|\text{Den}^{(j)}|} \right) \geq \nu. \quad (8.17)$$

P2: Filter estimation by ℓ^1 minimisation:

For each source j , the matrix $\mathbf{B}_{\text{WB}}^{\Omega_j}$ is built using the set Ω_j as a restriction of the matrix \mathbf{B}_{WB} to the rows indexed by the set Ω_j . Then, the resulting convex optimisation problem (7.8) is solved with $\epsilon = 6 \times 10^{-4}$ to obtain the filter estimates $\tilde{\mathbf{a}}^{(j)}$.

The performance measure defined in Eq. (7.2.1.3) is used to assess the recovery performance. In the following section, we present the results.

Results

As in the case of the narrowband CR method, there are two important parameters that drive the performance of the proposed approach: a) The size of the STFT window chosen for the transforming the mixtures into the time-frequency domain and b) the value of the clustering parameter ν . Recovery experiments are performed for input data with various filter and source sparsity levels for various settings of the STFT window sizes and the clustering parameter ν .

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

Effect of STFT window size

Three Gabor sources of length $T = 30000$ samples are generated along with corresponding mixing filters with $\lambda = 1/512$ and $1 \leq N_F \leq 4$. Fig. 8.7 shows the performance of the WB CR based filter recovery approach for different STFT window sizes for four different values of N_F with clustering threshold $\nu = 30$ dB.

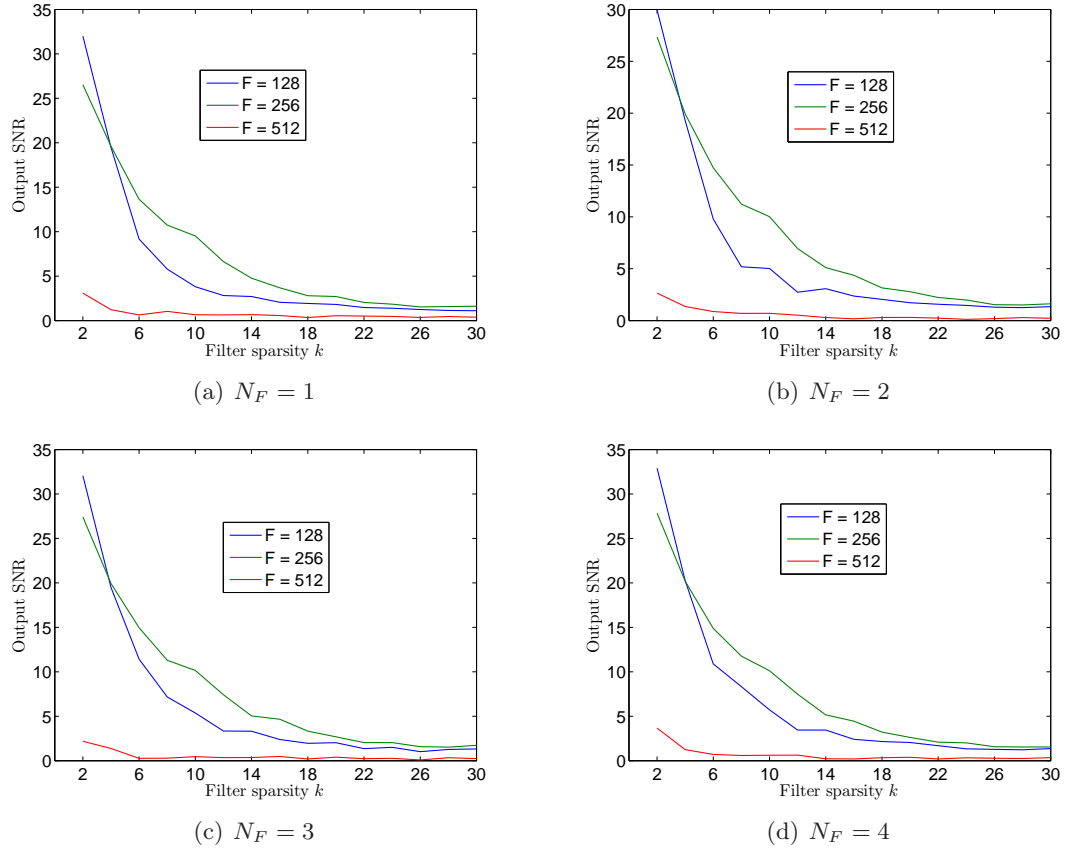


Figure 8.7: Effect of STFT window size on the performance of wideband CR approach.

As opposed to the narrowband CR approach, the performance of the wideband CR approach deteriorates when the STFT window size increases relative to the filter size. As described in Sec. 8.4.1, the wideband CR is satisfied for a source at a time-frequency point (τ, f) when there is only one active source in the neighbourhood of length L samples. When the STFT analysis window is longer than the filter length, then the STFT coefficient at (τ, f) will be computed considering a neighbourhood longer than L samples which results in the interference from outside the neighbourhood. This will affect the wideband CR and hence the performance goes down. Alternatively, when the STFT window size is smaller compared to the filter length, smaller neighbourhoods are considered for analysis and hence there is less chance of interference, which leads to a

good satisfaction of the wideband CR.

Effect of clustering threshold ν

Fig. 8.8 shows the recovery performance with four different values of clustering threshold ν for four different values of N_F with $F = 128$.

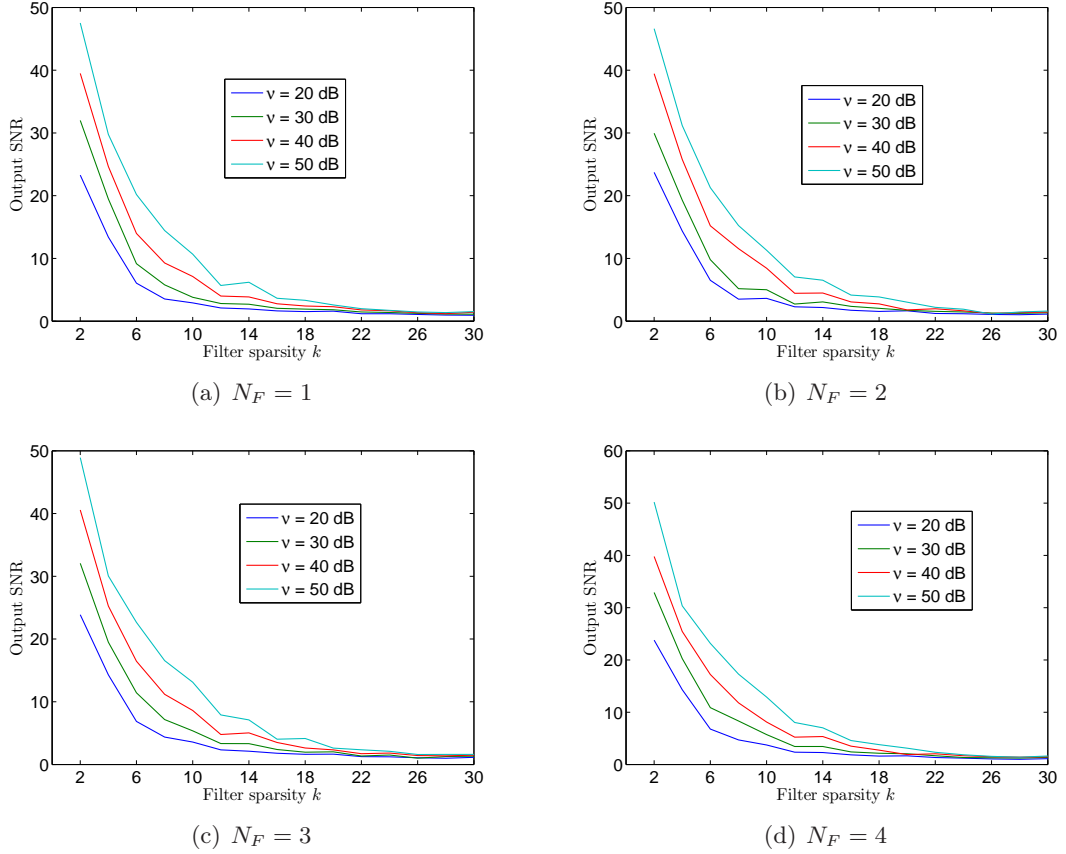


Figure 8.8: Effect of clustering threshold ν on wideband CR approach.

The performance of the wideband CR approach significantly improves when the clustering threshold ν is increased. In the anechoic ($k = 2$) case, the improvement is over 15 dB when the threshold is increased from 30 dB to 50 dB. The performance with $\nu = 50$ is better than that of $\nu = 30$ till $k = 18$ after which they are similar and poor.

Even though increasing the threshold decreases the number of time-frequency points available for filter recovery, the wideband CR recovers the filters with better output SNR compared to smaller thresholds.

Comparing Fig. 8.7 with Fig. 8.5 we can say that the wideband CR approach performs better than the narrowband CR approach when the STFT window size is smaller than the filter length, and viceversa. Similarly by comparing Figs. 8.8 and 8.6 we can

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

say that the wideband CR approach with a larger threshold and smaller STFT size performs better than the narrowband CR approach with a larger threshold and longer STFT size.

8.5 Blind clustering and filter estimation in a special convolutive setting

In the previous sections, the focus was on the experimental evaluation of the filter estimation stage, when the clustering is performed using side information, both in the narrowband and wideband CR approach. In this section, we will consider a simple setting of the sources that enables blind clustering of the time-frequency points and subsequent estimation of filters.

8.5.1 Setting and approach

We consider a scenario in which all the sources but one are mixed instantaneously. Let us denote the length of the associated filters by $L_j = 1$, and each such source is associated to a *Intensity Parameter* (IP) $\theta_j = \tan^{-1}(a_{2j}/a_{1j})$, introduced in Sec. 2.3.3.

As we discussed in Sec. 2.3.1.3, intensity parameter estimation by clustering the time-frequency points in linear instantaneous and anechoic settings has been studied extensively and many methods like DEMIX [7] can be used to do the same. Hence, if the time-frequency points associated to the “instantaneous” sources (sources that are instantaneously mixed) can be detected and removed from the time-frequency plane, then the remaining points can be considered as belonging to the “convoluted” source. With these points, we can form a CR and perform filter estimation.

For simplicity let us consider one “instantaneous” source s_1 and one “convoluted” source s_2 . Existing algorithms such as DEMIX [7] can both estimate the IP θ_1 of the instantaneous source s_1 and the TF region Ω_1 where it is prominently active. This is done simply by looking at the deviation between the estimated angle θ_1 of the source to the inverse tangent of the ratio of the mixtures at each time-frequency point. That is, given a threshold η

$$(\tau, f) \in \Omega_1 \iff \left| \tan^{-1} (|\hat{\mathbf{x}}_2(\tau, f)/\hat{\mathbf{x}}_1(\tau, f)|) - \theta_1 \right| < \eta. \quad (8.18)$$

We can then build a set $\overline{\Omega}_1$ containing all TF points closed in time or in frequency to the points in Ω_1 , and obtain the set Ω_2 as the complement of $\overline{\Omega}_1$.

Fig. 8.9 shows an example of STFT of two sources, and Fig. 8.10(a) displays the STFT of one of the mixtures x_1 (black corresponds to high energy, white to low energy). Figure 8.10(b) illustrates the set Ω_2 obtained with the described approach (white indicates points in Ω_2), and we can see that as expected Ω_2 only contains TF points where source s_1 is not very active.

8.5 Blind clustering and filter estimation in a special convolutive setting




Figure 8.9: Spectrograms of the two sources: (a) source s_1 is a flute sound and (b) source s_2 is a guitar sound.




Figure 8.10: (a) Spectrogram of mixture x_1 . (b) Time-frequency mask. The white pixels correspond to the points in the set Ω_2

8.5.2 Experiments with audio sources

Experiments are conducted to assess the recovery performance of the wideband and narrowband CR based methods in a special setting with two real audio sources: s_1 a flute sound and s_2 a guitar sound. The flute sound is mixed using a pair of anechoic filters with a known intensity parameter θ_1 , and the guitar sound is mixed with filters of sparsity k and length $L = 256$ generated according to the procedure described in Sec. 7.2.1.1.

Experimental protocol

For each sparsity level k , 20 sets of filters are generated and the mixtures are obtained. The STFT of the mixtures are then computed using a Blackman-Harris window of size $F = 512$. We then blindly build the sets Ω_1 and Ω_2 with $\eta = 0.1$ using Eq. (8.18). The

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

set Ω_2 is used to then build the matrices $\mathbf{B}_{\text{NB}}^{\Omega_2}$ and $\mathbf{B}_{\text{WB}}^{\Omega_2}$ as described in Secs. 8.1 and 8.4 respectively.

To reduce the computations, the sizes of the matrices $\mathbf{B}_{\text{NB}}^{\Omega_2}$ and $\mathbf{B}_{\text{WB}}^{\Omega_2}$ are reduced by merging the rows corresponding to the identical frequency bins f . This merging is done by averaging for each frequency bin f the normalized rows of the matrices corresponding to f . Also, the sources are simultaneously active in the time segment between $t = 4000$ and $t = 5000$, and hence we keep only the points in this segment.

The filters are then estimated by solving the ℓ^1 minimisation problem (7.8) with $\mathbf{B} := \mathbf{B}_{\text{NB}}^{\Omega_2}$ and $\mathbf{B} := \mathbf{B}_{\text{WB}}^{\Omega_2}$ for the narrowband and wideband approach respectively with $\epsilon = 6 \times 10^{-4}$.

Debiasing step

Suppose the sparsity k of the filters is available to us as side-information, then a least squares problem can be solved by selecting the support of the k largest coefficients (in magnitude) of the estimated filters and the corresponding columns of the matrix \mathbf{B} , in both narrowband and wideband cases. We shall refer to this step as the *debiasing* (DB) step.

Let Γ_j be the set of indices of k largest coefficients of the estimated vector $\tilde{\mathbf{a}}^{(j)}$. To estimate the vector $\tilde{\mathbf{a}}^{(j)}$, the optimisation problem in Eq. 7.8 is solved with a constraint $\mathbf{a}_1(L/2) = 1$ and hence $L/2 \in \Gamma_j$. Let $\mathbf{b}_{L/2}$ be the $(L/2)^{\text{th}}$ column vector of the matrix \mathbf{B} and let $\mathbf{B}_{\Gamma'}$ be a matrix built using the columns of \mathbf{B} whose indices are in the set Γ' , where $\Gamma' = \Gamma - \{L/2\}$. The debiasing step computes the least squares solution $\tilde{\mathbf{a}}^{(j)}$ using

$$\tilde{\mathbf{a}}^{(j)} = \mathbf{B}_{\Gamma'}^\dagger \cdot -\mathbf{b}_{L/2}. \quad (8.19)$$

where $\mathbf{B}_{\Gamma'}^\dagger$ is the Moore-Penrose pseudoinverse of the matrix $\mathbf{B}_{\Gamma'}$.

Results

Fig. 8.11 shows the performance curves for the narrowband and wideband CR based approach with and without the debiasing step. We notice that the wideband approach outperforms the narrowband approach by between 5 and 20 dB, both with and without debiasing.

A comparison with the state of the art method for anechoic filter estimation using GCC-PHAT [106], is also done. The delays associated with the anechoic filters are estimated using GCC-PHAT and the magnitudes of the peaks are estimated by averaging the intensity parameter of all the time-frequency points in Ω_1 .

GCC -PHAT works better than the wideband and narrowband CR based approach but we notice that the wideband CR based approach with debiasing has around 10 dB improved performance compared to GCC-PHAT for the anechoic $k = 2$ case. The comparison of the performance of GCC-PHAT for $k > 2$ is not fair because GCC-PHAT can be used only for estimating the delays associated with anechoic filters.




Figure 8.11: Performance of filter recovery using narrowband and wideband CR approach for $\theta_1 = 0.2$ radian.

8.6 Summary

In this chapter, we developed a framework for estimating multiple sparse mixing filters from convolutive mixtures based on the time-frequency domain cross-relation (CR). We first developed the framework based on the narrowband approximation of the time-domain cross-relation. The framework consists of two steps: 1) a clustering step to identify time-frequency points of the mixtures where only one source is active and b) a filter estimation step using the clusters obtained in step 1.

As a consequence of the time-frequency domain sparsity of the sources, the mixtures will contain several time-frequency points where only one source contributes to the mixture and the objective of the first stage of the framework is to identify these sets of time-frequency points. These sets of time-frequency points are used to obtain the CR for each source and the second step uses these cross-relations to formulate and solve an ℓ^1 minimisation problem for each of the sources to estimate the filters.

Experiments are conducted to evaluate the performance of the filter estimation stage by assuming that the clustering step is performed using some side information. A heuristic approach is developed using information about the sources to perform the clustering of time-frequency points. The filter estimation stage is experimentally evaluated for different factors that drive the complexity of the problem and different parameters that influence the performance of the filter estimation.

Subsequently, a wideband version of the time-frequency domain cross-relation is developed which retains the exactness of the time-domain cross-relation. The application of wideband CR for filter estimation in the multiple source setting is presented and the two-stage framework is appropriately modified to suit the wideband CR. The wideband CR requires a neighbourhood of points in the mixtures where only one source is active,

8. A CONVEX OPTIMISATION FRAMEWORK FOR MULTIPLE SOURCE LOCALISATION IN SPARSE CONVOLUTIVE SETTING

as opposed to a single point in the narrowband case. Hence, the wideband CR is more stringent than the narrowband CR.

Experimental evaluation of the filter estimation stage using wideband CR is done by performing the clustering using a heuristic which requires the knowledge of the true filters as side-information. The wideband CR approach is shown to be better than the narrowband CR approach under certain settings.

In the last part of the chapter, a special setting where only one source is convolutively mixed is considered and a simple approach to blindly cluster the time-frequency points is presented. Filter recovery experiments with two real audio sources are presented and it is shown that the wideband CR based approach combined with a post-processing step performs very well compared to the narrowband approach.

In the following chapter, we conclude the thesis and discuss certain possibilities for the extensions and generalisations of the current work.

Part III

Conclusion and perspectives

Chapter 9

Conclusions and perspectives

The work presented in this thesis is based centrally on the *sparsity* of the objects involved: filters and sources. The following section summarises and concludes the thesis.

Conclusions

Chapter 5 shows that filter sparsity can be effectively used as a consistency measure to correct the filter permutation problem in the absence of scaling. The sparsity of the filters and the number of permuted sub-bands have a nice interplay which is demonstrated in Fig. 5.1. A simple combinatorial ℓ^1 minimisation algorithm such as Algorithm 1 can effectively recover correct permutations. In noiseless conditions, the experiments show that as long as the sum of filter sparsity and number of permutations does not cross half the filter length, the permutations can be successfully recovered. Even under noisy conditions, an ℓ^1 minimisation based approach for correcting permutations can recover filters upto the input SNR level. Hence, we can conclude that as long as the filters are sparse and the number of permutations is not too large, minimising the ℓ^1 norm of the filters is indeed a good way to correct the permutations in the absence of scaling, even under noisy conditions.

Chapter 6 attempts to theoretically establish the connections between filter sparsity and permutations and explain the experimental observations in chapter 5. In the case where the filters have mutually disjoint supports theorem 6.1 showed that permutations can only increase the ℓ^1 norm of the filters and hence ℓ^1 minimisation recovers the correct permutations. In case of two filters with arbitrary supports, theorem 6.2 provides a condition on the recoverability of the filters in terms of the product of filter sparsity and number of permutations. However, this result is pessimistic compared to the experimental observations.

Chapter 7 presents an approach to estimate sparse filters from stereo convolutive observations in the single source setting. The idea is to exploit the time-frequency domain cross-relation between the channels and formulate an ℓ^1 minimisation problem

9. CONCLUSIONS AND PERSPECTIVES

for filter estimation. The time-frequency domain CR is developed by the narrowband approximation of the time-domain CR.

It is shown experimentally with synthetically generated Gaussian sources that the time-frequency domain narrowband CR method can recover filters having upto 15 spikes with more than 20 dB output SNR (as defined in (7.12)). The single Gaussian source case is the ideal problem setting where the time-frequency CR is valid over the entire time-frequency plane, and the corresponding recovery performance is the best we can have with the narrowband CR approach.

However, the advantage of using time-frequency domain CR in place of the time-domain CR is that the time-frequency domain CR method is capable of recovering filters even with missing frequency information about the filters. This has been demonstrated experimentally and filters having upto 15 spikes can be recovered with an output SNR of at least 20 dB even when only 50% of the frequency information, relative to the size of the filters, is available. It is due to the sparsity of filters that it is possible to estimate them even with missing information, and this fact plays an important role in the task of estimating filters in the presence of multiple sources.

Chapter 8 builds upon the idea of time-frequency cross-relation presented in chapter 7 to propose a framework for estimating multiple sparse filters from stereo convolutive mixtures of several sources. In the presence of multiple sources which are disjoint in the time-frequency domain, a clustering step identifies the time-frequency points in the mixtures where only one particular source contributes and these points are used to form the cross-relation. Using the CR, the filters are estimated as it is done in the single source case.

Firstly, the sparsity of sources is helping us to form the CR for each source and secondly the sparsity of filters is helping in estimating them even with missing information. The clustering step is crucial because it is this step which is inherently solving the permutation ambiguity. The way in which the filter estimation problem is formulated takes care of the scaling ambiguity. Therefore, two important problems associated with convolutive source localisation are completely taken care of by the proposed framework. However, it should be mentioned that clustering time-frequency points is a difficult task in itself.

The experimental evaluation primarily focusses on the filter estimation stage and the clustering is performed heuristically using side information about the sources themselves. The STFT window size influences the performance of the algorithm: larger the window size relative to the filters, better the narrowband approximation and this results in improved performance. Also, with the presented heuristic clustering scheme, the clustering threshold parameter plays an important role: for a fixed source size, a large threshold selects few good quality points and hence one can recover only highly sparse filters with very good quality, whereas a smaller threshold selects many points, possibly not all of them satisfying the CR strictly which results in reduced performance for highly sparse filters, but an improvement for filters with relatively larger number of spikes.

To avoid the problem of approximation, a wideband version of the time-frequency

domain CR is proposed. This retains the exact nature of the time-domain CR even in the time-frequency domain, but places additional restrictions on the time-frequency points in order to satisfy the CR. The framework for multiple filter estimation tailored for wideband CR also contains a clustering step and a filter estimation step.

Experimental evaluation of the wideband CR based filter estimation approach shows that the wideband CR has good performance when the STFT window sizes are smaller compared to the filter length. More importantly, even though larger values of the clustering threshold selects fewer points, the wideband CR method performs better compared to the smaller thresholds. If the STFT window size and the threshold are chosen carefully, the wideband CR based method outperforms the narrowband CR based method for same sparsity levels of the filters.

In a special setting with two real audio sources, with only one of them being convolutively mixed, we presented a method to blindly cluster the time-frequency points of the mixtures. With these clusters, the wideband CR based method is shown to be superior compared to the narrowband CR method, and also a post-processing step can boost the results significantly.

For highly sparse filters ($k \leq 6$ when filter length $L = 256$), the performance of both the narrowband and wideband approaches are comparable when the STFT window sizes are appropriately chosen (longer for narrowband and shorter for wideband). When the number of spikes increases, then the wideband CR method with shorter STFT window size and larger clustering threshold gives better performance than the narrowband one.

In the following section we shall present some of the possible extensions and generalisations of the work presented in this thesis.

Perspectives

Sparsity based method to solve the scaling ambiguity

In chapter 5, we have presented an ℓ^1 norm minimisation based approach for solving the permutation problem for source localisation in convolutive setting. Although we considered solving the permutation problem in isolation, it is often accompanied by the scaling problem as well. Hence, we can explore the possibility of using sparsity as a tool to solve the scaling problem too.

If sparsity based methods proves to be successful for solving the scaling problem, then we can explore the possibility of developing an ℓ^1 minimisation approach for solving the permutation and scaling problems simultaneously.

Efficient ℓ^1 minimisation algorithm

The ℓ^1 minimisation algorithm to solve the permutation problem that is presented in chapter 5 is combinatorial in nature. The algorithm is not scalable when the number

9. CONCLUSIONS AND PERSPECTIVES

of sources is large. Hence, an efficient algorithm to minimise the ℓ^1 norm of the filters is essential.

As a simple improvement to the proposed algorithm, one can replace the combinatorial brute force search method for finding a permutation in each sub-band f to minimise the ℓ^1 norm by an optimisation step which searches for a matrix $\mathbf{P}(f)$ that minimises the ℓ^1 norm, under the constraint that $\mathbf{P}(f)$ is a permutation matrix.

Permutations, scaling and sparsity: Theoretical connections

Theoretical connections between filter permutations/scaling and sparsity will help us to understand the regime under which the ℓ^1 minimisation based approach will succeed in recovering the permutations and scaling. It will be interesting to generalise the work that we have presented in chapter 5. Understanding the connections between these ambiguities and sparsity could also enable us to come up with efficient algorithms for solving the ℓ^1 minimisation problem.

Improved methods for blind clustering of time-frequency points

In chapter 7, we proposed a framework for multiple sparse filter estimation from stereo mixtures, which consists of a time-frequency point clustering stage and a filter recovery stage. A fully blind way of clustering the time-frequency points was proposed only for a special setting where only one source is convolutively mixed and another source is mixed instantaneously. The clustering stage has to be improved and generalised to handle multiple sources that are convolutively mixed.

A fully blind system to cluster the time-frequency points involves two steps: 1) identification of points where only one source is active and 2) association of those points to the sources. We can either extend the existing time-frequency points identification methods proposed for instantaneous and anechoic mixtures such as [7], or use the methods proposed for convolutive mixtures such as [49, 48]

Once the points where a single source is active are identified, then subspace based method such as the Generalized Principal Component Analysis (GPCA) [107, 108] can be used to cluster the time-frequency points efficiently. GPCA is an algebraic method which involves solving a set of linear equations for clustering the points.

Another interesting avenue to explore is clustering of time-frequency points via dictionary learning. The idea is to learn dictionaries from the mixtures which allow sparse representations of sources and at the same time provide discrimination between the sources.

Alternation of clustering and filter estimation stages

At the end of chapter 7, we mentioned about the possibility of using an alternating scheme to simultaneously estimate the filters and clusters. We can assume that anechoic

approximations of the mixing filters are given to us and starting with these approximate filters, we can alternatively estimate the clusters of time-frequency points and the actual sparse mixing filters. It will be interesting to study whether such an alternating scheme converges to true clusters and true filters.

Efficient optimisation algorithm for filter estimation

In all our experiments, the convex program in the filter estimation step was solved using a readymade convex optimisation library **CVX**. While **CVX** is an easy to use package for modelling and solving small scale convex problems, it is unable to efficiently handle large scale problems. Hence, an efficient optimisation algorithm has to be designed to solve the filter estimation problem, with which we can study the performance of the filter estimation stage for large problem sizes.

Iterated threshold and shrinkage algorithms [70] have been gaining popularity in the last few years for solving ℓ^1 minimisation algorithms, for its simple structure and efficiency. It would be an interesting line of work to study the use of one such algorithm for filter estimation purpose.

From source localisation to source separation

As mentioned in the beginning of the thesis, often the goal of mixing filter estimation is to recover the sources. So, this will be a natural extension of our work.

Estimation of sources by exploiting their sparse representations and by using the explicit knowledge of mixing filters has been studied by Kowalski et al. [2]. If we have good estimates of the filters using our framework, then it will be interesting to use these filters for source estimation.

Extension to multiple channel setting

The work we presented in this thesis deals with the case of stereo channels (Number of channels $M = 2$). We can generalise the framework to multiple channels setting by extending the multiple channel cross-relation (CR) based approach reported in [89].

9. CONCLUSIONS AND PERSPECTIVES

Bibliography

- [1] P. Comon and C. Jutten, eds., *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press, 2010. [2](#), [13](#), [29](#), [35](#)
- [2] M. Kowalski, E. Vincent, and R. Gribonval, “Beyond the narrowband approximation: wideband convex methods for under-determined reverberant audio source separation,” *Trans. Audio, Speech and Lang. Proc.*, vol. 18, pp. 1818–1829, September 2010. [8](#), [20](#), [47](#), [131](#)
- [3] M. Plumbley, T. Blumensath, L. Daudet, R. Gribonval, and M. Davies, “Sparse representations in audio and music: From coding to source separation,” *Proceedings of the IEEE*, vol. 98, pp. 995–1005, jun. 2010. [8](#), [20](#), [38](#)
- [4] A. Jourjine, S. Rickard, and O. Yilmaz, “Blind separation of disjoint orthogonal signals: demixing n sources from 2 mixtures,” in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 2985–2988, 2000. [8](#), [21](#), [23](#), [35](#), [39](#), [45](#), [104](#)
- [5] P. Bofill and M. Zibulevsky, “Underdetermined blind source separation using sparse representations,” *Signal Processing*, vol. 81, no. 11, pp. 2353–2362, 2001. [8](#), [21](#), [35](#), [43](#)
- [6] M. Zibulevsky and B. A. Pearlmutter, “Blind source separation by sparse decomposition in a signal dictionary,” *Neural Computation*, vol. 13, no. 4, pp. 863–882, 2001. [8](#), [21](#), [35](#)
- [7] S. Arberet, R. Gribonval, and F. Bimbot, “A robust method to count and locate audio sources in a multichannel underdetermined mixture,” *Signal Processing, IEEE Transactions on*, vol. 58, pp. 121–133, jan. 2010. [8](#), [21](#), [120](#), [130](#)
- [8] H. Liu, G. Xu, and L. Tong, “A deterministic approach to blind identification of multi-channel FIR systems,” in *Proc. of ICASSP*, (Washington, DC, USA), pp. 581–584, 1994. [9](#), [57](#), [59](#), [65](#)
- [9] M. Kocic, D. Brady, and M. Stojanovic, “Sparse equalization for real-time digital underwater acoustic communications,” in *OCEANS ’95. MTS/IEEE. Challenges*

BIBLIOGRAPHY

- of Our Changing Global Environment. Conference Proceedings.*, vol. 3, pp. 1417–1422 vol.3, Oct. 1995. [20](#)
- [10] N. R. Chapman and I. Barrodale, “Deconvolution of marine seismic data using the l1 norm,” *Geophysical Journal International*, vol. 72, pp. 93–100, Jan. 1983. [20](#)
- [11] S. Ariyavisitakul, N. Sollenberger, and L. Greenstein, “Tap-selectable decision-feedback equalization,” *Communications, IEEE Transactions on*, vol. 45, pp. 1497–1500, Dec. 1997. [20](#)
- [12] J.-F. Cardoso, “Blind signal separation: statistical principles,” *Proceedings of the IEEE*, vol. 90, no. 8, pp. 2009–2026, 1998. [31](#)
- [13] T. M. Cover and J. A. Thomas, *Elements of information theory*. New York, NY, USA: Wiley-Interscience, 1991. [31](#)
- [14] P. Comon, “Independent component analysis, a new concept?,” *Signal Processing*, vol. 36, no. 3, pp. 287–314, 1994. [32](#)
- [15] C. Jutten and J. Héroult, “Blind separation of sources, part i: An adaptive algorithm based on neuromimetic architecture,” *Signal Processing*, vol. 24, no. 1, pp. 1 – 10, 1991. [32](#)
- [16] A. J. Bell and T. J. Sejnowski, “An information-maximization approach to blind separation and blind deconvolution,” *Neural Computation*, vol. 7, no. 6, pp. 1129–1159, 1995. [32](#), [33](#)
- [17] A. Hyvärinen, “Fast and robust fixed-point algorithms for independent component analysis,” *IEEE Transactions on Neural Networks*, vol. 10, no. 3, pp. 626 – 634, 1999. [32](#)
- [18] J.-F. Cardoso, “Blind identification of independent components with higher-order statistics,” in *Proc. of Workshop on Higher-Order Spectral Analysis*, pp. 157–160, 1989. [32](#)
- [19] J.-F. Cardoso, “High-order contrasts for independent component analysis,” *Neural Computation*, vol. 11, no. 1, pp. 157–192, 1999. [32](#)
- [20] J.-F. Cardoso, “Blind separation of instantaneous mixture of sources via the gaussian mutual information criterion,” *Signal Processing*, vol. 81, no. 4, pp. 855 – 870, 2001. [32](#)
- [21] J. Cardoso and A. Souloumiac, “Jacobi angles for simultaneous diagonalization,” *SIAM Journal on Matrix Analysis and Applications*, vol. 17, no. 1, pp. 161–164, 1996. [32](#)
- [22] M. S. Pedersen, J. Larsen, U. Kjems, and L. C. Parra, *A survey of convolutive blind source separation methods*. Springer Verlag, 2007. [33](#)

- [23] K. Torkkola, “Blind separation of convolved sources based on information maximization,” in *Proc. of IEEE Signal Processing Society Workshop on Neural Networks for Signal Processing*, pp. 423–432, 1996. [33](#)
- [24] T. won Lee, A. J. Bell, and R. H. Lambert, “Blind separation of delayed and convolved sources,” *Advances in neural information processing systems*, vol. 9, pp. 758–764, 1997. [33](#)
- [25] S. Mallat, *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*. Academic Press, 2008. [33](#), [110](#)
- [26] H. Sawada, S. Araki, R. Mukai, and S. Makino, “Grouping separated frequency components by estimating propagation model parameters in frequency-domain blind source separation,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 15, pp. 1592–1604, july 2007. [34](#), [92](#)
- [27] D. T. Pham, C. Servière, and H. Boumaraf, “Blind separation of convolutive audio mixtures using nonstationarity,” in *Proceedings of ICA*, pp. 975–980, 2003. [34](#)
- [28] L. Parra and C. Spence, “Convolutive blind separation of non-stationary sources,” *IEEE Transactions on Speech and Audio Processing*, vol. 8, no. 3, pp. 320–327, 2000. [35](#), [73](#)
- [29] N. Mitianoudis and M. E. Davies, “Audio source separation of convolutive mixtures,” *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 5, pp. 489–497, 2003. [35](#), [67](#)
- [30] S. Amari, A. Cichocki, and H. H. Yang, “A new learning algorithm for blind signal separation,” in *Advances in Neural Information Processing Systems*, pp. 757–763, MIT Press, 1996. [35](#)
- [31] T.-W. Lee, M. Lewicki, M. Girolami, and T. Sejnowski, “Blind source separation of more sources than mixtures using overcomplete representations,” *IEEE Signal Processing Letters*, vol. 6, pp. 87–90, April 1999. [35](#)
- [32] O. Yilmaz and S. Rickard, “Blind separation of speech mixtures via time-frequency masking,” *IEEE Transactions on Signal Processing*, vol. 52, pp. 1830–1847, July 2004. [35](#), [39](#), [45](#), [104](#)
- [33] P. Georgiev, F. Theis, and A. Cichocki, “Sparse component analysis and blind source separation of underdetermined mixtures,” *Neural Networks, IEEE Transactions on*, vol. 16, pp. 992–996, july 2005. [35](#)
- [34] B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM journal on computing*, vol. 24, pp. 227–234, April 1995. [36](#)
- [35] A. M. Bruckstein, D. L. Donoho, and M. Elad, “From sparse solutions of systems of equations to sparse modeling of signals and images,” *SIAM Review*, vol. 51, no. 1, pp. 34–81, 2009. [37](#), [38](#)

BIBLIOGRAPHY

- [36] M. Van Hulle, "Clustering approach to square and non-square blind source separation," in *Proceedings of the IEEE Signal Processing Society Workshop Neural Networks for Signal Processing IX*, pp. 315–323, August 1999. [37](#)
- [37] M. Elad and A. M. Bruckstein, "A generalized uncertainty principle and sparse representation in pairs of bases," *IEEE Trans. on Inf. Theory*, vol. 48, pp. 2558–2567, September 2002. [38](#), [87](#)
- [38] A. Feuer and A. Nemirovski, "On sparse representation in pairs of bases," *IEEE Transactions on Information Theory*, vol. 49, pp. 1579–1581, jun. 2003. [38](#)
- [39] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Transactions on Information Theory*, vol. 49, pp. 3320–3325, dec. 2003. [38](#)
- [40] M. Aharon, M. Elad, and A. M. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear Algebra and its Applications*, vol. 416, no. 1, pp. 48–67, 2006. [38](#)
- [41] F. Jaillet, R. Gribonval, M. Plumbley, and H. Zayyani, "An l1 criterion for dictionary learning by subspace identification," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pp. 5482–5485, mar. 2010. [38](#)
- [42] P. Sprechmann and G. Sapiro, "Dictionary learning and sparse coding for unsupervised clustering," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pp. 2042–2045, march 2010. [38](#)
- [43] R. Xu and I. Wunsch, D., "Survey of clustering algorithms," *IEEE Transactions on Neural Networks*, vol. 16, pp. 645–678, May 2005. [40](#), [43](#)
- [44] P. O'Grady, B. Pearlmutter, and S. Rickard, "Survey of sparse and non-sparse methods in source separation," *IJIST*, vol. 15, no. 1, pp. 18–33, 2005. [42](#)
- [45] F. Abrard and Y. Deville, "Blind separation of dependent sources using the "time-frequency ratio of mixtures" approach," in *Proceedings of the Seventh International Symposium on Signal Processing and Its Applications*, vol. 2, pp. 81–84, July 2003. [43](#)
- [46] Y. Deville and M. Puigt, "Temporal and time-frequency correlation-based blind source separation methods. part i: Determined and underdetermined linear instantaneous mixtures," *Signal Processing*, vol. 87, no. 3, pp. 374–407, 2007. [44](#)
- [47] S. Arberet, R. Gribonval, and F. Bimbot, "A robust method to count and locate audio sources in a stereophonic linear instantaneous mixture," *Independent Component Analysis and Blind Signal Separation*, pp. 536–543, 2006. [44](#)
- [48] V. Reju, S. N. Koh, and I. Y. Soon, "Underdetermined convolutive blind source separation via time-frequency masking," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 18, pp. 101–116, jan. 2010. [46](#), [130](#)

- [49] H. Sawada, S. Araki, and S. Makino, “A two-stage frequency-domain blind source separation method for underdetermined convolutive mixtures,” in *Applications of Signal Processing to Audio and Acoustics, 2007 IEEE Workshop on*, pp. 139–142, oct. 2007. [46](#), [130](#)
- [50] H. Sawada, S. Araki, and S. Makino, “Underdetermined convolutive blind source separation via frequency bin-wise clustering and permutation alignment,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2010. [46](#)
- [51] M. Mandel, R. Weiss, and D. Ellis, “Model-based expectation-maximization source separation and localization,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 18, pp. 382–394, feb. 2010. [47](#)
- [52] K. Hoffman and R. Kunze, *Linear Algebra*. Prentice-Hall, Englewood Cliffs, New Jersey, USA, 2 ed., 1971. [49](#)
- [53] J. Tropp, “Just relax: convex programming methods for identifying sparse signals in noise,” *Information Theory, IEEE Transactions on*, vol. 52, pp. 1030–1051, march 2006. [50](#)
- [54] S. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” *Signal Processing, IEEE Transactions on*, vol. 41, pp. 3397–3415, dec 1993. [50](#)
- [55] Y. Pati, R. Rezaiifar, and P. Krishnaprasad, “Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition,” in *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, pp. 40–44 vol.1, nov 1993. [50](#)
- [56] D. L. Donoho, Y. Tsaig, I. Drori, and J. luc Starck, “Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit,” tech. rep., Stanford University, 2006. [50](#)
- [57] D. Needell and R. Vershynin, “Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit,” *Foundations of Computational Mathematics*, vol. 9, pp. 317–334, April 2009. [50](#)
- [58] D. Needell and R. Vershynin, “Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit,” *Selected Topics in Signal Processing, IEEE Journal of*, vol. 4, pp. 310–316, april 2010. [50](#)
- [59] S. Krstulovic and R. Gribonval, “MPTK: Matching Pursuit made tractable,” in *Proc. Int. Conf. Acoust. Speech Signal Process. (ICASSP’06)*, vol. 3, (Toulouse, France), pp. III–496 – III–499, May 2006. [50](#)
- [60] “Sparsify.” <http://www.personal.soton.ac.uk/tb1m08/sparsify/sparsify.html>. [50](#)

BIBLIOGRAPHY

- [61] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, December 1998. 51
- [62] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, March 2004. 51
- [63] “ ℓ^1 -magic.” <http://www.acm.caltech.edu/l1magic/>. 51
- [64] “Sparselab.” <http://sparselab.stanford.edu/>. 51
- [65] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming (web page and software).” <http://www.stanford.edu/~boyd/cvx>, June 2009. 51, 60, 94
- [66] I. Gorodnitsky and B. Rao, “Sparse signal reconstruction from limited data using FOCUSS: a re-weighted minimum norm algorithm,” *Signal Processing, IEEE Transactions on*, vol. 45, pp. 600–616, mar 1997. 51
- [67] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. pp. 267–288, 1996. 51
- [68] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *The Annals of Statistics*, vol. 32, no. 2, pp. pp. 407–451, 2004. 51
- [69] I. Daubechies, R. Devore, M. Fornasier, and C. S. Gšntšrk, “Iteratively reweighted least squares minimization for sparse recovery,” *Comm. Pure Appl. Math*, vol. 63, pp. 1–38, January 2010. 51
- [70] M. Elad, B. Matalon, J. Shtok, and M. Zibulevsky, “A wide-angle view at iterated shrinkage algorithms,” in *in SPIE Wavelet XII*, pp. 26–29, 2007. 51, 131
- [71] I. Damnjanovic, M. Davies, and M. Plumbley, “Smallbox - an evaluation framework for sparse representations and dictionary learning algorithms,” in *Latent Variable Analysis and Signal Separation*, vol. 6365 of *Lecture Notes in Computer Science*, pp. 418–425, 2010. 51
- [72] “SMALLBox.” <http://www.small-project.eu/software-data/smallbox>. 51
- [73] E. Candès and T. Tao, “Decoding by linear programming,” *IEEE Trans. Inf. Theory*, vol. 51, pp. 4203–4215, December 2005. 52
- [74] J. Tropp, “Greed is good: algorithmic results for sparse approximation,” *Information Theory, IEEE Transactions on*, vol. 50, pp. 2231 – 2242, oct. 2004. 52
- [75] R. Gribonval and P. Vandergheynst, “On the exponential convergence of matching pursuits in quasi-incoherent dictionaries,” *Information Theory, IEEE Transactions on*, vol. 52, pp. 255 – 261, jan. 2006. 52

- [76] R. Gribonval, H. Rauhut, K. Schnass, and P. Vandergheynst, “Atoms of all channels, unite! average case analysis of multi-channel sparse recovery using greedy algorithms,” *Journal of Fourier Analysis and Applications*, vol. 14, pp. 655–687, 2008. [52](#)
- [77] M. Elad, *Sparse and Redundant Representations*. Springer, 2010. [52](#)
- [78] Z. Ben-Haim, Y. Eldar, and M. Elad, “Coherence-based performance guarantees for estimating a sparse vector under random noise,” *Signal Processing, IEEE Transactions on*, vol. 58, pp. 5030–5043, oct. 2010. [52](#)
- [79] D. Donoho, ed., *On minimum entropy deconvolution*, in: *Applied Time-Series Analysis II*. Academic Press, 1981. [57](#)
- [80] J. Tugnait, “Identification and deconvolution of multichannel linear non-gaussian processes using higher order statistics and inverse filter criteria,” *IEEE Transactions on Signal Processing*, vol. 45, pp. 658–672, Mar. 1997. [57](#)
- [81] J. Tugnait, “Cumulant-based blind identification of linear multi-input-multi-output systems driven by colored inputs,” *IEEE Transactions on Signal Processing*, vol. 45, pp. 1543–1552, June 1997. [57](#)
- [82] Y. Sato, “A method of self-recovering equalization for multilevel amplitude-modulation systems,” *IEEE Transactions on Communications*, vol. 23, pp. 679–682, June 1975. [57](#)
- [83] G. Xu, H. Liu, L. Tong, and T. Kailath, “A least-squares approach to blind channel identification,” *IEEE Transactions on Signal Processing*, vol. 43, no. 12, pp. 2982–2993, 1995. [57](#), [58](#)
- [84] E. Moulines, P. Duhamel, J. Cardoso, and S. Mayrargue, “Subspace methods for the blind identification of multichannel FIR filters,” *IEEE Transactions on Signal Processing*, vol. 43, pp. 516–525, 1995. [57](#)
- [85] C. Avendano, J. Benesty, and D. R. Morgan, “A least squares component normalization approach to blind channel identification,” in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, (Washington, DC, USA), pp. 1797–1800, 1999. [58](#)
- [86] Y. A. Huang and J. Benesty, “Adaptive multi-channel least mean square and newton algorithms for blind channel identification,” *Signal Processing*, vol. 82, no. 8, pp. 1127–1138, 2002. [59](#)
- [87] Y. Huang and J. Benesty, “A class of frequency-domain adaptive approaches to blind multichannel identification,” *IEEE Transactions on Signal Processing*, vol. 51, pp. 11–24, Jan. 2003. [59](#)

BIBLIOGRAPHY

- [88] A. Aïssa-El-Bey, M. Grebici, K. Abed-Meraim, and A. Belouchrani, “Blind system identification using cross-relation methods: further results and developments,” in *Proceedings of the ISSPA*, vol. 1, pp. 649–652, 2003. [59](#)
- [89] A. Aïssa-El-Bey and K. Abed-Meraim, “Blind simo channel identification using a sparsity criterion,” in *Proc. of SPAWC*, pp. 271 – 275, 2008. [60](#), [131](#)
- [90] A. Aïssa-El-Bey, K. Abed-Meraim, and Y. Grenier, “Blind separation of underdetermined convolutive mixtures using their time - frequency representation,” *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 15, pp. 1540 –1550, jul. 2007. [60](#), [61](#)
- [91] W. Baumann, B. uwe Köhler, D. Kolossa, and R. Orglmeister, “Real time separation of convolutive mixtures,” in *Proc. of ICA, 2001*, pp. 65–69, 2001. [71](#), [73](#)
- [92] P. Smaragdis, “Blind separation of convolved mixtures in the frequency domain,” *Neurocomputing*, vol. 22, no. 1-3, pp. 21 – 34, 1998. [71](#), [73](#)
- [93] F. Asano, S. Ikeda, M. Ogawa, H. Asoh, and N. Kitawaki, “Combined approach of array processing and independent component analysis for blind separation of acoustic signals,” *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 3, pp. 204–215, 2003. [71](#), [73](#)
- [94] M. Z. Ikram and D. R. Morgan, “A beamforming approach to permutation alignment for multichannel frequency-domain blind source separation,” in *Proc. of ICASSP*, pp. 881–884, May 2002. [72](#), [73](#)
- [95] S. Kurita, H. Saruwatari, S. Kajita, K. Takeda, and F. Itakura, “Evaluation of blind signal separation method using directivity pattern under reverberant conditions,” in *Proceedings of IEEE International Conference the Acoustics, Speech, and Signal Processing*, pp. 3140–3143, 2000. [72](#), [73](#)
- [96] N. Mitianoudis and M. Davies, “Permutation alignment for frequency domain ica using subspace beamforming methods,” in *Independent Component Analysis and Blind Signal Separation* (C. G. Puntonet and A. Prieto, eds.), vol. 3195 of *Lecture Notes in Computer Science*, pp. 669–676, Springer Berlin / Heidelberg, 2004. [72](#), [73](#)
- [97] L. Parra and C. Alvino, “Geometric source separation: Merging convolutive source separation with geometric beamforming,” *IEEE Transactions on Speech and Audio Processing*, vol. 10, no. 6, pp. 352–362, 2002. [72](#), [73](#)
- [98] W. Wang, D. Cosker, Y. Hicks, S. Saneit, and J. Chambers, “Video assisted speech source separation,” in *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 425–428, March 2005. [72](#), [73](#)

- [99] W. Baumann, D. Kolossa, and R. Orglmeister, “Maximum likelihood permutation correction for convolutive source separation,” *Proc. of ICA ’03*, Jan 2003. [72](#), [73](#)
- [100] H. Gotanda, K. Nobu, T. Koya, K. Kaneda, and T. Ishibashi, “Permutation correction and speech extraction based on split spectrum through fastica,” *Proc. of ICA ’03*, Jan 2003. [72](#), [73](#)
- [101] M. S. Pedersen, J. Larsen, U. Kjems, and L. C. Parra, “A survey of convolutive blind source separation,” *Springer Handbook on Speech Processing and Speech Communication*, 2006. [72](#)
- [102] V. Soon, L. Tong, Y. Huang, and R. Liu, “A wideband blind identification approach to speech acquisition using a microphone array,” in *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, vol. 1, pp. 293 –296 vol.1, mar 1992. [73](#)
- [103] W. Baumann, D. Kolossa, and R. Orglmeister, “Beamforming-based convolutive source separation,” in *Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP ’03). 2003 IEEE International Conference on*, vol. 5, pp. V – 357–60 vol.5, april 2003. [73](#)
- [104] K. Diamantaras, A. Petropulu, and B. Chen, “Blind two-input-two-output FIR channel identification based on frequency domain second-order statistics,” *Signal Processing, IEEE Transactions on*, vol. 48, pp. 534 –542, feb 2000. [73](#)
- [105] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete information,” *IEEE Trans. Inf. Theory*, vol. 52, pp. 489–509, February 2006. [94](#)
- [106] C. Knapp and G. Carter, “The generalized correlation method for estimation of time delay,” *IEEE Trans. on Acoust., Speech and Signal Proc.*, vol. 24, no. 4, pp. 320 – 327, 1976. [122](#)
- [107] R. Vidal, Y. Ma, and J. Piazzi, “A new gpca algorithm for clustering subspaces by fitting, differentiating and dividing polynomials,” in *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 1, pp. I–510 – I–517 Vol.1, june- july 2004. [130](#)
- [108] R. Vidal, Y. Ma, and S. Sastry, “Generalized principal component analysis (GPCA),” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, pp. 1945 –1959, dec. 2005. [130](#)

Abstract

Blind source separation from underdetermined mixtures is usually a two-step process: the estimation of the mixing filters, followed by that of the sources. An enabling assumption is that the sources are sparse and disjoint in the time-frequency domain. For convolutive mixtures, the solution is not straightforward due to the permutation and scaling ambiguities. The sparsity of the filters in the time-domain is also an enabling factor for blind filter estimation approaches that are based on cross-relation. However, such approaches are restricted to the single source setting.

In this thesis, we jointly exploit the sparsity of the sources and mixing filters for blind estimation of sparse filters from stereo convolutive mixtures of several sources. First, we show why the sparsity of the filters can help solve the permutation problem in convolutive source separation, in the absence of scaling. Then, we propose a two-stage estimation framework, which is primarily based on the time-frequency domain cross-relation and an ℓ^1 minimisation formulation: a) a clustering step to group the time-frequency points where only one source is active, for each source; b) a convex optimisation step which estimates the filters. The resulting algorithms are assessed on audio source separation and filter estimation problems.

Résumé

La séparation aveugle de sources à partir de mélanges sous-déterminés se fait traditionnellement en deux étapes: l'estimation des filtres de mélange, puis celle des sources. L'hypothèse de parcimonie temps-fréquence des sources facilite la séparation, qui reste cependant difficile dans le cas de mélanges convolutifs à cause des ambiguïtés de permutation et de mise à l'échelle. Par ailleurs, la parcimonie temporelle des filtres facilite les techniques d'estimation aveugle de filtres fondées sur des corrélations croisées, qui restent cependant limitées au cas où une seule source est active.

Dans cette thèse, on exploite conjointement la parcimonie des sources et des filtres de mélange pour l'estimation aveugle de filtres parcimonieux à partir de mélanges convolutifs stéréophoniques de plusieurs sources. Dans un premier temps, on montre comment la parcimonie des filtres permet de résoudre le problème de permutation, en l'absence de problème de mise à l'échelle. Ensuite, on propose un cadre constitué de deux étapes pour l'estimation, basé sur des versions temps-fréquence de la corrélation croisée et sur la minimisation de norme ℓ^1 : a) un *clustering* qui regroupe les points temps-fréquence où une seule source est active; b) la résolution d'un problème d'optimisation convexe pour estimer les filtres. La performance des algorithmes qui en résultent est évaluée numériquement sur des problèmes de filtre d'estimation de filtres et de séparation de sources audio.