Analysis, algorithms and applications of compressed sensing

Yang, Zai

2014

Yang, Z. (2014). Analysis, algorithms and applications of compressed sensing. Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/59537>

<https://doi.org/10.32657/10356/59537>

ANALYSIS, ALGORITHMS AND APPLICATIONS OF COMPRESSED SENSING

YANG ZAI

YANG ZAI

**SCHOOL OF ELECTRICAL AND ELECTRONIC
ENGINEERING**

A thesis submitted to the Nanyang Technological University in fulfillment of
the requirement for the degree of Doctor of Philosophy

2014

Statement of Originality

I hereby certify that the work embodied in this thesis is the result of original research and has not been submitted for a higher degree to any other University or Institution.

.....

Date

.....

YANG ZAI

Acknowledgements

There were many people who helped me during my Ph.D, and I would like to take this opportunity to thank them.

First and foremost, I would like to express my appreciation and gratitude to my supervisors, Professor XIE Lihua and Professor ZHANG Cishen (currently at Swinburne University of Technology, Australia), for years of invaluable guidance, encouragement, and financial support. Professor Zhang was the first person who introduced me to the scientific world. Thanks to Professor Xie for his endless patience, thoughtful discussions, and the many opportunities and the freedom he has offered that allowed me to find the research topics I was most interested in. This thesis would not have been possible without their enthusiastic help.

Thanks also to Dr. LU Wenmiao, Dr. Arian Maleki and Professor Richard G. Baraniuk for their collaborations and fruitful discussions.

I would also like to thank several fellow graduate students, postdocs and visiting scholars for their help and support in both research and life. In particular, thanks to YOU Keyou, LIU Shuai, XIAO Nan, GAO Tingting, HU Jinwen, MENG Wei, YU Chengpu, DENG Jun, GUO Yuqian, Fang Jianyin, ZHOU Siwang, ZHAO Dongya, QU Xiaomei and LI Wuquan.

In the end, I would like to thank my parents and my wife, ZHANG Xia, for their endless love, support and self-sacrifices. They have always believed in me, even more than I myself do. Special thanks to Xia for her encouragement whenever I was frustrated. This thesis would not have been possible without her support.

Abstract

Compressed sensing (CS) is an emerging research area which studies the problem of recovering a high dimensional sparse signal from its low dimensional linear samples. While CS can acquire a signal with a sub-Nyquist sampling rate, it requires a nonlinear signal reconstruction strategy which is computationally expensive compared to the simple linear reconstruction procedure in the traditional Nyquist sampling case. Given the exact information of the sampling/sensing system, existing results have shown that the signal of interest can be accurately recovered via convex relaxation (ℓ_1 minimization) and/or other approaches provided the signal is sufficiently sparse.

To understand the sparsity-undersampling tradeoff is essential in CS, which is the first task of this thesis. Phase transition is currently the most precise tool to describe the tradeoff in the sense that it provides a condition that is both necessary and sufficient. While existing results are focused on the real valued signal and system case, we extend it to the complex case, discovering a new phase transition curve of ℓ_1 minimization which is positioned well above the real one. On the other hand, the sampling system may not be exactly known *a priori*, the signal recovery performance in the presence of system perturbations is also studied in this thesis. In particular, a structured sensing matrix perturbation is studied which has practical relevance. Under mild sufficient conditions, it is shown that a sparse signal can be recovered by ℓ_1 minimization and the recovery error is at most proportional to the measurement noise level, which is similar to the standard CS result. In the special noise free case, the recovery is exact provided that the signal is sufficiently sparse with respect to the perturbation level.

Secondly, computationally efficient algorithms for CS are developed in the thesis. ℓ_1 minimization has favorable guarantees of signal recovery accuracy but most solvers are slower than expected for practical problems of high dimension. In the thesis, two first-order algorithms are proposed for the ℓ_1 minimization problem in CS: one exactly solves the problem and the other is a relaxed version of the first one. The latter can be considered as a modified iterative soft thresholding algorithm and is easy to implement. Sparse Bayesian learning (SBL) is another popular approach to the sparse signal recovery. In SBL, the signal sparsity information is exploited by assuming a sparsity-inducing prior for the signal that is then estimated using Bayesian inference. In the thesis, a new sparsity-inducing prior is introduced and efficient algorithms are developed for the signal recovery. The main algorithm is shown to produce a sparser solution than existing SBL methods while preserving their desirable properties. Moreover, a Bayesian framework is introduced in the thesis for the CS problem with quantized measurements. All proposed algorithms are compared with state-of-the-art algorithms through numerical simulations.

At last, applications of CS to direction of arrival (DOA) estimation and motion correction in magnetic resonance imaging (MRI) are presented by exploiting the spatial sparsity of the source signals and the image sparsity in an appropriate domain, respectively. The proposed approach to the DOA estimation resolves the problem that the true DOAs may not lie on the discretized sampling grid which is typically required in existing methods. The new method can maintain high estimation accuracy even under a very coarse sampling grid. While existing applications of CS to MRI are focused on accelerating the scanning process by undersampling k -space data, a sparsity-driven approach to the motion correction in MRI is presented for the first time in the thesis.

Symbols and Acronyms

Symbols

$\mathbb{C}$	set of complex numbers
$\mathbb{R}$	set of real numbers
$\mathbb{R}^+$	set of positive real numbers
$\mathbb{R}^N$	set of real N dimensional vectors
$\mathbb{R}^{M \times N}$	set of real $M \times N$ dimensional matrices
$\mathbf{A}, \mathbf{B}, \Phi, \dots$	matrices
$\mathbf{x}, \mathbf{y}, \dots$	vectors
$x, y, \dots$	scalars
$\mathbf{A}_i$	i th column of $\mathbf{A}$
A_{ij}	ij th entry of $\mathbf{A}$
$\mathbf{A}^T$	transpose of $\mathbf{A}$
$\mathbf{A}^H$	complex transpose of $\mathbf{A}$
$\mathbf{A}'$	transpose for real-valued $\mathbf{A}$ and complex transpose for complex-valued $\mathbf{A}$
x_i	i th entry of $\mathbf{x}$
$\ \cdot\ _2$	spectral norm for a matrix, or Euclidean norm for a vector
$ \cdot $	determinant for a matrix, or absolute value for a scalar, or cardinality for a set
$\text{diag}(\cdot)$	a diagonal matrix if applied to a vector with its diagonal entries formed by the vector, or if applied to a matrix, a vector composed of diagonal entries of the matrix

$\text{supp}\{\mathbf{x}\}$	support of $\mathbf{x}$ defined as $\{i : x_i \neq 0\}$
$\ \mathbf{x}\ _p$	ℓ_p norm ($p \geq 1$) or pseudo-norm ($0 < p < 1$) of $\mathbf{x}$ defined as $(\sum_i x_i ^p)^{\frac{1}{p}}$
$\ \mathbf{x}\ _0$	ℓ_0 pseudo-norm of $\mathbf{x}$ defined as the number of nonzero entries
$\ \mathbf{x}\ _\infty$	ℓ_∞ norm of $\mathbf{x}$ defined as $\max\{ x_i \}$
$\langle \mathbf{x}, \mathbf{y} \rangle$	inner product of $\mathbf{x}$ and $\mathbf{y}$
$\mathbf{x} \succeq \mathbf{y}$	$x_i \geq y_i$ for all i
$E\{\cdot\}$	expectation
$Var\{\cdot\}$	variance

Acronyms

CS	compressed sensing
SSR	sparse signal representation
DOA	direction of arrival
MRI	magnetic resonance imaging
BP	basis pursuit
BPDN	basis pursuit denoising
OMP	orthogonal matching pursuit
IHT	iterative hard thresholding
IST	iterative soft thresholding
SBL	sparse Bayesian learning
OGSBL	off-grid SBL
RAAR	relaxed averaged alternating reflections
SRAAR	sparse RAAR
MAP	maximum <i>a posteriori</i>
AWGN	additive white Gaussian noise
i.i.d.	independently and identically distributed

Contents

Acknowledgements	i
Abstract	iii
Symbols and Acronyms	v
List of Contents	vii
List of Figures	xv
List of Tables	xxi
List of Algorithms	xxiii
1 Introduction	1
1.1 Motivation and Objective	1
1.2 Main Contributions	2
1.3 Outline of the Thesis	4

2 Literature Review	7
2.1 Deterministic Results	8
2.1.1 Convex Relaxation	10
2.1.2 Other Approaches	11
2.1.3 Number of Measurements for Guaranteed Performance	12
2.2 Phase Transition	14
2.3 Sparse Bayesian Learning	16
2.4 CS with Practical Considerations	18
2.4.1 Measurement Quantization	19
2.4.2 Sensing Matrix Perturbation	21
2.5 Applications of CS	22
2.5.1 DOA Estimation	22
2.5.2 MRI	23
2.6 Conclusion	24
 Part I Analysis and Algorithms	 25
3 Orthonormal Expansion ℓ_1-Minimization Algorithms	27
3.1 Preliminaries	29
3.1.1 Soft Thresholding Operator	29
3.1.2 Iterative Soft Thresholding	29
3.2 ONE-L1 Algorithms	30
3.2.1 Problem Description	30
3.2.2 ALM Based ONE-L1 Algorithms	31
 NANYANG TECHNOLOGICAL UNIVERSITY	 SINGAPORE

3.2.3	Relationship Between rONE-L1 and IST	37
3.2.4	Implementation Details	38
3.3	Numerical Simulations	38
3.3.1	Sparsity-Undersampling Tradeoff	38
3.3.2	Comparison with IST	39
3.3.3	Comparison with AMP, FPC-AS and NESTA in Noise-free Case	40
3.3.4	Recovery of 2D Image	42
3.4	Conclusion	43
4	Phase Transition in the Complex Domain	45
4.1	ONE-L1 Algorithms for CVCS	46
4.2	Phase Transition of CVCS	48
4.2.1	Evaluation of Complex Phase Transition via ONE-L1	48
4.2.2	Universality of Phase Transitions of CVCS	49
4.2.3	Connection to Block-Sparse CS	51
4.3	Rigorous Derivation of the Complex Phase Transition	54
4.4	Conclusion	55
5	Robustly Stable Signal Recovery with Structured Matrix Perturbation	57
5.1	Preliminary Results	59
5.1.1	Definitions	59
5.1.2	Stable Signal Recovery of Standard CS	61
5.1.3	Robust Signal Recovery in Perturbed CS	62
5.2	SP-CS: CS Subject to Structured Perturbation	63

5.2.1	Problem Description	63
5.2.2	Main Results	64
5.2.3	Interpretation of the Main Results	66
5.2.4	When is the D-RIP satisfied?	68
5.2.5	Relaxation of the Optimal Solution	69
5.3	Algorithms for P-BPDN	71
5.3.1	Special Case: Positive Signals	71
5.3.2	AA-P-BPDN: Alternating Algorithm for P-BPDN	72
5.3.3	Effectiveness of AA-P-BPDN	75
5.4	Numerical Simulations	76
5.5	Conclusion	79
6	Bayesian Compressed Sensing with New Sparsity-Inducing Prior	81
6.1	G-STG Prior	82
6.1.1	Mathematical Formulation	82
6.1.2	Intuitive Interpretation and Threshold Parameter Setting	84
6.2	Sparse Bayesian Learning for Signal Recovery	86
6.2.1	Bayesian Formulation	86
6.2.2	Bayesian Inference	87
6.2.3	Analysis of Global and Local Maxima	90
6.3	Fast Algorithm	92
6.3.1	Fast Greedy Algorithm	92
6.3.2	Analysis of Basis Selection Condition	95
6.4	Numerical Simulations	97
6.4.1	One-Dimensional Synthetic Signals	98
6.4.2	Images	105
6.5	Conclusion	106

7 Variational Bayesian Algorithm for Quantized Compressed Sensing 109

7.1	A New Framework for Quantized CS	110
7.1.1	A Unified Observation Model	110
7.1.2	Bayesian Formulation of Quantized CS	112
7.1.3	Relations with Existing Methods in Quantized CS	114
7.2	Q-VMP: Variational Message Passing for Quantized CS	117
7.2.1	Model Selection	117
7.2.2	Q-VMP Algorithm	118
7.2.3	The Noise Free Case	124
7.2.4	Pruning a Basis Function	125
7.3	Numerical Simulations	125
7.3.1	Experimental Setup	125
7.3.2	Model Efficiency	126
7.3.3	Performance Comparison in Multi-bit CS	128
7.3.4	Performance Comparison in 1-bit CS	131
7.4	Conclusion	133

Part II Applications 135

8 Off-Grid Direction of Arrival Estimation 137

8.1	Problem Description and Off-Grid Model	138
8.1.1	Problem Description	138
8.1.2	Off-Grid DOA Estimation Model	139
8.2	Off-Grid DOA Estimation Using P-BPDN	140

8.2.1	Off-Grid DOA Estimation Under SP-CS	140
8.2.2	Numerical Simulations	142
8.3	Off-Grid DOA Estimation Using SBL	143
8.3.1	Sparse Bayesian Formulation	145
8.3.2	Bayesian Inference	147
8.3.3	OGSBL-SVD	150
8.3.4	Source Power and DOA Estimation	152
8.3.5	Implementation Details	152
8.3.6	Numerical Simulations	153
8.4	Conclusion	160
9	Sparse MRI for Motion Correction	161
9.1	Problem Formulation of Motion Correction	162
9.2	Sparsity-Driven Motion Correction	163
9.2.1	Sparsity-Based Formulation	163
9.2.2	An Alternating Algorithm	164
9.2.3	Computation of $\mathcal{P}_1$	165
9.2.4	Approximate Computation of $\mathcal{P}_2$	167
9.2.5	Sparse RAAR for Motion Correction	168
9.3	Numerical Simulations	168
9.4	Conclusion	169
10	Conclusion and Future Work	171
10.1	Conclusion	171
10.2	Future Work	172

Appendix	175
A Proofs of Results in Chapter 5	175
A.1 Proof of Theorem 5.3	175
A.2 Proof of Theorems 5.4 and 5.5	175
A.3 Proof of Lemma 5.1	178
A.4 Proof of Theorem 5.7	179
A.5 Proof of Theorem 5.8	180
Author's Publications	183
Bibliography	187

List of Figures

2.1	Upper left: original Lena image. Upper right: Lena image reconstructed from the largest (in amplitude) 10% of the wavelet coefficients. Lower: sorted absolute values of the wavelet coefficients of the original Lena image.	9
2.2	Phase transition curves of CS with three different settings derived in [31–33].	15
3.1	Observed phase transitions of rONE-L1, and comparison with those of IST. Notice that, 1) the observed phase transitions of rONE-L1 with the recommended r strongly agree with the theoretical calculation based on BP; 2) rONE-L1 has significantly enlarged success phases compared with IST.	40
3.2	An example of 2D image reconstruction from noise-contaminated partial-DCT measurements. Upper left: original Mondrian image; upper right: sampling pattern. The lower three are reconstructed images respectively by rONE-L1 (lower left, error: 0.0741, time: 11.1 s), FPC-AS (lower middle, error: 0.0809, time: 11.4 s) and NESTA (lower right, error: 0.0649, time: 29.4 s).	44
4.1	Phase transition curves of CS with four different settings. The ones in the real settings are derived in [31–33] (see details in Section 2.2) and the one in the complex setting is discovered due to our work presented in this chapter.	46

- 4.2 Observed success rates of partial-Fourier sampling. Four images show results of rONE-L1 with $N = 1024, 2048, 4096$ and 8192 , respectively. Each pixel refers to the success rate S/M at the corresponding coordinate (δ, ρ) . Sampling ratio δ ranges from 0.02 (left) to 0.98 (right) with an interval 0.03 ; Sparsity ratio ρ ranges from $\rho^R(\delta) - 0.1$ (bottom) to $\rho^R(\delta) + 0.1$ (top) with an interval 0.01 with respect to each δ (only part of the experiment results are shown). The middle line in each image refers to the real phase transition. 50
- 4.3 Observed phase transitions of partial-Fourier sampling. The upper five curves are observed phase transitions estimated by eONE-L1 with $N = 1024$ and rONE-L1 with $N = 1024, 2048, 4096$ and 8192 , respectively. The lower is the theoretical real phase transition of BP. . . 50
- 4.4 Observed universality of phase transitions of BP in the complex domain. The signal length $N = 8192$ for partial-Fourier matrix ensemble and $N = 1000$ for the other three matrix ensembles. 52
- 4.5 Coincidence between our observed complex phase transition and the block-sparse phase transition with block size 2. The indirect phase transition curve refers to minimizing $\|\mathbf{x}_r\|_1$ subject to $\mathbf{A}\mathbf{x} = \mathbf{y}$ 53
- 5.1 Illustration of Corollary 5.2. The shaded band area refers to the feasible domain of BPDN. The triangular area, the intersection of the feasible domain and the ℓ_1 ball $\{\mathbf{x} : \|\mathbf{x}\|_1 \leq \|\mathbf{x}^o\|_1\}$, is the set of all good recoveries. 70
- 5.2 Signal and perturbation recovery errors with respect to the noise level ϵ with parameter settings $(N, M, K, r) = (200, 80, 10, 0.1)$. Both signal and β^o recovery errors of AA-P-BPDN for P-BPDN in SP-CS are proportional to ϵ 77

5.3	Signal and perturbation recovery errors with respect to the perturbation level in terms of r with parameter settings $(N, M, K, \epsilon) = (200, 80, 10, 0.5)$. The error of AA-P-BPDN for P-BPDN in SP-CS slowly increases with the perturbation level and is quite close to that of the ideal case of O-BPDN for a moderate perturbation.	78
5.4	Signal and perturbation recovery errors with respect to the number of measurements with parameter settings $(N, K, r, \epsilon) = (200, 10, 0.1, 0.2)$. AA-P-BPDN for P-BPDN in SP-CS has the best performance except the ideal case of O-BPDN.	79
5.5	Exact recovery of a positive sparse signal from noise-free measurements with $(M, N, K, r, \epsilon) = (200, 50, 10, 0.1, 0)$. PP-BPDN is solved by solving (P_1) . β^o and its recovery are shown only on the support of $\mathbf{x}^o$. Black circles: original signal and β^o ; red stars: recoveries. . . .	80
6.1	PDFs of the G-STG prior in the case of (a) $\epsilon = 0.1$ and varying τ and (b) $\tau = 1 \times 10^{-8}$ and varying ϵ with $\eta = 1$	85
6.2	Performance of the proposed algorithm with respect to different settings of τ with (a) & (b) Gaussian ensemble, and (c) & (d) uniform spherical ensemble.	100
6.3	Performance of the proposed algorithm with respect to different settings of ϵ	101
6.4	Performance comparison of the proposed algorithm with existing ones with $(N, K, \text{SNR}) = (512, 20, 25\text{dB})$	103
6.5	Performance comparison of the proposed algorithm with existing ones with $(N, M, K) = (512, 120, 20)$	104

6.6	The 512×512 Mondrian image (a) and its reconstructions using (b) linear reconstruction (RMSE = 0.1333) from $N = 4096$ wavelet samples and a multiscale CS scheme from $M = 2713$ linear measurements by (c) BP (RMSE = 0.1391, time = 44.7s and # nonzeros = 4096), (d) StOMP with FDR thresholding (RMSE = 0.1751, time = 10.3s and # nonzeros = 2014), (e) StOMP with FAR thresholding (RMSE = 0.1529, time = 15.7s and # nonzeros = 1088), (f) BCS (RMSE = 0.1448, time = 24.1s and # nonzeros = 1293), (g) Laplace (RMSE = 0.1427, time = 25.5s and # nonzeros = 1229) and (h) our proposed method (RMSE = 0.1440, time = 19.8s and # nonzeros = 1033).	107
7.1	Directed graphical model that encodes the joint PDF in (7.23) of the Bayesian model. Nodes denoted with circles correspond to random variables, while nodes denoted with squares correspond to parameters of the model. Doubly circled $\mathbf{z}$ is the observation while single circled nodes represent hidden variables.	118
7.2	An illustration of Lemma 7.1 with nonnegative entries of $\mathbf{e}$. The unit circle in the first quadrant composes of $\mathcal{D}_e$. Projections of four possible $\mathbf{v}$'s are shown.	123
7.3	Reconstruction SNRs of VMP algorithms implemented respectively based on the proposed observation model in (7.8), an existing one that couples the quantization error and measurement noise, and conventional CS (oracle-aided quantized CS) as an upper boundary. . . .	128
7.4	Performance comparison of Q-VMP, BPDN and L1RML with bit depth $B = 4$. SNR = 30dB. (a) Averaged reconstruction SNR; (b) Averaged support size of recovered signal; (c) Averaged CPU time. . .	129
7.5	Reconstruction SNRs of Q-VMP and L1RML with a saturated quantizer, as well as those with the unsaturated quantizer in Fig. 7.4. . . .	131

7.6	Performance comparison of Q-VMP, BIHT and CVXP in 1-bit CS. SNR = 10dB. (a) Averaged reconstruction SNR; (b) Averaged support size of recovered signal; (c) Averaged CPU time.	132
8.1	Histogram of θ estimation error for both sources using P-BPDN for SP-CS. Statistics including mean, variance and mean squared error (MSE) are shown.	144
8.2	Performance comparison of SP-CS and standard CS (SCS) on DOA estimation. Upper left: signal recovery in SP-CS; upper right: β^o recovery in SP-CS (shown only on the signal support); lower left: signal recovery in standard CS; lower right: signal amplitude versus θ (near the location of source 1) in SP-CS and standard CS.	144
8.3	Bias comparison of ℓ_1 -SVD, SBL-SVD and OGSBL-SVD in localizing two on-grid sources as a function of angular separation with source 1 fixed at 50° and SNR = 10 dB. Significant bias exhibits when using ℓ_1 -SVD for closely spaced sources while small asymptotic biases occur in SBL-SVD and OGSBL-SVD.	155
8.4	Histograms of the DOA estimation errors of (a) ℓ_1 -SVD, (b) SBL-SVD and (c) OGSBL-SVD for source 2. Source 1: 50° ; source 2: 60° . SNR = 10 dB. Grid interval $r = 1^\circ$. The errors of ℓ_1 -SVD range from -4° to 2° while SBL-SVD finds the DOAs without error in all the trials. All errors of OGSBL-SVD are within half the grid interval without exploiting the information that the true DOAs are on the sampling grid.	156
8.5	Bias comparison of OGSBL-SVD and ℓ_1 -SVD in localizing two sources as a function of angular separation with source 1 fixed at 50.3° and SNR = 10 dB. A small asymptotic bias occurs in OGSBL-SVD while a significant bias exhibits all the time when using ℓ_1 -SVD.	157
8.6	MSEs of OGSBL-SVD and ℓ_1 -SVD. The lower bound is for ℓ_1 -SVD regardless of the SNR.	158

9.1	Simulation results of sparsity-based motion correction on a Shepp- Logan phantom (row 1, 100 iterations), simulated human brain (row 2, 200 iterations), and noisy brain with small motions (row 3, 200 iterations) and large motions (row 4, 400 iterations).	170
-----	---	-----

List of Tables

3.1	Comparison Results of ONE-L1 Algorithms with State-of-the-art Methods. The column of “# calls $\mathbf{A}$ & $\mathbf{A}'$ ” presents the values of $\text{mean}(\text{minimum}, \text{maximum})$	42
4.1	Number of Failures and CPU Time Consumptions. The “Number of failures” column presents $\# \text{failures}(\# \text{problem instances solved})$	49
6.1	Averaged Relative MSEs, CPU Times and Number of Nonzero Entries (mean $\pm$ standard deviation) for Multiscale CS Reconstruction of the Mondrian Image.	106
8.1	Averaged Time Consumptions of ℓ_1 -SVD and OGSBL-SVD with Respect to SNR and r . Time unit: sec.	158
8.2	Averaged MSEs and CPU Times of STLS and OGSBL in the SMV Case with Respect to r	160

List of Algorithms

2.1	Orthogonal matching pursuit (OMP)	12
3.1	Exact ONE-L1 Algorithm via ALM Method	34
3.2	Relaxed ONE-L1 Algorithm	36
7.1	Q-VMP	124

Chapter 1

Introduction

1.1 Motivation and Objective

Compressed sensing (CS) [1, 2], also known as compressive sensing or compressive sampling, is an emerging technique which has attracted a lot of researchers due to its potentially wide applications. In conventional wisdom, a signal of interest needs to be sampled at the Nyquist rate for its complete recovery. Since most natural signals of interest are sparse/compressible under an appropriate basis, the sampled signal is then represented in the associated basis and only a few significant coefficients are retained for signal reconstruction with little energy loss, which is known as the compression technique (consider the JPEG 2000 standard for image compression). While the separate “sampling” and “compression” steps can prove to be wasteful of sensing and sampling resources, CS attempts to do them at a single step as its name suggests and accesses the success of signal recovery with the sparsity/compressibility. In particular, CS acquires a much smaller number of compressive linear measurements for reconstructing a high dimensional signal. The CS theory guarantees exact or accurate signal reconstruction given that the signal of interest is sufficiently sparse.

While progresses have been made in the past few years, many problems in CS are still open or need further studies. For example, current results in CS are mainly

focused on the ideal situation where the sensing system is exactly known. This may not be the case in practice. The theoretical performance of CS in the non-ideal case needs to be studied. Since CS can acquire a signal with a sub-Nyquist sampling rate, it requires a nonlinear signal reconstruction strategy which is typically computationally expensive. Hence, computationally efficient and accurate algorithms are needed for sparse signal recovery in CS. The inherent reason of the success of CS is the sparsity while sparse signals are ubiquitous in practice. That means, CS has potential applications in many research areas. In summary, objectives of the thesis include:

1. To analyze the performance of CS with/without system uncertainties;
2. To propose accurate and efficient algorithms for sparse signal recovery in CS;
3. To apply CS to solve practical problems.

1.2 Main Contributions

The main contributions of the thesis are listed as follows:

1. We discover a new phase transition curve, which provides precise sparsity-undersampling tradeoff of ℓ_1 minimization, for the complex valued CS problem where both the signal of interest and the sensing matrix are complex valued.
2. We analyze the CS problem subject to a structured perturbation in the sensing matrix which has practical relevance. Under mild conditions, we show that a sparse signal can be recovered by solving an ℓ_1 minimization problem and the recovery error is at most proportional to the measurement noise level, which is similar to the standard CS result. In the special noise free case, the recovery is exact provided that the signal is sufficiently sparse with respect to the perturbation level.

3. We study two popular approaches to sparse signal recovery in CS including convex relaxation and sparse Bayesian learning (SBL). For the convex relaxation method, a first-order algorithm is proposed which can be considered as a modified iterative soft thresholding algorithm and is easy to implement. For SBL, we introduce a new sparsity-inducing prior and develop efficient algorithms for signal recovery. We show that the main algorithm produces a sparser solution than existing SBL methods while preserving their desirable properties. Their performances are demonstrated for 1D synthetic data and 2D practical images.
4. We study sparse signal recovery from quantized compressive measurements. A novel variational Bayesian inference based algorithm is presented, which unifies the multi- and 1-bit CS processing and is applicable to various cases of noiseless/noisy environment and unsaturated/saturated quantizer. The quantization error is decoupled from the measurement noise, modeled as a random variable and then estimated jointly with the signal being recovered. The novel characterization of the quantization error results in improved performance of the algorithm.
5. We present an accurate and efficient algorithm for applying CS to direction of arrival (DOA) estimation. The new method resolves the problem that the true DOAs are not exactly on the discretized sampling grid and can maintain high estimation accuracy under a very coarse sampling grid.
6. We propose a novel application of CS to motion correction in magnetic resonance imaging (MRI). An iterative algorithm is proposed to jointly reconstruct the image content and estimate the motion which solves the motion correction problem without acquiring additional measurement data or modifying the sampling sequence.

Moreover, for the purpose of reproducible research, the Matlab codes of the algorithms developed in this thesis have been made downloadable online at the author's personal website, see <https://sites.google.com/site/zaiyang0248>.

1.3 Outline of the Thesis

In Chapter 2 we review some of the existing results on theories, algorithms and applications of CS. The theoretical results include theoretical guarantees of existing sparse signal recovery algorithms, which are deterministic in the sense that they can be applied as long as the sensing matrix satisfies some condition, and the phase transition theory that is specialized for random matrices. The algorithms mentioned in this chapter include those with theoretical guarantees, e.g., convex relaxation and greedy pursuits, and SBL. Applications are focused on the DOA estimation and MRI that are closely related to the thesis.

The following chapters that present our main contributions are divided into two parts: Part I (Chapters 3–7) is on analysis and algorithms of CS and Part II (Chapters 8 and 9) is on applications.

In Chapter 3 we propose first-order algorithms for convex relaxation (ℓ_1 minimization) approach to CS. We introduce a method called orthonormal expansion to reformulate the basis pursuit problem and present two algorithms (named as ONE-L1) based on convex optimization: one exactly solves the problem and the other is a relaxed version of the first one. The relationship between the relaxed version and iterative soft thresholding algorithm is studied. This chapter is mainly based on the paper that has been published in *IEEE Transactions on Signal Processing* [3].

In Chapter 4 we study the phase transition of ℓ_1 minimization in the complex setting where both the signal and sensing system are complex valued while existing phase transition theory is focused on the popular real setting. By extending the algorithms presented in Chapter 3 to the complex case and extensive numerical experiments we discover a new phase transition curve for the complex valued CS problem which complements the existing phase transition results. This chapter is mainly based on the paper that has been published in *IEEE Signal Processing Letters* [4]. Some of the results mentioned in this chapter were published in [5, 6].

In Chapter 5 we study the CS problem subject to a structured matrix perturbation. We propose a nonconvex ℓ_1 minimization problem to recover sparse signals in which

the parameters specifying the perturbation are jointly recovered. We provide theoretical guarantees of the proposed problem for recovery of sparse and compressible signals in both noisy and noiseless environments. An alternating direction algorithm is presented to solve the ℓ_1 minimization problem. This chapter is mainly based on the paper that has been published in *IEEE Transactions on Signal Processing* [7]. Some of the results mentioned in this chapter were published in [8].

In Chapter 6 we present a new SBL approach to sparse signal recovery in CS, where we introduce a new sparsity-inducing prior and develop efficient algorithms for signal recovery. We analyze the main algorithm and reveal that it produces a sparser solution than its existing SBL peers while preserving their desirable properties.

In Chapter 7 SBL is used to study the quantized CS problem where the new challenge is the existence of the quantization error. We present a Bayesian framework which unifies the multi- and 1-bit CS problems and is able to deal with various cases of noiseless/noisy environment and unsaturated/saturated quantizer. After that we propose a novel algorithm for signal recovery within the framework based on variational Bayesian inference. This chapter is mainly based on the paper that has been published in *IEEE Transactions on Signal Processing* [9]. Some of the results mentioned in this chapter were published in [10].

In Chapter 8 we study the DOA estimation problem from a sparse representation perspective and are concerned about the practical situation where the true DOAs are not on the discretized sampling grid. An off-grid model is obtained using linear approximation and then two approaches to off-grid DOA estimation are proposed based on the model. In the first approach, we show that the off-grid DOA estimation problem can be formulated as the perturbed CS problem studied in Chapter 5 and thus the ℓ_1 optimization approach presented in Chapter 5 can be applied straightforward. In the other, SBL is adopted to do the estimation in an iterative manner to reduce the computational workload. This chapter is mainly based on the papers that have been published in *IEEE Transactions on Signal Processing* [7, 11].

In Chapter 9 we solve the motion correction problem in MRI by exploiting the MR image sparsity/compressibility. We formulate the translational motions as unknown

parameters in the sensing system and seek for a motion estimate such that the motion-compensated image has the maximal sparsity/compressibility. This chapter is inspired by Chapter 5. Some of the results mentioned in this chapter have been published in [12].

In Chapter 10 we conclude the thesis and highlight some potential future research directions.

Chapter 2

Literature Review

The recently developed compressed sensing (CS) theory and methods [1, 2] can achieve acquisition of information contained within a huge volume of data using only a small number of measurement samples. Different from the classical Shannon-Nyquist sampling theorem which requires that the sampling frequency be twice as high as the bandwidth of a signal in order to reconstruct its complete information, the CS theory accesses the success of signal recovery with the sparsity. There are different definitions of sparse signals. We consider a simple one which corresponds to so-called exactly sparse signals. Approximately sparse (or compressible) signals will be defined elsewhere when needed. A signal $\mathbf{x}^o \in \mathbb{R}^N$ of length N is called K -sparse under a basis $\Psi \in \mathbb{R}^{N \times N}$ if all but at most a number of $K \ll N$ entries of its coefficient vector $\boldsymbol{\theta} \in \mathbb{R}^N$ are zeros with $\mathbf{x}^o = \Psi\boldsymbol{\theta}$. Sparse/compressible signals are ubiquitous in reality. As an example shown in Fig. 2.1, a natural image can be compressively represented in a wavelet domain with most of the energy concentrating at a few number of wavelet coefficients. Without loss of generality we assume that Ψ is an identity matrix, i.e., $\mathbf{x}$ is sparse in the canonical basis, since for a general basis Ψ it can be absorbed into the following introduced sensing matrix $\mathbf{A}$. Rather than observing directly the original sparse signal $\mathbf{x}^o$, a number of M , $K < M \ll N$, linear measurements are acquired in CS as

$$\mathbf{y} = \mathbf{A}\mathbf{x}^o, \quad (2.1)$$

where $\mathbf{y} \in \mathbb{R}^M$ is the measurement vector and $\mathbf{A} \in \mathbb{R}^{M \times N}$ denotes the sensing/measurement matrix. Without considering the signal structure, the recovery of $\mathbf{x}^o$ from $\mathbf{y}$ is generally ill-posed since there are less linear equations than the unknown variables. By accounting for the signal sparsity, naturally, we want to seek for the sparsest solution by solving the combinatorial optimization problem

$$\min \|\mathbf{x}\|_0, \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x}, \quad (2.2)$$

where $\|\mathbf{x}\|_0$ denotes the pseudo ℓ_0 norm which counts the number of nonzero entries of $\mathbf{x}$. Unfortunately, the ℓ_0 optimization cannot be used in practice since it is non-deterministic polynomial-time (NP) hard. Alternative, efficient approaches are needed. Chen *et al.* [13] propose to solve the basis pursuit (BP) problem

$$\min \|\mathbf{x}\|_1, \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{x}, \quad (2.3)$$

which is a tight convex relaxation of the ℓ_0 optimization problem and can be solved in a polynomial time. In fact, the use of ℓ_1 optimization for obtaining a sparse solution can date back to [14] for seismic data recovery. BP was empirically observed to give good performance, however, rigorous analysis had not been provided until the last decade.

2.1 Deterministic Results

We review some of the existing deterministic results in this section which mainly fall into the following two categories: one based on mutual coherence [15] and the other based on restricted isometry property (RIP) [16]. By deterministic results we mean that they can be applied once the sensing matrix satisfies some condition.

Definition 2.1. *The mutual coherence of a matrix $\mathbf{A}$, $\mu(\mathbf{A})$, is the largest absolute*

Figure 2.1: Upper left: original Lena image. Upper right: Lena image reconstructed from the largest (in amplitude) 10% of the wavelet coefficients. Lower: sorted absolute values of the wavelet coefficients of the original Lena image.

correlation between any two columns of $\mathbf{A}$, i.e.,

$$\mu(\mathbf{A}) = \max_{i \neq j} \frac{|\langle \mathbf{a}_i, \mathbf{a}_j \rangle|}{\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2}. \quad (2.4)$$

Definition 2.2. The k -restricted isometry constant (RIC) of a matrix $\mathbf{A}$, $\delta_k(\mathbf{A})$, is the smallest number such that the inequality

$$(1 - \delta_k(\mathbf{A})) \|\mathbf{v}\|_2^2 \leq \|\mathbf{A}\mathbf{v}\|_2^2 \leq (1 + \delta_k(\mathbf{A})) \|\mathbf{v}\|_2^2$$

holds for all k -sparse vectors $\mathbf{v}$. $\mathbf{A}$ is said to satisfy the k -RIP with constant $\delta_k(\mathbf{A})$ if $\delta_k(\mathbf{A}) < 1$.

Note that it is usually easy to compute the coherence of a matrix. But it may become infeasible to compute the RIC in the case of a large k (which is related to and typically larger than the signal sparsity) since the computational complexity increases exponentially with respect to k . Since the coherence considers only the

relationship between any two columns of the sensing matrix while the RIP is on much more (a number adaptive to the signal sparsity) columns, the RIP may provide better results for the sparse signal recovery, which will be demonstrated later. Without ambiguity, we write μ and δ_k instead of $\mu(\mathbf{A})$ and $\delta_k(\mathbf{A})$ for short in the remaining of the chapter.

2.1.1 Convex Relaxation

Convex relaxation (or ℓ_1 optimization) has been a most studied method for the sparse signal recovery. The following theorem is the first result which demonstrates the equivalence between the BP and the ℓ_0 optimization.

Theorem 2.1 ([15]). *Assume that $\|\mathbf{x}^o\|_0 \leq K$ and $\mu < \frac{1}{2K-1}$. Then $\mathbf{x}^o$ is the unique solution of the ℓ_0 optimization and the BP.*

Theorem 2.1 states that under some conditions the solution of the NP-hard ℓ_0 optimization can be efficiently obtained by solving the convex BP problem. An RIP-based result is as follows.

Theorem 2.2 ([17]). *Assume that $\|\mathbf{x}^o\|_0 \leq K$ and $\delta_{2K} < \sqrt{2} - 1$. Then $\mathbf{x}^o$ is the unique solution of the ℓ_0 optimization and the BP.*

After the work by Candès [17], the sufficient RIP condition has been improved, e.g., to $\delta_{2K} < \frac{3}{4+\sqrt{6}}$ by Foucart [18]. Other types of RIP conditions are also provided, e.g., $\delta_K < 0.307$ in [19]. Moreover, in the presence of measurement noises and/or when the signal of interest is not exactly sparse the convex relaxation can produce faithful signal estimates by solving the basis pursuit denoising (BPDN) problem

$$\min \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{Ax} - \mathbf{y}\|_2 \leq \epsilon, \quad (2.5)$$

where ϵ is an upper bound of the noise energy. Readers are referred to [17] or Chapter 5 for more details.

2.1.2 Other Approaches

At the beginning of the development of CS, convex relaxation methods (BP and BPDN) were solved using linear programming or second order cone programming while these solvers are usually slower than expected when dealing with practical problems of high dimension, e.g., the recovery of a 512×512 image. Other efficient approaches are developed as well for the sparse signal recovery and those with guaranteed performance include orthogonal matching pursuit (OMP) [20], compressive sampling matching pursuit (CoSaMP) [21], subspace pursuit (SP) [22] and iterative hard thresholding (IHT) [23]. We focus on OMP and IHT.

OMP is a greedy pursuit algorithm and has been studied by many researchers due to its concise implementation which is presented in Algorithm 2.1 below. It is proven in [24] that OMP provides the same theoretical guarantee as BP in terms of the mutual coherence, i.e., a K -sparse signal can be exactly recovered by OMP if $\mu < \frac{1}{2K-1}$. The first RIP-based analysis of OMP appears in [25], showing that if $\delta_{K+1} < \frac{1}{3\sqrt{K}}$, then OMP can exactly recover a K -sparse signal within K iterations. The RIP condition has been relaxed in later studies. For example, [26] shows that $\delta_{K+1} < \frac{1}{\sqrt{K+1}}$ is sufficient for the exact recovery within K iterations. However, it is clear that such RIP conditions where the upper bounds decay with respect to the sparsity level K are inferior to that for BP in Theorem 2.2 where an RIC bounded by a constant is sufficient. By allowing OMP to perform more than K iterations, [27] proves that the exact recovery of a K -sparse signal can be obtained within $30K$ iterations if $\delta_{31K} < \frac{1}{3}$, where the coefficients before K are not optimal and expected to be decreased in the future.

On the other hand, IHT starts with $\mathbf{x}^0 = \mathbf{0}$ and carries out the following iteration

$$\mathbf{x}^{l+1} = H_K \left(\mathbf{x}^l + \mathbf{A}^T (\mathbf{y} - \mathbf{A}\mathbf{x}^l) \right), \quad (2.6)$$

where $H_K(\mathbf{v})$ denotes a hard thresholding operator which sets all but the largest (in amplitude) K entries of a vector $\mathbf{v}$ to zero. The following result is proven in [23]:

Algorithm 2.1: Orthogonal matching pursuit (OMP)

Input: $\mathbf{A}, \mathbf{y}$
 1. initialize: $\mathbf{r}^0 = \mathbf{y}, \Omega = \emptyset, l = 0$
 2. **while** not converged **do**
 3. match: $\mathbf{h}^l = \mathbf{A}^T \mathbf{r}^l$
 4. identify: $\Omega^{l+1} = \Omega^l \cup \arg \max_j \{ |h_j^l| \}$
 5. update: $\mathbf{x}^{l+1} = \arg \min_{\mathbf{z}, \text{supp}(\mathbf{z}) \subset \Omega^{l+1}} \|\mathbf{y} - \mathbf{A}\mathbf{z}\|_2$
 $\mathbf{r}^{l+1} = \mathbf{y} - \mathbf{A}\mathbf{x}^{l+1}$
 $l = l + 1$
 6. **end while**
 Output: $\hat{\mathbf{x}} = \mathbf{x}^l = \arg \min_{\mathbf{z}, \text{supp}(\mathbf{z}) \subset \Omega^l} \|\mathbf{y} - \mathbf{A}\mathbf{z}\|_2$

if $\delta_{2K} < \frac{1}{\sqrt{8}}$, then IHT can exactly recover a K -sparse signal. Moreover, IHT has stable performance in the presence of measurement noises and when the signal of interest is not exactly sparse. Though IHT has comparable theoretical guarantees as BP with a more concise implementation, one disadvantage of IHT is that it requires the signal sparsity level in advance while it is not needed by BP. In fact, the same disadvantage is also shared among the above mentioned OMP, CoSaMP and SP.

2.1.3 Number of Measurements for Guaranteed Performance

In the last two subsections, various sufficient conditions are provided for BP and other approaches to guarantee the recovery of the sparse signal of interest, however, it is unclear how many measurements are sufficient for the sparse signal recovery. Random sensing matrices, especially sub-gaussian ones, are typically considered to answer the question.

Definition 2.3. A random variable X is called sub-gaussian iff there exist constants $c, C > 0$ such that $P(|X| > t) < C \exp(-ct^2)$.¹

Examples of sub-gaussian random variables include Gaussian and all bounded random variables. A sub-gaussian matrix means that all entries of the matrix are drawn

¹The notations c and C will be repetitively used in the thesis whose values may change from instance to instance.

independently and identically from a sub-gaussian distribution. The following theorems reveal the relationship between the mutual coherence as well as the RIP and the number of measurements.

Theorem 2.3 ([28]). *Suppose that the elements of matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ are drawn i.i.d. from a sub-gaussian distribution with mean zero and variance M^{-1} . Then there exists a constant $C > 0$ such that the following holds for every number $\bar{\mu} \in (0, 1)$ and $\epsilon \in (0, 1)$:*

$$\text{if } M \geq C\bar{\mu}^{-2} \log \frac{N}{\epsilon}, \text{ then } P(\mu < \bar{\mu}) \geq 1 - \epsilon^2.$$

Theorem 2.4 ([29]). *Suppose that the elements of matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ are drawn i.i.d. from a sub-gaussian distribution with mean zero and variance M^{-1} . Then there exists constants $c, C > 0$ such that the following holds for every sparsity level $1 \leq K \leq N$ and every number $\bar{\delta} \in (0, 1)$:*

$$\text{if } M \geq C\bar{\delta}^{-2}K \log \frac{eN}{K}, \text{ then } P(\delta_K < \bar{\delta}) \geq 1 - 2 \exp(-c\bar{\delta}^2 M).$$

For both convex relaxation and OMP, the results presented previously show that the mutual coherence of the sensing matrix is required to be on the order of K^{-1} , denoted by $\sim K^{-1}$, for recovering a sparse signal with the sparsity level of K . According to Theorem 2.3, a number of $M \sim K^2 \log N$ random sub-gaussian measurements are required to guarantee its recovery. Meanwhile, a (moderately small) constant RIC can guarantee the signal recovery according to the RIP-based results for all of the BP, OMP and IHT. As a result, a number of $M \sim K \log \frac{N}{K}$ random measurements are sufficient following from Theorem 2.4. So, the RIP-based results are superior to the ones based on the coherence in terms of the number of measurements required for guaranteed performance. Before closing this section, we note that the RIP-based

results provide only sufficient conditions and it is shown in [30] that such conditions can be quite conservative. In the next section, a sufficient and necessary condition will be introduced in the case of Gaussian sensing matrices.

2.2 Phase Transition

The theory of CS is mainly focused on studying how aggressively a sparse signal can be undersampled while still preserving all information for its recovery. The existing results based on the mutual incoherence and the RIP are introduced in the last section which provide sufficient conditions for the sparse signal recovery and are conservative in practice [30]. So far the theory of phase transition [4, 6, 31–33] is the most precise one owing to its necessary and sufficient condition for measuring the sparsity-undersampling tradeoff performance. Before our work [4–6], the study of phase transition is focused on the real setting where the sensing matrix and the sparse signal are real valued. Our main interest is the sparsity-undersampling tradeoff of the BP problem in (2.3). Some variants will be discussed later. Consider sensing matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ with i.i.d. random Gaussian entries. Denote by $\delta = M/N$ the sampling ratio and by $\rho = K/M$ the sparsity ratio, where K denotes the sparsity level of the signal $\mathbf{x}$ of interest. Then the sparsity-undersampling tradeoff of the BP can be precisely described by δ and ρ when the problem dimension approaches infinity. In particular, the plane of (δ, ρ) is divided by a phase transition curve into two phases, a ‘success’ phase where BP successfully recovers the sparse signal and a ‘failure’ phase where the original signal cannot be recovered by solving BP, both with an overwhelming probability (see Fig. 2.2).

So far three different approaches have been developed to calculate the phase transition curve of BP, including combinatorial geometry [31], null space method [32] and state evolution [33]. Different formulae are derived based on different approaches, however, they coincide with each other numerically. Since the formulae are complicated, their expressions are omitted. An interesting scenario is when

Figure 2.2: Phase transition curves of CS with three different settings derived in [31–33].

$\delta \rightarrow 0$, i.e., in the extreme undersampling case. It is shown in [31] and [33] that $\rho^R(\delta) \sim [2 \ln(1/\delta)]^{-1}$ as $\delta \rightarrow 0$, where $\rho^R(\delta)$ denotes the real phase transition of BP and by the expression $a \sim b$ we mean $a = b + o(b)$. This means that, the sparsity ratio stays positive whenever $\delta > 0$, i.e., for any given degree of undersampling, there exists a signal of sufficient sparsity leading to successful reconstruction using BP. Informally, we can conclude that a K -sparse signal can be reconstructed provided that the number of random Gaussian measurements $M \geq 2K \ln \frac{N}{M}$. Notice that the number of measurements required is about as the same order as that derived based on the RIP in Subsection 2.1.3 but with a significantly reduced coefficient. Numerical simulations have shown that the observed phase transition matches the theoretical curve even for modestly large N , e.g., $N = 1000$ [3, 34], where the phase transition for finite- N is defined as the value of ρ at which the original signal is successfully recovered with the probability of 50%. Moreover, it is practically observed that the Gaussian condition on $\mathbf{A}$ can be considerably relaxed, resulting in the well known observed universality of phase transitions [34].

Another two phase transition curves have also been derived respectively in other two

different settings. One imposes that the sparse signal of interest is positive valued and thus the following BP problem is solved:

$$\min_{\mathbf{x} \succeq \mathbf{0}} \sum_i x_i, \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{y}, \quad (2.7)$$

where $\succeq$ denotes $\geq$ with an elementwise operation. The other assumes that all of the nonzero entries of the sparse signal fall into the interval $(0, 1]$ and the signal is recovered by solving the following feasibility problem:

$$\text{find } \mathbf{x} \in [0, 1]^N, \text{ subject to } \mathbf{A}\mathbf{x} = \mathbf{y}. \quad (2.8)$$

The phase transition curves are displayed in Fig. 2.2. Since additional knowledge is taken into the problem solving, the phase transition of (2.7) is positioned well above that of the standard BP. But they are equal asymptotically when $\delta \rightarrow 0$ [35], i.e., the value of positivity constraints evaporates at extreme undersampling. The phase transition of (2.8) has a simple form: $\rho(\delta) = \max(0, 2 - \delta^{-1})$, i.e., the reconstruction is possible only when $\delta > 0.5$.

Complex data are involved in numerous CS applications, e.g., MRI [36], radar imaging [37] and source localization [11], where both the signal and sensing matrix are complex valued. Hence, sparsity-undersampling tradeoff in the complex setting deserves further studies. Chapter 4 presents our first result on the phase transition in the complex domain, where a new phase transition is discovered which is positioned well above the one in the real setting. Its theoretical derivation is later carried out in [6].

2.3 Sparse Bayesian Learning

Bayesian approaches to CS, known as Bayesian CS [38], was originated from the area of machine learning and introduced by Tipping [39] for obtaining sparse solutions to regression and classification tasks that use models which are linear in the parameters,

coined as relevance vector machine (RVM) or sparse Bayesian learning (SBL). SBL is built upon a statistical perspective where the sparsity information is exploited by assuming a sparsity-inducing prior for the signal of interest that is then estimated via Bayesian inference. Its theoretical performance is analyzed by Wipf and Rao [40]. After being introduced into CS by Ji *et al.* [38], this technique has become a popular approach to CS and other sparsity-related problems, see a list of references [11, 41–51]. It is noted that Chapters 6, 7 and 8 of the thesis are built upon this approach. Main research topics in SBL for CS include (a) developing efficient sparsity-inducing priors, (b) incorporating additional signal structures in the prior besides sparsity and (c) designing fast and accurate inference algorithms.

Many sparsity-inducing priors have been studied in the literature. In [44, 50], a spike-and-slab prior [52] is applied which is a mixture of a point mass at zero and a continuous distribution elsewhere and fits naturally for sparse signals. A typical inference scheme for such a prior is a Markov chain Monte Carlo (MCMC) method [53] due to the lack of closed-form expressions of Bayesian estimators. As a result, the inference process may suffer from computational difficulties because a large number of samples are required to approximate the posterior distribution and the convergence is typically slow. A popular class of sparsity-inducing priors is introduced in a hierarchical framework where a complex prior is composed of two or more simple distributions. For example, a Student’s t -prior (or Gaussian-inverse gamma prior) is used in the basic SBL [39] that is composed of a Gaussian prior in the first layer and a gamma prior in the second. A Laplace (Gaussian-exponential) prior is used in [46]. A Gaussian-gamma prior is recently studied in [54] that generalizes the Laplace prior. Two popular inference methods for the hierarchical priors are evidence procedure [55], e.g., in [39, 46], and variational Bayesian inference [56], e.g., in [48]. Both the methods are approximations of Bayesian inference since the exact inference is intractable. In an evidence procedure, the signal estimator has a simple expression in which some unknown hyperparameters are involved and estimated iteratively by maximizing their evidence. In a variational Bayesian inference method, the posterior distribution is approximated using some family of tractable

distributions followed by computation of an optimal distribution within the family. To circumvent high-dimensional matrix inversions, a fast algorithm framework is developed in [57] for evidence procedure and also adopted in [46]. When additional signal structures besides sparsity are known *a priori*, they can be incorporated into the signal prior for possibly improved performance. For example, [44] exploits the wavelet structure of images and [50] studies cluster structured sparse signals. SBL is also applied to the multiple-measurement-vectors (MMV) case where joint sparsity which exists among the two or more signals of interest is exploited [41, 43, 49].

Though formulated from a different perspective, SBL is related to other approaches to CS. Consider the observation model $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ where $\mathbf{e}$ represents an additive white Gaussian noise (AWGN). Let $p(\mathbf{x})$ be the prior for $\mathbf{x}$. Then, a maximum *a posteriori* (MAP) estimator of $\mathbf{x}$ coincides with a solution to a regularized least-squares problem with $-\log p(\mathbf{x})$ (up to a scale) being the regularization term, which bridges SBL and optimization methods. For example, a Laplace prior corresponds to the widely studied ℓ_1 minimization. A prior corresponding to the nonconvex ℓ_p ($0 < p < 1$) norm is studied in [45]. The fast algorithm in [57] is related to the greedy pursuit method. In fact, it is a greedy algorithm using a different support modification criterion. Unlike OMP and StOMP, it allows deletion of irrelevant basis vectors that may have been added to the solution support in earlier steps.

2.4 CS with Practical Considerations

The observation model is $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ in standard CS. The main problem to explore is under what conditions and how the signal $\mathbf{x}$ can be recovered given the measured samples $\mathbf{y}$, the sensing matrix $\mathbf{A}$ and possibly the noise information as well. It is noted that $\mathbf{y}$ and $\mathbf{A}$ are assumed to be exactly known, which may not be satisfied in practice. We focus on the gap between standard CS and practical situations in this section and review some existing results on these issues.

2.4.1 Measurement Quantization

The real-valued measurement $\mathbf{y}$ is assumed in the standard CS framework which has infinite bit precision. But it is known that measurement quantization is inevitable for practical considerations, e.g., data storage and transmission. During the quantization process, each continuous-valued measurement is quantized into some value in a finite set. A new challenge for the sparse signal recovery is thus the existence of quantization errors. The noise free case with a uniform unsaturated quantizer is studied in [1, 58, 59]. A solver with quantization consistency is recommended in [1] that corresponds to replacing the ℓ_2 norm in the BPDN problem by the ℓ_∞ norm. A BPDN solver is used in [58] which treats the quantization errors as additive noises with bounded energy. A family of solvers, named as basis pursuit dequantizer of moment p (BPDQ $_p$), which includes BPDN and that in [1] as special cases is studied in [59] where the ℓ_2 norm in BPDN is replaced by an ℓ_p norm with $2 \leq p \leq +\infty$. By characterizing the quantization errors as independent random variables uniformly distributed in a common interval, it is shown in [59] that the optimal signal recovery accuracy is obtained at some finite $p \geq 2$. But unfortunately, the optimal p cannot be explicitly given in practice. Note that the common uniform distribution assumption is crucial to obtain the results in [59]. As a result, it is unclear whether the results in [59] can be extended to a general quantizer case where such an assumption fails. It is obvious that both BPDN and BPDQ $_p$ are inappropriate in the case of a saturated quantizer since data saturation may lead to large or even unbounded quantization errors which deteriorate their performance. To deal with the data saturation, Laska *et al.* [60] propose two modified versions of BPDN which either reject saturated measurements or incorporate them into signal recovery. While quantization errors and measurement noises are coupled in most existing methods (some methods, e.g., BPDQ $_p$, consider only the noise free case to avoid such a problem), e.g., in [60], they are separately studied by Zymnis *et al.* [61] where the authors seek to find a signal estimate that maximizes the likelihood of the quantized measurements while the ℓ_1 norm is used to promote the signal sparsity. The resulting

algorithm is quoted as ℓ_1 -regularized maximum likelihood (L1RML).

An extreme case of quantized CS is so-called 1-bit CS where each quantized measurement keeps only the sign information of the real measurement and thus uses just one bit. The 1-bit CS framework is proposed in [62] and has attracted many research interests because it possesses many merits. For example, a 1-bit quantizer is a simple comparator that tests whether the measurement is above or below zero, leading to an easy implementation and a fast quantization process. A measurement noise can be neglected in 1-bit CS as long as it does not change the sign of the measurement. It is shown in [63] that to acquire just one bit for each measurement is optimal in the presence of heavy noises. The 1-bit case is quite different from the multi-bit case since all measurements are saturated in 1-bit CS and the signal scaling information is lost. A common approach to the signal scaling problem is to impose that the signal to be recovered has a fixed unit norm and then search for the signal on the unit hyper-sphere rather than in the whole space. Such a constraint is nonconvex and brings new challenges to algorithm design. Existing algorithms based on this constraint include renormalized fixed point iteration (RFPI) [62], matching sign pursuit (MSP) [64], restricted-step shrinkage (RSS) [65] and binary iterative hard thresholding (BIHT) [66]. Convex formulations of the 1-bit CS problem have been recently proposed by Plan and Vershynin [67, 68]. They show in [67] that a linear program can decode the noiseless case with guaranteed signal recovery accuracy under similar mild conditions as in conventional CS. In [68] they introduce a seemingly unrelated convex program for the noisy case and show similar results. It is noted that both the BIHT and the convex program in [68] which deal with the noisy 1-bit CS problem require the signal sparsity information (BIHT needs the sparsity level K and CVXP requires a proper upper bound for the signal's ℓ_1 norm). While the multi- and 1-bit CS problems are studied separately in the literature due to their big difference, we present a Bayesian framework that unifies both the cases in Chapter 7.

2.4.2 Sensing Matrix Perturbation

The sensing matrix $\mathbf{A}$ is assumed known *a priori* in standard CS, which is, however, not always the case in practical situations. For example, a matrix perturbation can be caused by quantization during implementation. In source separation [69, 70] the sensing matrix (or mixing system) is usually unknown and needs to be estimated, and thus estimation errors exist. In source localization such as direction of arrival (DOA) estimation [11, 71] and radar imaging [37, 72], the sensing matrix (or overcomplete dictionary) is constructed via discretizing one or more continuous parameters, and errors exist typically in the sensing matrix since the true source locations may not be on a discretized sampling grid.

There have been recent active studies on the CS problem where the sensing matrix is unknown or subject to an unknown perturbation. Gleichman and Eldar [73] introduce a concept named as blind CS where the sensing matrix is assumed unknown. Herman and Strohmer [74] analyze the effect of a general matrix perturbation and show that the signal recovery is robust to the perturbation in the sense that the recovery error grows linearly with the perturbation level. Similar robust recovery results are reported in [75, 76]. It is demonstrated in [76, 77] that the signal recovery may suffer from a large error under a large perturbation. In addition, the existence of recovery error caused by the perturbed sensing matrix is independent of the sparsity of the original signal. Algorithms have also been proposed to deal with sensing matrix perturbations. For example, Zhu *et al.* [78] propose a sparse total least-squares approach to alleviating the effect of perturbation where they explore the structure of the perturbation to improve recovery performance. In Chapter 8, we formulate the off-grid DOA estimation problem from a sparse Bayesian inference perspective and iteratively recover the source signal and the matrix perturbation. It is noted that existing algorithmic results provide no guarantees on the signal recovery accuracy when there exist perturbations in the sensing matrix. In Chapter 5, we study the compressed sensing problem subject to a structured perturbation in the sensing matrix. Theoretical guarantees are provided for the signal recovery

performance via an ℓ_1 optimization approach. Its practical relevance to the DOA estimation is studied in Chapter 8.

2.5 Applications of CS

It is always of great interest to acquire a signal using reduced measurements which can still guarantee its recovery since each measurement may have a potential time and/or power cost. The inherent reason of the success of CS is the signal sparsity. Fortunately, most natural signals can be sparsely represented in some basis, rendering that CS has been successfully applied to many problems in the past few years. Current applications include single-pixel camera [79], medical imaging [36,80], communications [81,82] and computational biology [83], to name just a few. In this section we focus on the applications of CS to array processing, specifically, direction of arrival (DOA) estimation [84], and magnetic resonance imaging (MRI) [85] which are closely related to the thesis.

2.5.1 DOA Estimation

Source localization using sensor arrays [84] has been an active research area for decades. We focus on the narrowband far-field source case where the wave front is assumed to be planar and the angle/direction information is to be estimated, known as the DOA estimation problem. MUSIC is the most successful method among conventional DOA estimation techniques in the case of a large number of snapshots and uncorrelated source signals [86,87]. The research on the DOA estimation has been advanced in recent years owing to the development of methods based on sparse signal representation (SSR) or CS [88,89]. Note that the DOAs of interest lie in a continuous range while CS is applicable to recovery of discrete signals. Hence, discretization of the DOA range is used in [88,89], where a fixed sampling grid is selected which serves as the set of all candidates of DOA estimates. Then by assuming that all true (unknown) DOAs are exactly on the selected grid, an SSR

problem can be formulated where the DOAs of interest constitute the support of the sparse signal to be recovered.

Though existing CS-based approaches have shown their improvements in the DOA estimation, e.g., their success in the case of limited snapshots, there are still difficulties in practical situations where the true DOAs are not on the sampling grid. On one hand, a dense sampling grid is necessary for accurate DOA estimation to reduce the gap between the true DOA and its nearest grid point since the estimated DOAs are constrained on the grid. On the other hand, too dense a sampling grid leads to a highly coherent matrix that violates the condition for the sparse signal recovery. Chapter 8 concerns this problem where we propose novel approaches to off-grid DOA estimation.

2.5.2 MRI

MRI is a noninvasive medical imaging technique to display detailed internal structure of the body. High-quality MR images enable physicians to assess the health condition of the body and determine the presence of certain diseases. MRI uses no ionizing radiation and does not have the associated potential harmful effects. Instead, a strong magnetic field is utilized to align the nuclear magnetization of hydrogen atoms in water of the body. Then MRI adopts radio frequency fields to change the alignment of the magnetization, which causes the hydrogen nuclei to generate a magnetic field that can be detected by the scanner. Finally, additional magnetic fields are used to manipulate this signal to gather enough information, known as k -space, to construct an image of the body. The k -space data can be formulated as a 2D or 3D Fourier transform of the image, and then inverse Fourier transform will reconstruct the image.

The patient is required to stay still during the k -space data collection process. Imaging speed is thus important in many MRI applications since patient motion happens very likely in a long scanning process. However, there are many limitations, including physical and physiological constraints, to speed up data collection in MRI.

Hence, one research area is to seek for methods to reconstruct high-quality images from reduced (sub-Nyquist sampled) k -space data while each of which has a potential time cost. Current applications of CS to MRI are focused on this by exploiting the image sparsity under, for example, a wavelet domain. A short list of publications include [36, 90–94].

Motion correction is another active research topic in MRI since a small amount of motions in practice may cause severe imaging artifacts. Existing methods [95–99] resolve the problem by acquiring more k -space data which prolongs the scanning process and may possibly introduce more motions. While CS has been vastly applied to MRI for imaging acceleration, few results have appeared on its application to the motion correction problem. Chapter 9 presents our first sparsity-based approach to MRI motion correction.

2.6 Conclusion

In this chapter we have reviewed many of the existing results of CS, including main approaches to the sparse signal recovery in CS (e.g., convex relaxation, greedy pursuit and SBL), their theoretical guarantees (those based on RIP, coherence and the phase transition theory) and practical applications of CS (e.g., DOA estimation and MRI). We have also discussed some open problems or those that need to be further studied. The following chapters are focused on solving these problems.

Part I

Analysis and Algorithms

Chapter 3

Orthonormal Expansion

ℓ_1 -Minimization Algorithms

Many existing results of CS have been reviewed in Chapter 2. Among both the theoretical and algorithmic results, convex relaxation (or ℓ_1 minimization) plays an important role for the sparse signal recovery. Since the ℓ_1 norm is non-smooth at the origin, algorithm design remains a major problem with the convex relaxation methods for dealing with practical problems of high dimension. In this chapter we study practically useful algorithms for the sparse signal recovery based on ℓ_1 minimization. Notice that besides the basis pursuit (BP) formulation in the noise-free case which takes the form

$$(\text{BP}) \quad \min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \mathbf{Ax} = \mathbf{y}, \quad (3.1)$$

two other formulations in the noisy case are basis pursuit denoising (BPDN) and a regularized form (QP) which are, respectively,

$$(\text{BPDN}) \quad \min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - \mathbf{Ax}\|_2 \leq \epsilon, \quad (3.2)$$

$$(\text{QP}) \quad \min_{\mathbf{x}} \left\{ \lambda \|\mathbf{x}\|_1 + \frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2 \right\}. \quad (3.3)$$

In the above, ϵ is an upper bound of the noise level and $\lambda > 0$ is a regularization parameter. It is well known that BPDN and QP are equivalent with appropriate choices of ϵ and λ . In general, λ decreases as ϵ decreases. In the limiting case of $\lambda, \epsilon \rightarrow 0$, both BPDN and QP converges to the BP for the noiseless case. Since QP seems to be easier to solve than BPDN and BP, many algorithms have been proposed to solve QP or based on that. Convex optimization methods include ℓ_1 -magic [100], interior-point method [101], conjugate gradient method [36], fixed-point continuation [102] and NESTA [103] based on Nesterov's smoothing technique [104] with continuation.

This chapter is based on [3] and mainly focused on solving the BP in noiseless CS. We consider the case where $\mathbf{A}\mathbf{A}'$ is an identity matrix, i.e., the rows of $\mathbf{A}$ are orthonormal. This consists of most fast transforms in CS, e.g., discrete cosine transform (DCT), discrete Fourier transform (DFT) and some wavelet transforms, e.g., Haar wavelet transform. This case has also been studied in other algorithms, e.g., NESTA. We introduce a novel method called orthonormal expansion to reformulate BP. The exact OrthoNormal Expansion ℓ_1 minimization (eONE-L1) algorithm is then proposed to exactly solve BP based on the augmented Lagrange multiplier (ALM) method [105, 106]. The relaxed ONE-L1 (rONE-L1) algorithm is further developed to simplify eONE-L1. We show that rONE-L1 converges at least exponentially and is in the form of modified iterative soft thresholding (see details in Subsection 3.1.2). In the case of strictly sparse signals and noise-free measurements, numerical simulations show that rONE-L1 has the same sparsity-undersampling tradeoff as BP. We further compare the proposed algorithms with existing methods including FPC-AS [107], AMP [33] and NESTA.

3.1 Preliminaries

3.1.1 Soft Thresholding Operator

For $w \in \mathbb{R}$, the soft thresholding of w with a threshold $\lambda \in \mathbb{R}^+$ is defined as:

$$S_\lambda(w) = \text{sgn}(w) \cdot (|w| - \lambda)^+, \quad (3.4)$$

where $(\cdot)^+ = \max(\cdot, 0)$ and

$$\text{sgn}(w) = \begin{cases} w/|w|, & w \neq 0; \\ 0, & w = 0. \end{cases} \quad (3.5)$$

The operator $S_\lambda(\cdot)$ can be extended to vector variables by its element-wise operation.

The soft thresholding operator can be applied to solve the following ℓ_1 -norm regularized least square problem [108], i.e.,

$$S_\lambda(\mathbf{w}) = \arg \min_{\mathbf{v}} \left\{ \lambda \|\mathbf{v}\|_1 + \frac{1}{2} \|\mathbf{w} - \mathbf{v}\|_2^2 \right\}, \quad (3.6)$$

where $\mathbf{v}, \mathbf{w} \in \mathbb{R}^M$, and $S_\lambda(\mathbf{w})$ is the unique minimizer.

3.1.2 Iterative Soft Thresholding

One popular approach to QP is iterative soft thresholding (IST) of the form (starting from $\mathbf{x}_0 = 0$) [108, 109]

$$\mathbf{x}_{t+1} = S_\lambda(\mathbf{x}_t + \mathbf{A}'\mathbf{z}_t), \quad \mathbf{z}_t = \mathbf{y} - \mathbf{A}\mathbf{x}_t, \quad (3.7)$$

where $'$ denotes the transpose operator. IST has a concise form and is easy to implement, but its convergence can be very slow [109], especially for small λ . To improve its speed, a fixed-point continuation (FPC) strategy is presented in [102], where λ is decreased in a continuation scheme and a q -linear convergence rate is

achieved. Further, FPC-AS [107] improves the performance of FPC by introducing an active set inspired by greedy pursuit methods. An alternative approach to improving the speed of IST is to use an aggressive continuation where λ is decreased in each iteration:

$$\mathbf{x}_{t+1} = S_{\lambda_t}(\mathbf{x}_t + \mathbf{A}'\mathbf{z}_t), \quad \mathbf{z}_t = \mathbf{y} - \mathbf{A}\mathbf{x}_t. \quad (3.8)$$

However, the algorithm of this form typically has a worse sparsity-undersampling tradeoff than BP [110]. In the case of recovering exactly sparse signals, such a disadvantage is overcome by approximate message passing (AMP) [33] which has the same sparsity-undersampling tradeoff as BP and is a modified IST:

$$\mathbf{x}_{t+1} = S_{\lambda_t}(\mathbf{x}_t + \mathbf{A}'\mathbf{z}_t), \quad \mathbf{z}_t = \mathbf{y} - \mathbf{A}\mathbf{x}_t + \frac{N \|\mathbf{x}_t\|_0}{M} \mathbf{z}_{t-1}. \quad (3.9)$$

3.2 ONE-L1 Algorithms

3.2.1 Problem Description

Consider the ℓ_1 -minimization problem BP with the sampling matrix $\mathbf{A}$ satisfying that $\mathbf{A}\mathbf{A}' = \mathbf{I}$, where $\mathbf{I}$ is an identity matrix. We say that $\mathbf{A}$ is a partially orthonormal matrix hereafter as its rows are usually randomly selected from an orthonormal matrix in practice, e.g. partial-DCT matrix. Hence, there exists another partially orthonormal matrix $\mathbf{B} \in \mathbb{R}^{(N-M) \times N}$, whose rows are orthogonal to those of $\mathbf{A}$, such that $\Phi = \begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}$ is orthonormal. Let $\mathbf{p} = \Phi\mathbf{x}$. The BP is then equivalent to

$$(\text{BP}^o) \quad \min_{\mathbf{x}, \mathbf{p}, \Gamma(\mathbf{p})=\mathbf{y}} \|\mathbf{x}\|_1, \quad \text{subject to } \Phi\mathbf{x} = \mathbf{p},$$

where $\Gamma(\mathbf{p})$ is an operator projecting the vector $\mathbf{p}$ onto its first M entries.

In (BP^o) , the sampling matrix $\mathbf{A}$ is expanded into an orthonormal matrix Φ . It corresponds to the scenario where the full sampling is carried out with the sampling

matrix Φ and $\mathbf{p}$ is the vector containing all measurements. Note that only $\mathbf{y}$, as a part of $\mathbf{p}$, is actually observed. To expand the sampling matrix $\mathbf{A}$ into an orthonormal matrix Φ is a key step to show that the ALM method exactly solves (BP^o) and, hence, BP. The next subsection describes the proposed algorithm, referred to as orthonormal expansion ℓ_1 -minimization (ONE-L1).

3.2.2 ALM Based ONE-L1 Algorithms

In this subsection we solve the ℓ_1 -minimization problem (BP^o) using the ALM method [105, 106]. The ALM method is similar to the quadratic penalty method except an additional Lagrange multiplier term. Compared with the quadratic penalty method, the ALM method has some salient properties, e.g., the ease of parameter tuning and the convergence speed. The augmented Lagrangian function is

$$\mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu) = \|\mathbf{x}\|_1 + \langle \mathbf{p} - \Phi \mathbf{x}, \boldsymbol{\nu} \rangle + \frac{\mu}{2} \|\mathbf{p} - \Phi \mathbf{x}\|_2^2, \quad (3.10)$$

where Lagrange multiplier $\boldsymbol{\nu} \in \mathbb{R}^N$, $\mu \in \mathbb{R}^+$ and $\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}'\mathbf{v} \in \mathbb{R}$ is the inner product of $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$. (3.10) can be expressed as follows:

$$\mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu) = \|\mathbf{x}\|_1 + \frac{\mu}{2} \|\mathbf{p} - \Phi \mathbf{x} + \mu^{-1} \boldsymbol{\nu}\|_2^2 - \frac{1}{2\mu} \|\boldsymbol{\nu}\|_2^2. \quad (3.11)$$

Subsequently, we have the following optimization problem (SP) :

$$(SP) \quad \min_{\mathbf{x}, \mathbf{p}, \Gamma(\mathbf{p})=\mathbf{y}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu).$$

Instead of solving (SP) , let us consider the two related problems

$$(SP_1) \quad \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu),$$

and

$$(SP_2) \quad \min_{\mathbf{p}, \Gamma(\mathbf{p})=\mathbf{y}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu).$$

Note that problem (SP_1) is similar to the ℓ_1 -regularized problem in (3.6). In general, (SP_1) cannot be directly solved using the soft thresholding operator as that in (3.6) since there is a matrix product of Φ and $\mathbf{x}$ in the term of ℓ_2 -norm. However, the soft thresholding operator does apply to the special case where Φ is orthonormal. Given $\|\Phi \mathbf{u}\|_2 = \|\mathbf{u}\|_2$ for any $\mathbf{u} \in \mathbb{R}^N$, we can apply the soft thresholding to obtain

$$S_{\mu^{-1}}(\Phi'(\mathbf{p} + \mu^{-1}\boldsymbol{\nu})) = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu). \quad (3.12)$$

To solve (SP_2) , we let $\partial_{\bar{\Gamma}(\mathbf{p})} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu) = 0$ to obtain $\bar{\Gamma}(\mathbf{p}) = \bar{\Gamma}(\Phi \mathbf{x} - \mu^{-1}\boldsymbol{\nu})$, i.e.,

$$\begin{bmatrix} \mathbf{y} \\ \bar{\Gamma}(\Phi \mathbf{x} - \mu^{-1}\boldsymbol{\nu}) \end{bmatrix} = \arg \min_{\mathbf{p}, \bar{\Gamma}(\mathbf{p})=\mathbf{b}} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu), \quad (3.13)$$

where $\bar{\Gamma}(\cdot)$ is the operator projecting the variable to its last $N - M$ entries. As a result, an iterative solution of (SP) is stated in the following lemma.

Lemma 3.1. *For fixed $\boldsymbol{\nu}$ and μ , the iterative algorithm given by*

$$\mathbf{x}^{j+1} = S_{\mu^{-1}}(\Phi'(\mathbf{p}^j + \mu^{-1}\boldsymbol{\nu})), \quad (3.14)$$

$$\mathbf{p}^{j+1} = \begin{bmatrix} \mathbf{y} \\ \bar{\Gamma}(\Phi \mathbf{x}^{j+1} - \mu^{-1}\boldsymbol{\nu}) \end{bmatrix} \quad (3.15)$$

converges to an optimal solution of (SP) .

Proof. Denote $\mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}, \mu)$ as $\mathcal{L}(\mathbf{x}, \mathbf{p})$, for simplicity. By the optimality and uniqueness of $\mathbf{x}^{j+1}$ and $\mathbf{p}^{j+1}$, we have $\mathcal{L}(\mathbf{x}^{j+1}, \mathbf{p}^{j+1}) \leq \mathcal{L}(\mathbf{x}^j, \mathbf{p}^j)$ and the equality holds if and only if $(\mathbf{x}^{j+1}, \mathbf{p}^{j+1}) = (\mathbf{x}^j, \mathbf{p}^j)$. Hence, the sequence $\{\mathcal{L}(\mathbf{x}^j, \mathbf{p}^j)\}$ is bounded and converges to a constant L^* , i.e., $\mathcal{L}(\mathbf{x}^j, \mathbf{p}^j) \rightarrow L^*$ as $j \rightarrow +\infty$. Since the sequence $\{\mathbf{x}^j\}$ is also bounded by $\|\mathbf{x}^j\|_1 \leq \mathcal{L}(\mathbf{x}^j, \mathbf{p}^j) + \frac{1}{2\mu} \|\boldsymbol{\nu}\|_2^2$, there exists a sub-sequence $\{\mathbf{x}^{j_i}\}_{i=1}^{+\infty}$ such that $\mathbf{x}^{j_i} \rightarrow \mathbf{x}_s^*$ as $i \rightarrow +\infty$, where $\mathbf{x}_s^*$ is an accumulation point of $\{\mathbf{x}^j\}$. Correspondingly, $\mathbf{p}^{j_i} \rightarrow \mathbf{p}_s^* = \begin{bmatrix} \mathbf{y} \\ \bar{\Gamma}(\Phi \mathbf{x}_s^* - \mu^{-1}\boldsymbol{\nu}) \end{bmatrix}$ and $\mathcal{L}(\mathbf{x}_s^*, \mathbf{p}_s^*) = L^*$.

We then show that $(\mathbf{x}_s^*, \mathbf{p}_s^*)$ is a fixed point of the algorithm. Since it holds that $\mathbf{p}_s^* = \arg \min_{\mathbf{p}, \Gamma(\mathbf{p})=\mathbf{y}} \mathcal{L}(\mathbf{x}_s^*, \mathbf{p})$, we need only to show $\mathbf{x}_s^* = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{p}_s^*)$. Suppose there exists $\bar{\mathbf{x}}_s \neq \mathbf{x}_s^*$ such that $\bar{\mathbf{x}}_s = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{p}_s^*)$. By (3.14), (3.15) and [108, Lemma 2.2], we have $\|\mathbf{x}^{j_i+1} - \bar{\mathbf{x}}_s\|_2 \leq \|\mathbf{x}^{j_i} - \mathbf{x}_s^*\|_2 \rightarrow 0$, i.e., $\mathbf{x}^{j_i+1} \rightarrow \bar{\mathbf{x}}_s$, as $i \rightarrow +\infty$. Meanwhile, $\mathbf{p}^{j_i+1} \rightarrow \bar{\mathbf{p}}_s$. Hence, $\mathcal{L}(\mathbf{x}^{j_i+1}, \mathbf{p}^{j_i+1}) \rightarrow \mathcal{L}(\bar{\mathbf{x}}_s, \bar{\mathbf{p}}_s) < L^*$, which is because $(\bar{\mathbf{x}}_s, \bar{\mathbf{p}}_s) \neq (\mathbf{x}_s^*, \mathbf{p}_s^*)$ and contradicts $\mathcal{L}(\mathbf{x}^j, \mathbf{p}^j) \rightarrow L^*$, resulting in that $(\mathbf{x}_s^*, \mathbf{p}_s^*)$ is a fixed point. Moreover, it follows from $\|\mathbf{x}^{j_i+q} - \mathbf{x}_s^*\|_2 \leq \|\mathbf{x}^{j_i} - \mathbf{x}_s^*\|_2 \rightarrow 0$ for any positive integer q , that $\mathbf{x}^j \rightarrow \mathbf{x}_s^*$, as $j \rightarrow +\infty$.

Note that orthonormal matrix $\Phi = \begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}$ and $\Phi' \Phi = \mathbf{A}' \mathbf{A} + \mathbf{B}' \mathbf{B} = \mathbf{I}$. We can obtain

$$\mathbf{x}_s^* = S_{\mu^{-1}}(\mathbf{x}_s^* + \mathbf{A}'(\mathbf{y} + \mu^{-1}\Gamma(\boldsymbol{\nu}) - \mathbf{A}\mathbf{x}_s^*)). \quad (3.16)$$

Meanwhile, (SP) is equivalent to

$$\begin{aligned} \min_{\mathbf{x}} \mathcal{L} \left(\mathbf{x}, \begin{bmatrix} \mathbf{y} \\ \bar{\Gamma}(\Phi \mathbf{x} - \mu^{-1}\boldsymbol{\nu}) \end{bmatrix} \right) \\ = \|\mathbf{x}\|_1 + \frac{\mu}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y} - \mu^{-1}\Gamma(\boldsymbol{\nu})\|_2^2 - \frac{1}{2\mu} \|\boldsymbol{\nu}\|_2^2. \end{aligned} \quad (3.17)$$

By [108, Proposition 3.10], $\mathbf{x}_s^*$ is an optimal solution of the problem in (3.17) and equivalently, $(\mathbf{x}_s^*, \mathbf{p}_s^*)$ is an optimal solution of (SP) . $\blacksquare$

Remark 3.1.

- (1) Lemma 1 shows that to solve problem (SP) is equivalent to solve, iteratively, problems (SP_1) and (SP_2) .
- (2) Reference [108, Proposition 3.10] only deals with the special case $\|\mathbf{A}\|_2 < 1$ and it is, in fact, straightforward to extend the result to arbitrary $\mathbf{A}$.

Following from the framework of the ALM method [106] and Lemma 3.1, the ALM based ONE-L1 algorithm is outlined in Algorithm 3.1, where $(\mathbf{x}_t^*, \mathbf{p}_t^*)$ is the optimal solution to (SP) in the t th iteration and $\boldsymbol{\nu}_t^*$ is the corresponding Lagrange multiplier.

The convergence of Algorithm 3.1 is stated in the following theorem.

Algorithm 3.1: Exact ONE-L1 Algorithm via ALM Method

Input: Expanded orthonormal matrix Φ and observed sample data $\mathbf{y}$.

1. $\mathbf{x}_0^* = \mathbf{0}$; $\mathbf{p}_0^* = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}$; $\boldsymbol{\nu}_0^* = \mathbf{0}$; $\mu_0 > 0$; $t = 0$.
 2. while not converged do
 3. Lines 4-9 solve $(\mathbf{x}_{t+1}^*, \mathbf{p}_{t+1}^*) = \arg \min_{(\mathbf{x}, \mathbf{p}, \Gamma(\mathbf{p})=\mathbf{y})} \mathcal{L}(\mathbf{x}, \mathbf{p}, \boldsymbol{\nu}_t^*, \mu_t)$;
 4. $\mathbf{x}_{t+1}^0 = \mathbf{x}_t^*$, $\mathbf{p}_{t+1}^0 = \mathbf{p}_t^*$, $j = 0$;
 5. while not converged do
 6. $\mathbf{x}_{t+1}^{j+1} = S_{\mu_t^{-1}}(\Phi'(\mathbf{p}_{t+1}^j + \mu_t^{-1}\boldsymbol{\nu}_t^*))$;
 7. $\mathbf{p}_{t+1}^{j+1} = \begin{bmatrix} \mathbf{y} \\ \bar{\Gamma}(\Phi \mathbf{x}_{t+1}^{j+1} - \mu_t^{-1}\boldsymbol{\nu}_t^*) \end{bmatrix}$;
 8. set $j = j + 1$;
 9. end while
 10. $\boldsymbol{\nu}_{t+1}^* = \boldsymbol{\nu}_t^* + \mu_t(\mathbf{p}_{t+1}^* - \Phi \mathbf{x}_{t+1}^*)$;
 11. choose $\mu_{t+1} > \mu_t$;
 12. set $t = t + 1$;
 13. end while
- Output: $(\mathbf{x}_t^*, \mathbf{p}_t^*)$.
-

Theorem 3.1. Any accumulation point $(\mathbf{x}^*, \mathbf{p}^*)$ of sequence $\{(\mathbf{x}_t^*, \mathbf{p}_t^*)\}_{t=1}^{+\infty}$ of Algorithm 3.1 is an optimal solution of (BP^0) and the convergence rate with respect to the outer iteration loop index t is at least $O(\mu_{t-1}^{-1})$ in the sense that

$$|\|\mathbf{x}_t^*\|_1 - \mathbf{x}^\dagger| = O(\mu_{t-1}^{-1}),$$

where $\mathbf{x}^\dagger = \|\mathbf{x}^*\|_1$.

Proof. We first show that the sequence $\{\boldsymbol{\nu}_t^*\}$ is bounded. By the optimality of $(\mathbf{x}_{t+1}^*, \mathbf{p}_{t+1}^*)$ we have

$$0 \in \partial_{\mathbf{x}} \mathcal{L}(\mathbf{x}_{t+1}^*, \mathbf{p}_{t+1}^*, \boldsymbol{\nu}_t^*, \mu_t) = \partial \|\mathbf{x}_{t+1}^*\|_1 - \Phi' \boldsymbol{\nu}_{t+1}^*, \quad (3.18)$$

$$0 = \partial_{\bar{\Gamma}(\mathbf{p})} \mathcal{L}(\mathbf{x}_{t+1}^*, \mathbf{p}_{t+1}^*, \boldsymbol{\nu}_t^*, \mu_t) = \bar{\Gamma}(\boldsymbol{\nu}_{t+1}^*), \quad (3.19)$$

where $\partial_{\mathbf{x}}$ denotes the partial differential operator with respect to $\mathbf{x}$ resulting in a set of subgradients. Hence, $\Phi' \boldsymbol{\nu}_{t+1}^* \in \partial \|\mathbf{x}_{t+1}^*\|_1$. It follows that $\|\Phi' \boldsymbol{\nu}_{t+1}^*\|_\infty \leq 1$ and

$\{\boldsymbol{\nu}_t^*\}$ is bounded. By $\mathbf{x}^\dagger \geq \mathcal{L}(\mathbf{x}_{t+1}^*, \mathbf{p}_{t+1}^*, \boldsymbol{\nu}_t^*, \mu_t)$,

$$\begin{aligned} \|\mathbf{x}_{t+1}^*\|_1 &= \mathcal{L}(\mathbf{x}_{t+1}^*, \mathbf{p}_{t+1}^*, \boldsymbol{\nu}_t^*, \mu_t) - \frac{1}{2\mu_t} \left(\|\boldsymbol{\nu}_{t+1}^*\|_2^2 - \|\boldsymbol{\nu}_t^*\|_2^2 \right) \\ &\leq \mathbf{x}^\dagger - \frac{1}{2\mu_t} \left(\|\boldsymbol{\nu}_{t+1}^*\|_2^2 - \|\boldsymbol{\nu}_t^*\|_2^2 \right). \end{aligned} \quad (3.20)$$

Since $\{\boldsymbol{\nu}_t^*\}$ is bounded,

$$\|\mathbf{x}_{t+1}^*\|_1 \leq \mathbf{x}^\dagger + O(\mu_t^{-1}). \quad (3.21)$$

For any accumulation point $\mathbf{x}^*$ of $\mathbf{x}_t^*$, without loss of generality, we have $\mathbf{x}_t^* \rightarrow \mathbf{x}^*$ as $t \rightarrow +\infty$. Hence, $\|\mathbf{x}^*\|_1 \leq \mathbf{x}^\dagger$. In the mean time, $\mathbf{p}_{t+1}^* = \Phi \mathbf{x}_{t+1}^* + \mu_t^{-1}(\boldsymbol{\nu}_{t+1}^* - \boldsymbol{\nu}_t^*) \rightarrow \mathbf{p}^*$ and $\Phi \mathbf{x}^* = \mathbf{p}^*$ result in that $(\mathbf{x}^*, \mathbf{p}^*)$ is an optimal solution to (BP^o) .

Moreover, by $\mathbf{x}_{t+1}^* = \Phi'[\mathbf{p}_{t+1}^* - \mu_t^{-1}(\boldsymbol{\nu}_{t+1}^* - \boldsymbol{\nu}_t^*)]$ and

$$\mathbf{x}^\dagger = \min_{\Phi \mathbf{x} = \mathbf{p}, \Gamma(\mathbf{p}) = \mathbf{y}} \|\mathbf{x}\|_1 = \min_{\mathbf{p}, \Gamma(\mathbf{p}) = \mathbf{y}} \|\Phi' \mathbf{p}\|_1 \leq \|\Phi' \mathbf{p}_{t+1}^*\|_1, \quad (3.22)$$

we have $\|\mathbf{x}_{t+1}^*\|_1 \geq \mathbf{x}^\dagger - O(\mu_t^{-1})$, which establishes the theorem with (3.21). $\blacksquare$

Algorithm 3.1 contains, respectively, an inner and an outer iteration loops. Theorem 3.1 presents only the convergence rate of the outer loop. A natural way to speed up Algorithm 3.1 is to terminate the inner loop without convergence and use the obtained inner-loop solution as the initialization for the next iteration. This is similar to a continuation strategy and can be realized with reasonably set precision and step size μ_t [106]. When the continuation parameter μ_t increases very slowly, in a few iterations, the inner loop can produce a solution with high accuracy. In particular, for the purpose of fast and simple computing, we may update the variables in the inner loop only once before stepping into the outer loop operation. This results in a relaxed version of exact ONE-L1 algorithm (eONE-L1), namely relaxed ONE-L1 algorithm (rONE-L1) outlined in Algorithm 3.2.

Theorem 3.2. *The iterative solution $(\mathbf{x}_t, \mathbf{p}_t)$ of Algorithm 3.2 converges to a feasible solution $(\mathbf{x}^f, \mathbf{p}^f)$ of (BP^o) if $\sum_{t=1}^{+\infty} \mu_t^{-1} < +\infty$. It converges at least exponentially to $(\mathbf{x}^f, \mathbf{p}^f)$ if $\{\mu_t\}$ is an exponentially increasing sequence.*

Algorithm 3.2: Relaxed ONE-L1 Algorithm

Input: Expanded orthonormal matrix Φ and observed sample data $\mathbf{y}$.

1. $\mathbf{x}_0 = \mathbf{0}$; $\mathbf{p}_0 = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}$; $\boldsymbol{\nu}_0 = \mathbf{0}$; $\mu_0 > 0$; $t = 0$.

2. while not converged do

3. $\mathbf{x}_{t+1} = S_{\mu_t^{-1}}(\Phi'(\mathbf{p}_t + \mu_t^{-1}\boldsymbol{\nu}_t))$;

4. $\mathbf{p}_{t+1} = \begin{bmatrix} \mathbf{y} \\ \bar{\Gamma}(\Phi\mathbf{x}_{t+1} - \mu_t^{-1}\boldsymbol{\nu}_t) \end{bmatrix}$;

5. $\boldsymbol{\nu}_{t+1} = \boldsymbol{\nu}_t + \mu_t(\mathbf{p}_{t+1} - \Phi\mathbf{x}_{t+1})$;

6. choose $\mu_{t+1} > \mu_t$;

7. set $t = t + 1$;

8. end while

Output: $(\mathbf{x}_t, \mathbf{p}_t)$.

Proof. We show first that sequences $\{\hat{\boldsymbol{\nu}}_t\}$ and $\{\boldsymbol{\nu}_t\}$ are bounded, where $\hat{\boldsymbol{\nu}}_t = \boldsymbol{\nu}_{t-1} + \mu_{t-1}(\mathbf{p}_{t-1} - \Phi\mathbf{x}_t)$. By the optimality of $\mathbf{x}_{t+1}$ and $\mathbf{p}_{t+1}$ we have

$$0 \in \partial_{\mathbf{x}} \mathcal{L}(\mathbf{x}_{t+1}, \mathbf{p}_t, \boldsymbol{\nu}_t, \mu_t) = \partial \|\mathbf{x}_{t+1}\|_1 - \Phi' \hat{\boldsymbol{\nu}}_{t+1}, \quad (3.23)$$

$$0 = \partial_{\bar{\Gamma}(\mathbf{p})} \mathcal{L}(\mathbf{x}_{t+1}, \mathbf{p}_{t+1}, \boldsymbol{\nu}_t, \mu_t) = \bar{\Gamma}(\boldsymbol{\nu}_{t+1}). \quad (3.24)$$

Hence, $\|\Phi' \hat{\boldsymbol{\nu}}_{t+1}\|_{\infty} \leq 1$ and it follows that $\{\hat{\boldsymbol{\nu}}_t\}$ is bounded. Since $\boldsymbol{\nu}_{t+1} = \hat{\boldsymbol{\nu}}_{t+1} + \mu_t(\mathbf{p}_{t+1} - \mathbf{p}_t)$, we obtain $\bar{\Gamma}(\boldsymbol{\nu}_{t+1}) = \bar{\Gamma}(\hat{\boldsymbol{\nu}}_{t+1})$. This together with $\bar{\Gamma}(\boldsymbol{\nu}_{t+1}) = 0$ results in $\|\boldsymbol{\nu}_{t+1}\|_2 \leq \|\hat{\boldsymbol{\nu}}_{t+1}\|_2$ and the boundedness of $\{\boldsymbol{\nu}_t\}$. By $\mathbf{p}_{t+1} - \mathbf{p}_t = \mu_t^{-1}(\boldsymbol{\nu}_{t+1} - \hat{\boldsymbol{\nu}}_{t+1})$, we have $\|\mathbf{p}_{t+1} - \mathbf{p}_t\|_2 \leq C\mu_t^{-1}$ with C being a constant. Then $\{\mathbf{p}_t\}$ is a Cauchy sequence if $\sum_{t=1}^{+\infty} \mu_t^{-1} < +\infty$, resulting in $\mathbf{p}_t \rightarrow \mathbf{p}^f$ as $t \rightarrow +\infty$. In the mean time, $\mathbf{x}_t \rightarrow \mathbf{x}^f$, $\Phi\mathbf{x}^f = \mathbf{p}^f$. Hence, $(\mathbf{x}^f, \mathbf{p}^f)$ is a feasible solution of (BP^o). Suppose that $\{\mu_t\}$ is an exponentially increasing sequence, i.e., $\mu_{t+1} = r\mu_t$ with $r > 1$. By the boundedness of $\{\boldsymbol{\nu}_t\}$ and $\{\hat{\boldsymbol{\nu}}_t\}$ we have

$$\begin{aligned} \|\mathbf{p}_t - \mathbf{p}^f\|_2 &= \left\| \sum_{i=t}^{+\infty} (\mathbf{p}_i - \mathbf{p}_{i+1}) \right\|_2 \leq \sum_{i=t}^{+\infty} \|\mathbf{p}_i - \mathbf{p}_{i+1}\|_2 \\ &\leq C\mu_t^{-1} \sum_{i=0}^{+\infty} r^{-i} = O(\mu_t^{-1}). \end{aligned} \quad (3.25)$$

Hence, $\{\mathbf{p}_t\}$ converges at least exponentially to $\mathbf{p}^f$ since $\{\mu_t^{-1}\}$ exponentially con-

verges to 0, and the same result holds for $\{\mathbf{x}_t\}$. ■

Remark 3.2. *It is shown in Theorem 3.2 that faster growth of $\{\mu_t\}$ can result in faster convergence of $\{\mathbf{x}_t\}$. Intuitively, the reduced number of iterations for the inner loop problem (SP) may result in some error from the optimal solution $\mathbf{x}_t^*$ of the inner loop. This will likely affect the accuracy of the final solution $\mathbf{x}^f$ for BP. Therefore, the growth speed of $\{\mu_t\}$ provides a tradeoff between the convergence speed of the algorithm and the precision of the final solution, which will be illustrated in Section 3.3 through numerical simulations.*

3.2.3 Relationship Between rONE-L1 and IST

The studies and applications of IST type algorithms have been very active in recent years because of their concise presentations. This subsection considers the relationship between rONE-L1 and IST. Note that $\bar{\Gamma}(\boldsymbol{\nu}_t) = 0$ in Algorithm 3.2 and $\boldsymbol{\Phi}'\boldsymbol{\Phi} = \mathbf{A}'\mathbf{A} + \mathbf{B}'\mathbf{B} = \mathbf{I}$. After some derivations, it can be shown that the rONE-L1 algorithm is equivalent to the following iteration (starting from $\mathbf{x}_t = 0$, as $t \leq 0$, and $\mathbf{z}_t = 0$, as $t < 0$):

$$\begin{aligned}\mathbf{x}_{t+1} &= S_{\lambda_t}(\mathbf{x}_t + \mathbf{A}'\mathbf{z}_t), \\ \mathbf{z}_t &= \mathbf{y} - \mathbf{A}[(1 + \kappa_t)\mathbf{x}_t - \kappa_t\mathbf{x}_{t-1}] + \kappa_t\mathbf{z}_{t-1},\end{aligned}\tag{3.26}$$

where $\lambda_t = \mu_t^{-1}$ and $\kappa_t = \frac{\mu_t - 1}{\mu_t}$. Compared with the general form of IST in (3.8), one more term $\kappa_t\mathbf{z}_{t-1}$ is added when computing the current residual $\mathbf{z}_t$ in rONE-L1. Moreover, a weighted sum $(1 + \kappa_t)\mathbf{x}_t - \kappa_t\mathbf{x}_{t-1}$ is used instead of the current solution $\mathbf{x}_t$. It will be shown later that these two changes essentially improve the sparsity-undersampling tradeoff.

Remark 3.3. *Equations in (3.26) show that the expansion from the partially orthonormal matrix $\mathbf{A}$ to orthonormal $\boldsymbol{\Phi}$ is not at all involved in the actual implementation and computation of rONE-L1. The same claim also holds for eONE-L1 algorithm. Nevertheless, the orthonormal expansion is a key instrumentation in the derivation and analysis of Algorithms 3.1 and 3.2.*

3.2.4 Implementation Details

As noted in Remark 3.3, the expansion from $\mathbf{A}$ to Φ is not involved in the computing of ONE-L1 algorithms. In our implementations, we consider using exponentially increasing μ_t , i.e., fixing $r > 1$ and $\mu_t = r^t \mu_0$. Let $Q_\alpha(\cdot)$ be an α -quantile operator and $\mu_0 = 1/Q_\alpha(|\mathbf{A}'\mathbf{y}|)$, with $|\cdot|$ applying to the vector variable elementwise, μ_0^{-1} being the threshold in the first iteration and $\alpha = 0.99$. In eONE-L1, a large r can speed up the convergence of the outer loop iteration according to Theorem 3.1. However, simulations show that a larger r can result in more iterations in the inner loop. We use $r = 1 + M/N$ as default. In rONE-L1, the value of r provides a tradeoff between the convergence speed of the algorithm and the precision of the final solution. Our recommendation of r to achieve the optimal sparsity-undersampling tradeoff is $r = \min(1 + 0.04M/N, 1.02)$, which will be illustrated in Section 3.3.1.

An iterative algorithm needs a termination criterion. The eONE-L1 algorithm is considered converged if $\frac{\|\mathbf{A}\mathbf{x}_t^* - \mathbf{y}\|_2}{\|\mathbf{b}\|_2} < \tau_1$ with τ_1 being a user-defined tolerance. The inner iteration is considered converged if $\frac{\|\mathbf{x}_t^{j+1} - \mathbf{x}_t^j\|_2}{\|\mathbf{x}_t^j\|_2} < \tau_2$. In our implementation, the default values are $(\tau_1, \tau_2) = (10^{-5}, 10^{-6})$. The rONE-L1 algorithm is considered converged if $\frac{\|\mathbf{A}\mathbf{x}_t - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} < \tau$, with $\tau = 10^{-5}$ as default.

3.3 Numerical Simulations

3.3.1 Sparsity-Undersampling Tradeoff

This subsection considers the sparsity-undersampling tradeoff of rONE-L1 in the case of strictly sparse signals and noise-free measurements. Phase transition is a measure of the sparsity-undersampling tradeoff in this case. Let the sampling ratio be $\delta = M/N$ and the sparsity ratio be $\rho = K/M$, where K is a measure of sparsity of $\mathbf{x}$, and we call that $\mathbf{x}$ is K -sparse if at most K entries of $\mathbf{x}$ are nonzero. As $K, M, N \rightarrow \infty$ with fixed δ and ρ , the behavior of the phase transition of BP is

controlled by (δ, ρ) (see more details in Section 2.2). We denote this theoretical curve by $\rho = \rho^R(\delta)$, which is plotted in Fig 3.1.

We estimate the phase transition of rONE-L1 using a Monte Carlo method as in [33, 34]. Two matrix ensembles are considered, including Gaussian with $N = 1000$ and partial-DCT with $N = 1024$. Here the finite- N phase transition is defined as the value of ρ at which the success probability to recover the original signal is 50%. We consider 33 equispaced values of δ in $\{0.02, 0.05, \dots, 0.98\}$. For each δ , 21 equispaced values of ρ are generated in the interval $[\rho^R(\delta) - 0.1, \rho^R(\delta) + 0.1]$. Then $M = 20$ random problem instances are generated and solved with respect to each combination of (δ, ρ) , where $M = \lceil \delta N \rceil$, $K = \lceil \rho M \rceil$, and nonzero entries of sparse signals are generated from the standard Gaussian distribution. Success is declared if the relative root mean squared error (relative RMSE) $\frac{\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2}{\|\mathbf{x}^o\|_2} < 10^{-4}$, where $\hat{\mathbf{x}}$ is the recovered signal. The number of success among M experiments is recorded. Finally, a generalized linear model is used to estimate the phase transition as in [34].

The experiment result is presented in Fig. 3.1. The observed phase transitions using the recommended value of r strongly agree with the theoretical result of BP. It shows that the rONE-L1 algorithm has the optimal sparsity-undersampling tradeoff in the sense of ℓ_1 minimization.

3.3.2 Comparison with IST

The rONE-L1 algorithm can be considered as a modified version of IST in (3.8). In this subsection we compare the sparsity-undersampling tradeoff and speed of these two algorithms. A similar method is adopted to estimate the phase transition of IST, which is implemented using the same parameter values as rONE-L1. Only nine values of δ in $\{0.1, 0.2, \dots, 0.9\}$ are considered with the partial-DCT matrix ensemble for time consideration. Another choice of $r = 1 + 0.2\delta$ is considered besides the recommended one. Correspondingly, the phase transition of rONE-L1 with $r = 1 + 0.2\delta$ is also estimated.

Figure 3.1: Observed phase transitions of rONE-L1, and comparison with those of IST. Notice that, 1) the observed phase transitions of rONE-L1 with the recommended r strongly agree with the theoretical calculation based on BP; 2) rONE-L1 has significantly enlarged success phases compared with IST.

The observed phase transitions are shown in Fig. 3.1. As a modified version of IST, obviously, rONE-L1 makes a great improvement over IST in the sparsity-undersampling tradeoff. Meanwhile, comparison of the averaged number of iterations of the two algorithms shows that rONE-L1 is also faster than IST. For example, as $\delta = 0.2$ and the recommended r is used, rONE-L1 is about 6 times faster than IST.

3.3.3 Comparison with AMP, FPC-AS and NESTA in Noise-free Case

In this subsection, we report numerical simulation results comparing rONE-L1 with state-of-the-art algorithms, including AMP, FPC-AS and NESTA, in the case of sparse signals and noise-free measurements. Our experiments used FPC-AS v.1.21, NESTA v.1.1, and AMP codes provided by the author. We choose parameter values

for FPC-AS and NESTA such that each method produces a solution with approximately the same precision as that produced by rONE-L1. In our experiments we consider the recovery of exactly sparse signals from partial-DCT measurements. We set $N = 2^{14}$ and $\delta = 0.2$, and an ‘easy’ case where $\rho = 0.1$ and a ‘hard’ case where $\rho = 0.22$ are considered, respectively.¹ Twenty random problems are created and solved for each algorithm with each combination of (δ, ρ) , and the minimum, maximum and averaged relative RMSE, number of calls of $\mathbf{A}$ and $\mathbf{A}'$, and CPU time usages are recorded. All experiments are carried on Matlab v.7.7.0 on a PC with a Windows XP system and a 3GHz CPU. Default parameter values are used in eONE-L1 and rONE-L1.

AMP: terminating if $\frac{\|\mathbf{A}\mathbf{x}_t - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} < 10^{-5}$.

FPC-AS: $\lambda = 2 \times 10^{-6}$ and $gtol = 1 \times 10^{-6}$, where $gtol$ is the termination criterion on the maximum norm of sub-gradient. FPC-AS solves the problem QP.

NESTA: $\lambda = 2 \times 10^{-6}$, $\epsilon = 0$ and the termination criterion $tolvar = 1 \times 10^{-8}$. NESTA solves BPDN using the Nesterov algorithm [104], with continuation.

Our experiment results are presented in Table 3.1. In both ‘easy’ and ‘hard’ cases, rONE-L1 is much faster than eONE-L1. In the ‘easy’ case, the proposed rONE-L1 algorithm takes the most number of calls of $\mathbf{A}$ and $\mathbf{A}'$, except that of eONE-L1, due to a conservative setting of r . But this number of calls (515.4) is very close to that of NESTA (468.9), and furthermore, the CPU time usage of rONE-L1 (2.14 s) is less than that of NESTA (2.70 s) because of its concise implementation. In the ‘hard’ case, rONE-L1 has the second best performance with significantly less CPU time than that of AMP and NESTA. AMP has the second worst CPU time and the worst accuracy as the dynamic threshold in each iteration depends on the mean squared error of the current iterative solution, which cannot be calculated accurately in the implementation.

¹Here ‘easy’ and ‘hard’ refer to the difficulty degree in recovering a sparse signal from a specific number of measurements. The setting $(\delta, \rho) = (0.2, 0.22)$ is close to the phase transition of BP.

Table 3.1: Comparison Results of ONE-L1 Algorithms with State-of-the-art Methods. The column of “# calls $\mathbf{A}$ & $\mathbf{A}'$ ” presents the values of $\text{mean}(\text{minimum}, \text{maximum})$.

ρ	Method	# calls $\mathbf{A}$ & $\mathbf{A}'$	CPU time (s)	Error (10^{-5})
0.1	eONE-L1	1819 (1522,2054)	5.62 (4.67,6.52)	0.42 (0.11,0.94)
	rONE-L1	515.4 (286,954)	2.14 (1.19,3.92)	1.08 (0.53,1.30)
	AMP	222.7 (216,234)	0.80 (0.76,0.86)	1.02 (0.85,1.15)
	FPC-AS	150.2 (135,170)	0.50 (0.44,0.56)	1.13 (1.07,1.23)
	NESTA	468.9 (458,484)	2.70 (2.55,2.98)	1.05 (0.99,1.13)
0.22	eONE-L1	9038 (7270,11194)	28.5 (22.0,35.8)	1.87 (0.46,2.66)
	rONE-L1	722.3 (440,972)	2.61 (1.63,3.93)	1.80 (1.37,3.05)
	AMP	1708 (1150,2252)	6.21 (4.19,9.11)	10.5 (6.96,15.8)
	FPC-AS	589.4 (476,803)	2.10 (1.65,2.80)	1.96 (1.46,3.60)
	NESTA	1084 (890,1244)	6.47 (5.22,7.50)	2.90 (1.62,3.98)

3.3.4 Recovery of 2D Image

This subsection demonstrates the efficiency of rONE-L1 in the general CS setting where the signal of interest is approximately sparse and measurements are contaminated with noise. We seek to reconstruct the Mondrian image of size 256×256 , shown in Fig. 3.2, from its noise-contaminated partial-DCT coefficients. This image presents a challenge as its wavelet expansion is compressible but not exactly sparse. The sampling pattern, which is inspired by magnetic resonance imaging (MRI) and is shown in Fig. 3.2, is adopted since most energy of the image concentrates at low-frequency components after the DCT transform. The measurement vector $\mathbf{y}$ contains $M = 7419$ DCT measurements ($\delta = 0.113$). White Gaussian noise with standard deviation $\sigma = 1$ is then added. We set $\epsilon = \sqrt{M + 2\sqrt{2M}\sigma}$. Haar wavelet with a decomposition level 4 is chosen as the sparsifying transform $\mathcal{W}$. Hence, the problem to be solved is BPDN with $\mathbf{A} = \mathcal{F}_p \mathcal{W}'$, where $\mathcal{F}_p$ is the partial-DCT transform. The reconstructed image is $\hat{\mathbf{H}} = \mathcal{W}' \hat{\mathbf{x}}$ with $\hat{\mathbf{x}}$ being the reconstructed wavelet coefficients and reconstruction error is calculated as $\frac{\|\hat{\mathbf{H}} - \mathbf{H}^o\|_F}{\|\mathbf{H}^o\|_F}$, where $\mathbf{H}^o$ is the original image and $\|\cdot\|_F$ denotes the Frobenius norm. We compare the performance of rONE-L1 with NESTA and FPC-AS.

Remark 3.4. AMP is omitted for its poor performance in this approximately-sparse-signal case. For AMP, the value of the dynamic threshold λ_t and the term $\|x_t\|_0$ in

(3.9) depend on the condition that the signal to reconstruct is strictly sparse.

In such a noisy measurement case, an exact solution for BPDN is not sought after in the rONE-L1 simulation. The computation of the rONE-L1 algorithm is set to terminate if $\frac{\|A\mathbf{x}_t - \mathbf{y}\|_2}{\|\mathbf{y}\|_2} \leq \tau = \frac{\epsilon}{\|b\|_2}$, i.e., rONE-L1 outputs the first $\mathbf{x}_t$ when it becomes a feasible solution of BPDN.

FPC-AS: $\lambda = 1 \times 10^{-3}$, $gtol = 1 \times 10^{-3}$, $gtol_scale_x = 1 \times 10^{-6}$ and the maximum number of iterations for subspace optimization $sub_mxitr = 10$. The parameters are set according to [107, Section 4.4].

NESTA: $\lambda = 1 \times 10^{-4}$, and $tolvar = 1 \times 10^{-6}$. The parameters are tuned to achieve the minimum reconstruction error.

Fig. 3.2 shows the experiment results where rONE-L1, FPC-AS and NESTA produce faithful reconstructions of the original image. The rONE-L1 algorithm produces a reconstruction error (0.0741) lower than that of FPC-AS (0.0809) with comparable computation times (11.1 s and 11.4 s, respectively). While NESTA results in a slightly lower reconstruction error (0.0649), it incurs about twice more computation time (29.4 s).

3.4 Conclusion

In this chapter, we presented novel algorithms for the sparse signal recovery in CS. The proposed rONE-L1 algorithm which is based on the augmented Lagrange multiplier method and heuristic simplification can be considered as a modified IST with an aggressive continuation strategy. Their performance was demonstrated for 1D sparse signals and 2D images and compared with existing methods. In the next chapter, the ONE-L1 algorithms will be extended to deal with CS problems involving complex valued data and applied to discover a new phase transition in the complex domain.

Figure 3.2: An example of 2D image reconstruction from noise-contaminated partial-DCT measurements. Upper left: original Mondrian image; upper right: sampling pattern. The lower three are reconstructed images respectively by rONE-L1 (lower left, error: 0.0741, time: 11.1 s), FPC-AS (lower middle, error: 0.0809, time: 11.4 s) and NESTA (lower right, error: 0.0649, time: 29.4 s).

Chapter 4

Phase Transition in the Complex Domain

As mentioned in Chapter 2, phase transition plays an important role in the CS theory since it describes the exact sparsity-undersampling tradeoff of the ℓ_1 minimization approach under some technical assumptions. The existing phase transition results have been reviewed in Chapter 2 which are focused on the real setting where the sparse signal of interest and the sensing matrix are both real-valued. Since complex data are involved in numerous CS applications, e.g., MRI [36], radar imaging [37] and source localization [11], where both the signal and sensing matrix are complex valued, sparsity-undersampling tradeoff in such a complex setting deserves further studies. For convenience the CS problem in the real and complex settings are referred to as real valued CS (RVCS) and complex valued CS (CVCS) hereafter, respectively. This chapter presents the first result of the phase transition in the complex setting. The significance of our result is twofold: 1) discovering a new phase transition curve for the complex setting so to complement the existing phase transition theory (see Fig. 4.1); 2) providing insights into performance of CS applications involving complex data. Our result is obtained by extending the ONE-L1 algorithms presented in Chapter 3 to the CVCS and applying them to the empirical evaluation of the phase transition of CVCS. Notice that the discovered complex phase transition is

Figure 4.1: Phase transition curves of CS with four different settings. The ones in the real settings are derived in [31–33] (see details in Section 2.2) and the one in the complex setting is discovered due to our work presented in this chapter.

independent of the specific algorithm applied, e.g., ONE-L1 algorithms that we use. However, the computationally efficient ONE-L1 algorithms provide an effective means for empirically exploring the phase transition of CVCS. This chapter is mainly based on [4–6].

4.1 ONE-L1 Algorithms for CVCS

The empirical complex phase transition is evaluated based on the ONE-L1 algorithms presented in Chapter 3. The ONE-L1 algorithms solve the BP problem in RVCS where $\mathbf{x}$, $\mathbf{A}$ and $\mathbf{y}$ are all real valued. Assume that $\mathbf{A}\mathbf{A}' = \mathbf{I}$. The BP is firstly transformed into an equivalent form and then the augmented Lagrange multiplier (ALM) method is applied to solve it. The obtained result leads to the exact ONE-L1 (eONE-L1) algorithm for iterative computation of the optimal solution. While eONE-L1 has an inner iteration loop embedded in the outer loop iteration, its relaxed version, rONE-L1, further speeds up the computation by simplifying the

inner loop iteration into a single update. The rONE-L1 algorithm is numerically optimal in terms of the sparsity-undersampling tradeoff under reasonable parameter settings, i.e., its phase transition result matches that of the BP. It has been shown that rONE-L1 is of iterative thresholding type and is very fast, with appropriate settings of the regulation variable, in comparison with other state-of-the-art algorithms.

An important operation in the ONE-L1 algorithms is the soft thresholding operator $S_\lambda(\mathbf{w})$ defined in (3.4) which solves an ℓ_1 regularized least squares problem (see (3.6)). To extend the ONE-L1 algorithms to CVCS, the complex sensing matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$ is also assumed to satisfy $\mathbf{A}\mathbf{A}' = \mathbf{I}$ with $'$ denoting the conjugate transpose. The soft thresholding operator $S_\lambda(\mathbf{w})$, for a complex vector $\mathbf{w}$, is defined the same as in the real case and is the optimal solution to the same ℓ_1 regularized least squares problem. Let $\Re$ and $\Im$ be operators taking, respectively, the real and imaginary parts of a variable. The augmented Lagrangian function in Chapter 3 is modified, for the CVCS, as

$$\mathcal{L}(\mathbf{x}, \mathbf{d}, \mathbf{y}, \mu) = \|\mathbf{x}\|_1 + \Re \langle \mathbf{d} - \Phi \mathbf{x}, \mathbf{y} \rangle + \frac{\mu}{2} \|\mathbf{d} - \Phi \mathbf{x}\|_2^2,$$

where $\mathbf{y} \in \mathbb{C}^N$ is the Lagrange multiplier vector, $\mu > 0$ and $\langle \mathbf{u}_1, \mathbf{u}_2 \rangle = \mathbf{u}_1' \mathbf{u}_2 \in \mathbb{C}$ is the inner product of $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{C}^N$. The modified augmented Lagrangian function allows a straightforward extension of the derivation and optimization steps in Chapter 3 to the complex BP problem, yielding the optimal solution in the same form as that of the real valued BP. As a result, the ONE-L1 algorithms can be directly applicable to the CVCS problem. Readers are referred to Chapter 3 for detailed algorithm steps and the optimality and convergence analysis.

4.2 Phase Transition of CVCS

4.2.1 Evaluation of Complex Phase Transition via ONE-L1

Section 4.1 has shown that the ONE-L1 algorithms can be extended and applied to the CVCS problem. The eONE-L1 achieves the optimal solution of BP and rONE-L1 is numerically optimal and exponentially converges, which are applied in this subsection to empirically explore the sparsity-undersampling tradeoff of BP. The implementations of ONE-L1 algorithms follow the same procedure as their real versions in Chapter 3. We fix $r > 1$ and let $\mu_{t+1} = r \cdot \mu_t$. The regulation parameter r is set to $r = 1 + \delta$ in eONE-L1 and $r = \min(1 + 0.04\delta, 1.02)$ is chosen in rONE-L1. The success of recovering the original signal is stated if the relative root mean squared error (RRMSE) $\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 / \|\mathbf{x}^o\|_2 < 10^{-4}$, where $\mathbf{x}^o$ and $\hat{\mathbf{x}}$ are the original and recovered signals, respectively. Meanwhile, the failure in solving BP using ONE-L1 is stated if $\|\hat{\mathbf{x}}\|_1 \geq (1 + 10^{-5})\|\mathbf{x}^o\|_1$ and $\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 / \|\mathbf{x}^o\|_2 \geq 10^{-4}$.

Inspired by the estimation of the real phase transition in [33, 34], we first set a complex matrix ensemble, e.g., Gaussian, and dimension N . A grid of (δ, ρ) is generated in the plane $[0, 1] \times [0, 1]$ with equispaced $\delta \in \{0.02, 0.05, \dots, 0.98\}$ and $\rho \in \{\rho^R(\delta) + 0.01(i - 21) : i = 1, 2, \dots, 41\}$ with respect to δ , where $\rho^R(\delta)$ denotes the theoretical real phase transition as shown in Fig. 4.3. For each combination of (δ, ρ) , $M = 20$ random problem instances are generated and solved with $M = \lceil \delta N \rceil$ and $K = \lceil \rho M \rceil$. The number of success S among M instances is recorded. After data acquisition, a generalized linear modal (GLM) with a logistic link is used to estimate the phase transition.

We now explore the sparsity-undersampling tradeoff of BP with partial-Fourier sampling. Four values of signal length N are considered, including 1024, 2048, 4096 and 8192. When $N = 1024$, both eONE-L1 and rONE-L1 are used to estimate the phase transition of BP. The rONE-L1 algorithm is used for other N . Few failures in solving the BP occur when using rONE-L1 (see Table 4.1).

Table 4.1: Number of Failures and CPU Time Consumptions. The “Number of failures” column presents #failures(#problem instances solved).

Ensemble	N	Method	Number of failures			CPU time (hours)
			$\delta = 0.02$	$\delta = 0.05$	Others	
Fourier	1024	exact	0(480)	0(340)	0(8180)	6.26
Fourier	1024	relaxed	14(440)	4(340)	0(8280)	0.327
Fourier	2048	relaxed	1(320)	1(320)	0(7460)	0.594
Fourier	4096	relaxed	4(240)	0(260)	0(6980)	1.01
Fourier	8192	relaxed	1(260)	0(240)	0(6660)	2.15
Gaussian	1000	relaxed	16(420)	0(340)	0(8240)	5.56
Bernoulli	1000	relaxed	9(420)	2(360)	0(8340)	5.60
Ternary	1000	relaxed	9(360)	2(280)	0(8240)	5.65

The observed success rates of our experiments are shown in Fig. 4.2, where the phase-transition performance can be observed. The phase transitions occur at about the same location and larger signal length N leads to sharper phase transition, which is consistent with the behavior of the real phase transition. Fig. 4.3 presents the estimated phase transitions of partial-Fourier sampling. The five observed phase transitions of BP, estimated respectively by eONE-L1 with $N = 1024$ and rONE-L1 with $N = 1024, 2048, 4096$ and 8192 , coincide with each other and are higher than the real phase transition of BP. To sum up the observations of Figs. 4.2 and 4.3 we can state:

Finding 4.1. *For complex signals and the partial-Fourier matrix ensemble with large dimension N , we observe that*

- I. BP exhibits phase transition in the plane of (δ, ρ) , and a larger N can result in a sharper phase transition.*
- II. the complex phase transition of BP is higher than the real phase transition with a considerably enlarged success phase.*

4.2.2 Universality of Phase Transitions of CVCS

The observed universality of phase transitions of BP across different matrix ensembles in the real case has been studied in [34] and the same result is stated in [33] and

Figure 4.2: Observed success rates of partial-Fourier sampling. Four images show results of rONE-L1 with $N = 1024, 2048, 4096$ and 8192 , respectively. Each pixel refers to the success rate S/M at the corresponding coordinate (δ, ρ) . Sampling ratio δ ranges from 0.02 (left) to 0.98 (right) with an interval 0.03; Sparsity ratio ρ ranges from $\rho^R(\delta) - 0.1$ (bottom) to $\rho^R(\delta) + 0.1$ (top) with an interval 0.01 with respect to each δ (only part of the experiment results are shown). The middle line in each image refers to the real phase transition.

Figure 4.3: Observed phase transitions of partial-Fourier sampling. The upper five curves are observed phase transitions estimated by eONE-L1 with $N = 1024$ and rONE-L1 with $N = 1024, 2048, 4096$ and 8192 , respectively. The lower is the theoretical real phase transition of BP.

Chapter 3. In this subsection, we examine whether the same property holds in the complex case. Apart from the partial-Fourier matrix ensemble, three other complex matrix ensembles are considered with signal length $N = 1000$, including Gaussian, Bernoulli and Ternary. All random matrices have i.i.d. real and imaginary parts following the corresponding distributions. Bernoulli refers to equally likely being 0 or 1, and Ternary is equally likely to be -1 , 0 or 1. Notice that a matrix generated from these matrix ensembles may not be partially orthonormal as required by the ONE-L1 algorithms. This problem can be resolved by left multiplying both sides of $\mathbf{Ax} = \mathbf{y}$ with an invertible matrix, e.g., using QR decomposition, so the transformed equation meets the partially orthonormal condition and preserves the same solution space.

Few failures are observed in our experiment when using rONE-L1 to solve BP (see Table 4.1). Fig. 4.4 presents the observed phase transitions of BP with different matrix ensembles. Like the universality of phase transitions in the real case, we have the following finding.

Finding 4.2. *For complex signals and a number of complex matrix ensembles with large dimension N , BP exhibits the same phase transition.*

4.2.3 Connection to Block-Sparse CS

The complex system $\mathbf{Ax} = \mathbf{y}$ can be rewritten into a real valued formulation as $\mathbf{A}_r \mathbf{x}_r = \mathbf{y}_r$ with $\mathbf{A}_r = \begin{bmatrix} \Re \mathbf{A} & -\Im \mathbf{A} \\ \Im \mathbf{A} & \Re \mathbf{A} \end{bmatrix}$, $\mathbf{x}_r = \begin{bmatrix} \Re \mathbf{x} \\ \Im \mathbf{x} \end{bmatrix}$ and $\mathbf{y}_r = \begin{bmatrix} \Re \mathbf{y} \\ \Im \mathbf{y} \end{bmatrix}$. By the K -sparsity of $\mathbf{x}$, $\Re \mathbf{x}$ and $\Im \mathbf{x}$ are jointly K -sparse in the sense that they are both K -sparse and share the same support. On the other hand, after proper permutations of entries of $\mathbf{x}_r$ as well as the corresponding columns of $\mathbf{A}_r$, $\mathbf{A}_r \mathbf{x}_r = \mathbf{y}_r$ can be recast into a block-sparse CS (BSCS) problem [111] with block size 2. The blocked signal of $\mathbf{x}_r$ is K -block-sparse with at most K nonzero block entries and its $\ell_{2,1}$ -norm as defined in [111] is equivalent to the ℓ_1 norm of $\mathbf{x}$ in the CVCS. Thus

Figure 4.4: Observed universality of phase transitions of BP in the complex domain. The signal length $N = 8192$ for partial-Fourier matrix ensemble and $N = 1000$ for the other three matrix ensembles.

the CVCS problem is strongly connected to the BSCS problem via the real valued reformation $\mathbf{A}_r \mathbf{x}_r = \mathbf{y}_r$. Their only difference is that $\mathbf{A}_r$ is subject to a structured constraint whereas entries of the sensing matrix of the BSCS problem are assumed to be independent in general.

A comparison of the phase transition of CVCS with the theoretical phase transition of BSCS with block size 2 in [111] is presented in Fig. 4.5. It is shown that the observed phase transition of CVCS coincides with that of BSCS.

In [111], the theoretical phase transition of BSCS is derived under the condition that the null space of the real valued random sensing matrix is distributed uniformly in the Grassmanian with respect to the Haar measure. It is known that the real Gaussian matrix ensemble satisfies such a condition [112]. We now provide an analysis in the following proposition to show that such a condition is not satisfied in the CVCS problem. It therefore clarifies that the phase transition of the CVCS studied in this chapter is not a special case of the phase transition result of BSCS in [111].

Figure 4.5: Coincidence between our observed complex phase transition and the block-sparse phase transition with block size 2. The indirect phase transition curve refers to minimizing $\|\mathbf{x}_r\|_1$ subject to $\mathbf{A}\mathbf{x} = \mathbf{y}$.

Proposition 4.1. *For any complex random matrix $\mathbf{A}$ associated with the CVCS problem, the null space of $\mathbf{A}_r$ is not distributed uniformly in the Grassmanian with respect to the Haar measure.*

Proof. Let $\mathcal{N}(\mathbf{A}_r)$ denote the null space of $\mathbf{A}_r$. Suppose $\begin{bmatrix} \mathbf{w} \\ \mathbf{v} \end{bmatrix} \in \mathcal{N}(\mathbf{A}_r)$ with $\mathbf{w}, \mathbf{v} \in \mathbb{R}^N$. It is obvious that $\begin{bmatrix} -\mathbf{v} \\ \mathbf{w} \end{bmatrix} \in \mathcal{N}(\mathbf{A}_r)$ and thus it is always possible to choose a basis for $\mathcal{N}(\mathbf{A}_r)$ in the form $\begin{bmatrix} \mathbf{W} & -\mathbf{V} \\ \mathbf{V} & \mathbf{W} \end{bmatrix}$ with $\mathbf{W}, \mathbf{V} \in \mathbb{R}^{N \times (N-M)}$ (columns of $\mathbf{W} + j\mathbf{V} \in \mathbb{C}^{N \times (N-M)}$ in fact compose a basis for the null space of $\mathbf{A}$). As a result, $\mathcal{N}(\mathbf{A}_r)$ is not distributed uniformly in the Grassmanian. ■

Proposition 4.1 shows that the existing phase transition theory of BSCS cannot explain our obtained empirical phase transition of the CVCS. Intuitively, the coincidence of the phase transition of CVCS with that of BSCS may be interpreted

as the universality of phase transitions of BSCS across different matrix ensembles. Fig. 4.5 also shows that the observed phase transition by minimizing $\|\mathbf{x}_r\|_1$, subject to $\mathbf{A}\mathbf{x} = \mathbf{y}$ and without taking into account the block-sparsity, matches the real phase transition of BP. It further explains that the higher successful rate of CVCS is obtained by incorporating the block-sparsity factor into its problem solving.

In the above, it has been shown that the CVCS problem is connected to BSCS with a special structure of the sensing matrix $\mathbf{A}_r$. It is interesting to consider that the sensing matrix in CVCS is further specialized into a block-diagonal matrix, i.e., $\Im \mathbf{A} = \mathbf{0}$ in $\mathbf{A}_r$ or, equivalently, $\mathbf{A} \in \mathbb{R}^{M \times N}$. Such a case refers to the jointly sparse CS or the case of multiple measurement vectors (MMV). In comparison with the standard real CS or the case of single measurement vector (SMV), the recovery performance can be improved in the MMV case under some assumptions on the distribution of entries of $\mathbf{x}$ [113]. On the other hand, the analysis of the null space property in [114] shows that there is little performance improvement in the worst case of signal recovery, such as an $\mathbf{x}$ with identical real and imaginary parts. So, in this special case of jointly sparse CS, the observed universality of phase transitions of BSCS does not hold.

4.3 Rigorous Derivation of the Complex Phase Transition

After we obtained the empirical results of the phase transition presented above, a rigorous analysis has been carried out in [6] with collaborators. In our method, we extend the approximate message passing (AMP) algorithm originally proposed in [33] for RVCS to solve complex valued BP and obtain the complex AMP algorithm (CAMP). We then generalize the state evolution framework which was introduced in [33] for the analysis of AMP, to the complex setting. Using the state evolution, we derive accurate formulae for the phase transition of both BP and CAMP. Our results are theoretically proven for the case of i.i.d. Gaussian sensing matrices and coincide

with our empirical observation in this chapter. In the extreme undersampling regime it holds that $\rho^C(\delta) \sim [\ln(1/2\delta)]^{-1}$, $\delta \rightarrow 0$, where $\rho^C(\delta)$ denotes the complex phase transition curve as a function of δ . As a result, $\lim_{\delta \rightarrow 0} \frac{\rho^C(\delta)}{\rho^R(\delta)} = 2$ by recalling that $\rho^R(\delta) \sim [2\ln(1/\delta)]^{-1}$ as presented in Section 2.2, where $\rho^R(\delta)$ denotes the real phase transition of the BP. This means that, in the extreme undersampling regime the complex BP can recover signals that are two times denser than the signals that are recovered by the real BP.

4.4 Conclusion

In this chapter, a new phase transition of BP for CS is discovered in the complex setting, which is positioned well above the real one with similar properties. Its connection to the existing result of block-sparse CS was studied and its theoretical derivation was discussed. The new phase transition complements the existing phase transition theory in the real setting and provides insights into performance of CS applications involving complex data.

Chapter 5

Robustly Stable Signal Recovery with Structured Matrix Perturbation

The sparsity-undersampling tradeoff of CS has been studied in Chapter 4 in terms of phase transition. Notice that the sensing matrix is assumed to be exactly known there. However, this is not always the case in practice as discussed in Subsection 2.4.2. This chapter is concerned with the latter case where perturbations (or uncertainties) exist in the sensing matrix. In standard formulation of CS, the signal $\mathbf{x}^o \in \mathbb{R}^N$ of interest is acquired by the linear measurements

$$\mathbf{y} = \Phi \mathbf{x}^o + \mathbf{e}, \quad (5.1)$$

where $\Phi \in \mathbb{R}^{M \times N}$ is the sensing matrix (or linear operator), typically with $M \ll N$, $\mathbf{y} \in \mathbb{R}^M$ is the vector of measurements, and $\mathbf{e} \in \mathbb{R}^M$ denotes the vector of measurement noises with bounded energy, i.e., $\|\mathbf{e}\|_2 \leq \epsilon$ for $\epsilon > 0$. Given Φ and ϵ , the task of CS is to recover $\mathbf{x}^o$ from a significantly reduced number of measurements $\mathbf{y}$. Candès *et al.* [17, 58] show that if $\mathbf{x}^o$ is sparse, then it can be stably recovered under mild conditions on Φ with the recovery error being at most proportional to

the measurement noise level ϵ by solving an ℓ_1 minimization problem. Similarly, the largest (in amplitude) entries of a compressible signal can be stably recovered. More details are presented in Subsection 5.1.2. There have been recent active studies on the CS problem where the sensing matrix Φ is unknown or subject to an unknown perturbation. A list of publications include [73] on unknown sensing matrix, [74–76] on general matrix perturbations and [78] considering also structures in perturbation.

This chapter is focused on the perturbed CS problem and mainly based on [7]. We study a structured matrix perturbation with each column of the perturbation matrix being a (unknown) constant times a (known) vector which defines the direction of perturbation. For certain structured matrix perturbation, we provide conditions for guaranteed signal recovery performance. Our analysis shows that robust stability (see definition in Subsection 5.1.1) can be achieved for a sparse signal under similar mild conditions as those for the standard CS problem by solving an ℓ_1 minimization problem incorporated with the perturbation structure. In the special noise free case, the recovery is exact for a sufficiently sparse signal with respect to the perturbation level. A similar result holds for a compressible signal under an additional assumption of small perturbation (depending on the number of largest entries to be recovered). A practical application problem of DOA estimation will be studied in Chapter 8, where the off-grid DOA estimation problem is formulated into our proposed signal recovery problem subject to the structured sensing matrix perturbation, showing the practical relevance of our proposed framework and solution. To verify the obtained results, two algorithms for positive-valued and general signals respectively are proposed to solve the resulting nonconvex ℓ_1 minimization problem. Numerical simulations confirm our robustly stable signal recovery results.

5.1 Preliminary Results

5.1.1 Definitions

In CS, a signal $\mathbf{x}^o \in \mathbb{R}^N$ of length N is called K -sparse if it has at most K nonzero entries, and it is called compressible if its entries obey a power law

$$|x^o|_{(j)} \leq C_q j^{-q}, \quad (5.2)$$

where $|x^o|_{(j)}$ is the j th largest entry (in amplitude) of $\mathbf{x}^o$ ($|x^o|_{(1)} \geq |x^o|_{(2)} \geq \dots \geq |x^o|_{(N)}$), $q > 1$ and C_q is a constant that depends only on q . Let $\mathbf{x}^K$ be a vector that keeps the K largest entries (in amplitude) of $\mathbf{x}^o$ with the rest being zero. If $\mathbf{x}^o$ is compressible, then it can be well approximated by the sparse signal $\mathbf{x}^K$ in the sense that

$$\|\mathbf{x}^o - \mathbf{x}^K\|_2 \leq C'_q K^{-q+1/2} \quad (5.3)$$

where C'_q is a constant.

For the purpose of clarification of expression, we define formally some terminologies for signal recovery used in this chapter, including stability in standard CS, robustness and robust stability in perturbed CS.

Definition 5.1 ([58]). *In standard CS where Φ is known a priori, consider a recovered signal $\hat{\mathbf{x}}$ of $\mathbf{x}^o$ from measurements $\mathbf{y} = \Phi \mathbf{x}^o + \mathbf{e}$ with $\|\mathbf{e}\|_2 \leq \epsilon$. We say that $\hat{\mathbf{x}}$ achieves stable signal recovery if*

$$\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 \leq C_1^{stb} K^{-q+1/2} + C_2^{stb} \epsilon$$

holds for compressible signal $\mathbf{x}^o$ obeying (5.2) and an integer K , or if

$$\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 \leq C_2^{stb} \epsilon$$

holds for K -sparse signal $\mathbf{x}^o$, with nonnegative constants C_1^{stb} , C_2^{stb} .

Definition 5.2. In perturbed CS where $\Phi = \mathbf{A} + \mathbf{E}$ with $\mathbf{A}$ known a priori and $\mathbf{E}$ unknown with $\|\mathbf{E}\|_F \leq \eta$, consider a recovered signal $\hat{\mathbf{x}}$ of $\mathbf{x}^o$ from measurements $\mathbf{y} = \Phi \mathbf{x}^o + \mathbf{e}$ with $\|\mathbf{e}\|_2 \leq \epsilon$. We say that $\hat{\mathbf{x}}$ achieves robust signal recovery if

$$\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 \leq C_1^{rbt} K^{-q+1/2} + C_2^{rbt} \epsilon + C_3^{rbt} \eta$$

holds for compressible signal $\mathbf{x}^o$ obeying (5.2) and an integer K , or if

$$\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 \leq C_2^{rbt} \epsilon + C_3^{rbt} \eta$$

holds for K -sparse signal $\mathbf{x}^o$, with nonnegative constants C_1^{rbt} , C_2^{rbt} and C_3^{rbt} .

Definition 5.3. In perturbed CS where $\Phi = \mathbf{A} + \mathbf{E}$ with $\mathbf{A}$ known a priori and $\mathbf{E}$ unknown with $\|\mathbf{E}\|_F \leq \eta$, consider a recovered signal $\hat{\mathbf{x}}$ of $\mathbf{x}^o$ from measurements $\mathbf{y} = \Phi \mathbf{x}^o + \mathbf{e}$ with $\|\mathbf{e}\|_2 \leq \epsilon$. We say that $\hat{\mathbf{x}}$ achieves robustly stable signal recovery if

$$\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 \leq C_1^{rs}(\eta) K^{-q+1/2} + C_2^{rs}(\eta) \epsilon$$

holds for compressible signal $\mathbf{x}^o$ obeying (5.2) and an integer K , or if

$$\|\hat{\mathbf{x}} - \mathbf{x}^o\|_2 \leq C_2^{rs}(\eta) \epsilon$$

holds for K -sparse signal $\mathbf{x}^o$, with nonnegative constants C_1^{rs} , C_2^{rs} depending on η .

Remark 5.1.

- (1) In the case where $\mathbf{x}^o$ is compressible, the defined stable, robust, or robustly stable signal recovery is in fact for its K largest (in amplitude) entries. The first term $O(K^{-q+1/2})$ in the error bounds above represents, by (5.3), the best approximation error (up to a scale) that can be achieved when we know everything about $\mathbf{x}^o$ and select its K largest entries.
- (2) The Frobenius norm of $\mathbf{E}$, $\|\mathbf{E}\|_F$, can be replaced by any other norm in Definitions 5.2 and 5.3 since the norms are equivalent.

(3) *By robust stability, we mean that the signal recovery is stable for any fixed matrix perturbation level η according to Definition 5.3.*

It should be noted that the stable recovery in standard CS and the robustly stable recovery in perturbed CS are exact in the noise free, sparse signal case while there is no such a guarantee for the robust recovery in perturbed CS.

5.1.2 Stable Signal Recovery of Standard CS

The task of standard CS is to recover the original signal $\mathbf{x}^o$ via an efficient approach given the sensing matrix Φ , acquired sample $\mathbf{y}$ and upper bound ϵ for the measurement noise. This chapter is focused on the ℓ_1 norm minimization approach. The restricted isometry property (RIP) (see Definition 2.2 of Chapter 2) has become a dominant tool to such analysis. Based on the RIP, the following theorem holds.

Theorem 5.1 ([17]). *Assume that $\delta_{2K}(\Phi) < \sqrt{2} - 1$ and $\|\mathbf{e}\|_2 \leq \epsilon$. Then an optimal solution $\mathbf{x}^*$ to the basis pursuit denoising (BPDN) problem*

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - \Phi\mathbf{x}\|_2 \leq \epsilon \quad (5.4)$$

satisfies

$$\|\mathbf{x}^* - \mathbf{x}^o\|_2 \leq C_0^{std} K^{-1/2} \|\mathbf{x}^o - \mathbf{x}^K\|_1 + C_1^{std} \epsilon \quad (5.5)$$

$$\text{where } C_0^{std} = \frac{2[1+(\sqrt{2}-1)\delta_{2K}(\Phi)]}{1-(\sqrt{2}+1)\delta_{2K}(\Phi)}, \quad C_1^{std} = \frac{4\sqrt{1+\delta_{2K}(\Phi)}}{1-(\sqrt{2}+1)\delta_{2K}(\Phi)}.$$

Theorem 5.1 states that a K -sparse signal $\mathbf{x}^o$ ($\mathbf{x}^K = \mathbf{x}^o$) can be stably recovered by solving a computationally efficient convex optimization problem provided $\delta_{2K}(\Phi) < \sqrt{2} - 1$. The same conclusion holds in the case of compressible signal $\mathbf{x}^o$ since

$$K^{-1/2} \|\mathbf{x}^o - \mathbf{x}^K\|_1 \leq C_q'' k^{-q+1/2} \quad (5.6)$$

according to (5.2) and (5.3) with C_q'' being a constant. In the special noise free, K -sparse signal case, such a recovery is exact. The RIP condition in Theorem 5.1 can

be satisfied provided $M \geq O(K \log(N/K))$ with a large probability if the sensing matrix Φ is i.i.d. subgaussian distributed [115]. Note that the RIP condition for the stable signal recovery in standard CS has been relaxed in [116, 117] but it is beyond the scope of this chapter.

5.1.3 Robust Signal Recovery in Perturbed CS

In standard CS, the sensing matrix Φ is assumed to be exactly known. Such an ideal assumption is not always the case in practice. Consider that the true sensing matrix is $\Phi = \mathbf{A} + \mathbf{E}$ where $\mathbf{A} \in \mathbb{R}^{M \times N}$ is the known nominal sensing matrix and $\mathbf{E} \in \mathbb{R}^{M \times N}$ represents the unknown matrix perturbation. Unlike the additive noise term $\mathbf{e}$ in the observation model in (5.1), a multiplicative “noise” $\mathbf{E}\mathbf{x}^o$ is introduced in perturbed CS and is more difficult to analyze since it is correlated with the signal of interest. Denote $\|\mathbf{E}\|_2^{(K)}$ the largest spectral norm taken over all K -column submatrices of $\mathbf{E}$, and similarly define $\|\Phi\|_2^{(K)}$. The following theorem is stated in [74].

Theorem 5.2 ([74]). *Assume that there exist constants $\varepsilon_{\mathbf{E}, \Phi}^{(K)}$, ϵ and $\epsilon_{\mathbf{E}, \mathbf{x}^o}$ such that $\frac{\|\mathbf{E}\|_2^{(K)}}{\|\Phi\|_2^{(K)}} \leq \varepsilon_{\mathbf{E}, \Phi}^{(K)}$, $\|\mathbf{e}\|_2 \leq \epsilon$ and $\|\mathbf{E}\mathbf{x}^o\|_2 \leq \epsilon_{\mathbf{E}, \mathbf{x}^o}$. Assume that $\delta_{2K}(\Phi) < \frac{\sqrt{2}}{(1 + \varepsilon_{\mathbf{E}, \Phi}^{(2K)})^2} - 1$ and $\|\mathbf{x}^o\|_0 \leq K$. Then an optimal solution $\mathbf{x}^*$ to the BPDN problem with the nominal sensing matrix $\mathbf{A}$, denoted by N -BPDN,*

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \epsilon + \epsilon_{\mathbf{E}, \mathbf{x}^o} \quad (5.7)$$

achieves robust signal recovery with

$$\|\mathbf{x}^* - \mathbf{x}^o\|_2 \leq C^{ptb} \epsilon + C^{ptb} \epsilon_{\mathbf{E}, \mathbf{x}^o} \quad (5.8)$$

$$\text{where } C^{ptb} = \frac{4\sqrt{1 + \delta_{2K}(\Phi)}(1 + \varepsilon_{\mathbf{E}, \Phi}^{(2K)})}{1 - (\sqrt{2} + 1) \left[(1 + \delta_{2K}(\Phi))(1 + \varepsilon_{\mathbf{E}, \Phi}^{(2K)})^2 - 1 \right]}.$$

Remark 5.2.

- (1) *The relaxation of the inequality constraint in (5.7) from ϵ to $\epsilon + \epsilon_{\mathbf{E}, \mathbf{x}^o}$ is to ensure that the original signal $\mathbf{x}^o$ is a feasible solution to N-BPDN. Theorem 5.2 is a little different from that in [74], where the multiplicative “noise” $\mathbf{E}\mathbf{x}^o$ is bounded using $\varepsilon_{\mathbf{E}, \Phi}^{(K)}$, $\delta_K(\Phi)$ and $\|\Phi\mathbf{x}^o\|_2$ rather than a constant $\epsilon_{\mathbf{E}, \mathbf{x}^o}$.*
- (2) *Theorem 5.2 is applicable only to the small perturbation case where $\varepsilon_{\mathbf{E}, \Phi}^{(2K)} < \sqrt[4]{2} - 1$ since $\delta_{2K}(\Phi) \geq 0$.*
- (3) *Theorem 5.2 generalizes Theorem 5.1 for the K -sparse signal case. As the perturbation $\mathbf{E} \rightarrow 0$, Theorem 5.2 coincides with Theorem 5.1 for the K -sparse signal case.*

Theorem 5.2 states that, for a small matrix perturbation $\mathbf{E}$, the signal recovery of N-BPDN that is based on the nominal sensing matrix $\mathbf{A}$ is robust to the perturbation with the recovery error growing at most linearly with the perturbation level. Note that, in general, the signal recovery in Theorem 5.2 is unstable according to the definition of stability in this chapter since the recovery error cannot be bounded within a constant (independent of the noise) times the noise level as some perturbation occurs. A result on general signals in [74] is omitted that shows the robust recovery of a compressible signal. The same problem is studied and similar results are reported in [75] based on the greedy algorithm CoSaMP [21].

5.2 SP-CS: CS Subject to Structured Perturbation

5.2.1 Problem Description

In this chapter we consider a structured perturbation in the form $\mathbf{E} = \mathbf{B}\Delta^o$ where $\mathbf{B} \in \mathbb{R}^{M \times N}$ is known *a priori*, $\Delta^o = \text{diag}(\boldsymbol{\beta}^o)$ is a bounded uncertain term with $\boldsymbol{\beta}^o \in [-r, r]^N$ and $r > 0$, i.e., each column of the perturbation is on a known

direction. In addition, we assume that each column of $\mathbf{B}$ has unit norm to avoid the scaling problem between $\mathbf{B}$ and Δ^o (in fact, the D-RIP condition on matrix $[\mathbf{A}, \mathbf{B}]$ in Subsection 5.2.2 implies that columns of both $\mathbf{A}$ and $\mathbf{B}$ have approximately unit norms). As a result, the observation model in (5.1) becomes

$$\mathbf{y} = \Phi \mathbf{x}^o + \mathbf{e}, \quad \Phi = \mathbf{A} + \mathbf{B} \Delta^o \quad (5.9)$$

with $\Delta^o = \text{diag}(\beta^o)$, $\beta^o \in [-r, r]^N$ and $\|\mathbf{e}\|_2 \leq \epsilon$. Given $\mathbf{y}$, $\mathbf{A}$, $\mathbf{B}$, r and ϵ , the task of SP-CS is to recover $\mathbf{x}^o$ and possibly β^o as well.

Remark 5.3.

- (1) *The perturbed CS model in (5.9) is inspired by the off-grid DOA estimation problem that will be studied in Chapter 8. While this chapter is focused on the theoretical analysis, its application to the DOA estimation will be presented in Chapter 8.*
- (2) *Without loss of generality, we assume that $\mathbf{x}$, $\mathbf{y}$, $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{e}$ are all in the real domain unless otherwise stated.*
- (3) *If $x_j^o = 0$ for some $j \in \{1, \dots, N\}$, then β_j^o has no contributions to the observation $\mathbf{y}$ and hence it is impossible to recover β_j^o . As a result, the recovery of β^o in this chapter refers only to the recovery on the support of $\mathbf{x}^o$.*

5.2.2 Main Results

In this chapter, a vector $\mathbf{v}$ is called $2K$ -duplicately (D-) sparse if $\mathbf{v} = [\mathbf{v}_1^T, \mathbf{v}_2^T]^T$ with $\mathbf{v}_1$ and $\mathbf{v}_2$ being of the same dimension and jointly K -sparse (each being K -sparse and sharing the same support). The concept of duplicate (D-) RIP is defined as follows.

Definition 5.4. *Define the $2K$ -duplicate (D-) RIC of a matrix Φ , denoted by $\bar{\delta}_{2K}(\Phi)$, as the smallest number such that*

$$(1 - \bar{\delta}_{2K}(\Phi)) \|\mathbf{v}\|_2^2 \leq \|\Phi \mathbf{v}\|_2^2 \leq (1 + \bar{\delta}_{2K}(\Phi)) \|\mathbf{v}\|_2^2$$

holds for all $2K$ - D -sparse vectors $\mathbf{v}$. Φ is said to satisfy the $2K$ - D -RIP with constant $\bar{\delta}_{2K}(\Phi)$ if $\bar{\delta}_{2K}(\Phi) < 1$.

According to the definition above, D-RIP is a special case of RIP but concerns about only a type of structured sparse signals. With respect to the perturbed observation model in (5.9), let $\Psi = [\mathbf{A}, \mathbf{B}]$. The main results of this chapter are stated in the following theorems. The proof of Theorem 5.3 is provided in Appendix A.1 and proofs of Theorems 5.4 and 5.5 are in Appendix A.2.

Theorem 5.3. *In the noise free case where $\mathbf{e} = \mathbf{0}$, assume that $\|\mathbf{x}^o\|_0 \leq K$ and $\bar{\delta}_{4K}(\Psi) < 1$. Then an optimal solution $(\mathbf{x}^*, \beta^*)$ to the perturbed combinatorial optimization problem*

$$\min_{\mathbf{x} \in \mathbb{R}^N, \beta \in [-r, r]^N} \|\mathbf{x}\|_0, \text{ subject to } \mathbf{y} = (\mathbf{A} + \mathbf{B}\Delta) \mathbf{x} \quad (5.10)$$

with $\Delta = \text{diag}(\beta)$ recovers $\mathbf{x}^o$ and β^o .

Theorem 5.4. *Assume that $\bar{\delta}_{4K}(\Psi) < \left(\sqrt{2(1+r^2)} + 1\right)^{-1}$, $\|\mathbf{x}^o\|_0 \leq K$ and $\|\mathbf{e}\|_2 \leq \epsilon$. Then an optimal solution $(\mathbf{x}^*, \beta^*)$ to the perturbed (P-) BPDN problem*

$$\min_{\mathbf{x} \in \mathbb{R}^N, \beta \in [-r, r]^N} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta) \mathbf{x}\|_2 \leq \epsilon \quad (5.11)$$

achieves robustly stable signal recovery with

$$\|\mathbf{x}^* - \mathbf{x}^o\|_2 \leq C\epsilon, \quad (5.12)$$

$$\|(\beta^* - \beta^o) \odot \mathbf{x}^o\|_2 \leq \mathcal{C}\epsilon, \quad (5.13)$$

where

$$C = \frac{4\sqrt{1 + \bar{\delta}_{4K}(\Psi)}}{1 - \left(\sqrt{2(1+r^2)} + 1\right) \bar{\delta}_{4K}(\Psi)},$$

$$\mathcal{C} = \frac{[2 + \sqrt{1+r^2} \|\Psi\|_2 C]}{\sqrt{1 - \bar{\delta}_{4K}(\Psi)}}.$$

Theorem 5.5. *Assume that $\bar{\delta}_{4K}(\Psi) < \left(\sqrt{2(1+r^2)} + 1\right)^{-1}$ and $\|\mathbf{e}\|_2 \leq \epsilon$. Then an optimal solution $(\mathbf{x}^*, \beta^*)$ to the P-BPDN problem in (5.11) satisfies that*

$$\|\mathbf{x}^* - \mathbf{x}^o\|_2 \leq (C_0 K^{-1/2} + C_1) \|\mathbf{x}^o - \mathbf{x}^K\|_1 + C_2 \epsilon, \quad (5.14)$$

$$\|(\beta^* - \beta^o) \odot \mathbf{x}^K\|_2 \leq (C_0 K^{-1/2} + C_1) \|\mathbf{x}^o - \mathbf{x}^K\|_1 + C_2 \epsilon, \quad (5.15)$$

where

$$C_0 = 2 \left[1 + \left(\sqrt{2(1+r^2)} - 1 \right) \bar{\delta}_{4K}(\Psi) \right] / a,$$

$$C_1 = 2\sqrt{2}r\bar{\delta}_{4K}(\Psi) / a,$$

$$C_0 = \sqrt{1+r^2} \|\Psi\|_2 C_0 / b,$$

$$C_1 = \left[\sqrt{1+r^2} C_1 + 2r \right] \|\Psi\|_2 / b$$

with $a = 1 - \left(\sqrt{2(1+r^2)} + 1\right) \bar{\delta}_{4K}(\Psi)$, $b = \sqrt{1 - \bar{\delta}_{4K}(\Psi)}$ and $C_2 = C$, $C_2 = \mathcal{C}$ with $C, \mathcal{C}$ as defined in Theorem 5.4.

Remark 5.4. *In general, the robustly stable signal recovery cannot be concluded for compressible signals since the error bound in (5.14) may be very large in the case of large perturbation by $C_1 = O(r)$. If the perturbation is small with $r = O(K^{-1/2})$, then the robust stability can be achieved for compressible signals by (5.6) provided that the D-RIP condition in Theorem 5.5 is satisfied.*

5.2.3 Interpretation of the Main Results

Theorem 5.3 states that for a K -sparse signal $\mathbf{x}^o$, it can be recovered by solving a combinatorial optimization problem provided $\bar{\delta}_{4K}(\Psi) < 1$ when the measurements are exact. Meanwhile, β^o can be recovered. Since the combinatorial optimization problem is NP-hard and that its solution is sensitive to measurement noise [118], a more reliable approach, ℓ_1 minimization, is explored in Theorems 5.4 and 5.5.

Theorem 5.4 states the robustly stable recovery of a K -sparse signal $\mathbf{x}^o$ in SP-CS with the recovery error being at most proportional to the noise level. Such robust

stability is obtained by solving an ℓ_1 minimization problem incorporated with the perturbation structure provided that the D-RIC is sufficiently small with respect to the perturbation level in terms of r . Meanwhile, the perturbation parameter β^o can be stably recovered on the support of $\mathbf{x}^o$. As the D-RIP condition is satisfied in Theorem 5.4, the signal recovery error of perturbed CS is constrained by the noise level ϵ , and the influence of the perturbation is limited to the coefficient before ϵ . For example, if $\bar{\delta}_{4K}(\Psi) = 0.2$, then $\|\mathbf{x}^* - \mathbf{x}^o\|_2 \leq 8.48\epsilon, 8.50\epsilon, 11.0\epsilon$ corresponding to $r = 0.01, 0.1, 1$, respectively. In the special noise free case, the recovery is exact. This is similar to that in standard CS but in contrast to the existing robust signal recovery result in Subsection 5.1.3 where the recovery error exists once a matrix perturbation appears. Another interpretation of the D-RIP condition in Theorem 5.4 is that the robustly stable signal recovery requires that $r < \sqrt{\frac{1}{2}(\bar{\delta}_{4K}(\Psi)^{-1} - 1)^2} - 1$ for a fixed matrix Ψ . Using the aforementioned example where $\bar{\delta}_{4K}(\Psi) = 0.2$, the perturbation is required to satisfy $r < \sqrt{7}$. As a result, our robustly stable signal recovery result of SP-CS applies to the case of large perturbation if the D-RIC of Ψ is sufficiently small while the existing result does not as demonstrated in Remark 5.2.

Theorem 5.5 considers general signals and is a generalized form of Theorem 5.4. In comparison with Theorem 5.1 in standard CS, one more term $C_1 \|\mathbf{x}^o - \mathbf{x}^K\|_1$ appears in the upper bound of the recovery error. The robust stability does not hold generally for compressible signals as illustrated in Remark 5.4 while it is true under an additional assumption $r = O(K^{-1/2})$.

The results in this chapter generalize that in standard CS. Without accounting for the symbolic difference between $\delta_{2K}(\Phi)$ and $\bar{\delta}_{4K}(\Psi)$, the conditions in Theorems 5.1 and 5.5 coincide, as well as the upper bounds in (5.5) and (5.14) for the recovery errors, as the perturbation vanishes or equivalently $r \rightarrow 0$. As mentioned before, the RIP condition for guaranteed stable recovery in standard CS has been relaxed. Similar techniques may be adopted to possibly relax the D-RIP condition in SP-CS. While this chapter is focused on the ℓ_1 minimization approach, it is also possible to modify other algorithms in standard CS and apply them to SP-CS to provide similar recovery guarantees.

5.2.4 When is the D-RIP satisfied?

Existing works studying the RIP are mainly focused on random matrices. In standard CS, Φ has the K -RIP with constant δ with a large probability provided that $M \geq C_\delta K \log(N/K)$ and Φ has properly scaled i.i.d. subgaussian distributed entries, where the constant C_δ depending on δ and the distribution [115]. The D-RIP can be considered as a model-based RIP introduced in [119]. Suppose that $\mathbf{A}$, $\mathbf{B}$ are mutually independent and both are i.i.d. subgaussian distributed (the true sensing matrix $\Phi = \mathbf{A} + \mathbf{B}\Delta^o$ is also i.i.d. subgaussian distributed if β^o is independent of $\mathbf{A}$ and $\mathbf{B}$). The model-based RIP is determined by the number of subspaces of the structured sparse signals that are referred to as the D-sparse ones in this chapter. For $\Psi = [\mathbf{A}, \mathbf{B}]$, the number of $2K$ -dimensional subspaces for $2K$ -D-sparse signals is $\binom{N}{K}$. Consequently, Ψ has the $2K$ -D-RIP with constant δ with a large probability also provided that $M \geq C_\delta K \log(N/K)$ by [119, Theorem 1] or [120, Theorem 3.3]. So, in the case of a high dimensional system and $r \rightarrow 0$, the D-RIP condition on Ψ in Theorem 5.4 or 5.5 can be satisfied when the RIP condition on Φ (after proper scaling of its columns) in standard CS is met. It means that the perturbation in SP-CS gradually strengthens the D-RIP condition for robustly stable signal recovery but there exists no gap between SP-CS and standard CS in the case of high dimensional systems.

It is noted that there is another way to stably recover the original signal $\mathbf{x}^o$ in SP-CS. Given the sparse signal case as an example where $\mathbf{x}^o$ is K -sparse. Let $\mathbf{z}^o = \begin{bmatrix} \mathbf{x}^o \\ \beta^o \odot \mathbf{x}^o \end{bmatrix}$, and it is $2K$ -sparse. The observation model can be written as $\mathbf{y} = \Psi \mathbf{z}^o + \mathbf{e}$. Then $\mathbf{z}^o$ and hence, $\mathbf{x}^o$, can be stably recovered from the problem¹

$$\min_{\mathbf{z}} \|\mathbf{z}\|_1, \text{ subject to } \|\mathbf{y} - \Psi \mathbf{z}\|_2 \leq \epsilon \quad (5.16)$$

provided that $\delta_{4K}(\Psi) < \sqrt{2} - 1$ by Theorem 5.1. It looks like that we transformed the perturbation into a signal of interest. Denote TPS-BPDN the problem in (5.16).

¹It is hard to incorporate the knowledge $\beta^o \in [-r, r]^N$ into the problem in (5.16).

In a high dimensional system, the condition $\delta_{4K}(\Psi) < \sqrt{2} - 1$ requires about twice as many as the measurements that makes the D-RIP condition $\bar{\delta}_{4K}(\Psi) < \sqrt{2} - 1$ hold by [119, Theorem 1] corresponding to the D-RIP condition in Theorem 5.4 or 5.5 as $r \rightarrow 0$. As a result, for a considerable range of perturbation level, the D-RIP condition in Theorem 5.4 or 5.5 for P-BPDN is weaker than that for TPS-BPDN since it varies slowly for a moderate perturbation (as an example, $\bar{\delta}_{4K}(\Psi) < 0.414, 0.413, 0.409$ corresponds to $r = 0, 0.1, 0.2$ respectively). Numerical simulations in Subsection 5.4 can verify our conclusion.

5.2.5 Relaxation of the Optimal Solution

In Theorem 5.5 (Theorem 5.4 is a special case), $(\mathbf{x}^*, \beta^*)$ is required to be an optimal solution to P-BPDN. Naturally, we would like to know if the requirement of the optimality is necessary for a “good” recovery in the sense that a good recovery validates the error bounds in (5.14) and (5.15) under the conditions in Theorem 5.5. Generally speaking, the answer is negative since, regarding the optimality of $(\mathbf{x}^*, \beta^*)$, only $\|\mathbf{x}^*\|_1 \leq \|\mathbf{x}^o\|_1$ and the feasibility of $(\mathbf{x}^*, \beta^*)$ are used in the proof of Theorem 5.5 in Appendix A.2. Denote $\mathcal{D}$ the feasible domain of P-BPDN, i.e.,

$$\mathcal{D} = \left\{ (\mathbf{x}, \beta) : \beta \in [-r, r]^N, \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta)\mathbf{x}\|_2 \leq \epsilon \text{ with } \Delta = \text{diag}(\beta) \right\}. \quad (5.17)$$

We have the following corollary.

Corollary 5.1. *Under the assumptions in Theorem 5.5, any $(\mathbf{x}, \beta) \in \mathcal{D}$ that meets $\|\mathbf{x}\|_1 \leq \|\mathbf{x}^o\|_1$ satisfies that*

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}^o\|_2 &\leq (C_0 K^{-1/2} + C_1) \|\mathbf{x}^o - \mathbf{x}^K\|_1 + C_2 \epsilon, \\ \|(\beta - \beta^o) \odot \mathbf{x}^K\|_2 &\leq (C_0 K^{-1/2} + C_1) \|\mathbf{x}^o - \mathbf{x}^K\|_1 + C_2 \epsilon \end{aligned}$$

with $C_j, \mathcal{C}_j$, $j = 0, 1, 2$, as defined in Theorem 5.5.

Figure 5.1: Illustration of Corollary 5.2. The shaded band area refers to the feasible domain of BPDN. The triangular area, the intersection of the feasible domain and the ℓ_1 ball $\{\mathbf{x} : \|\mathbf{x}\|_1 \leq \|\mathbf{x}^o\|_1\}$, is the set of all good recoveries.

Corollary 5.1 generalizes Theorem 5.5 and its proof follows directly from that of Theorem 5.5. It shows that a good recovery in SP-CS is not necessarily an optimal solution to P-BPDN. A similar result holds in standard CS that generalizes Theorem 5.1, and the proof of Theorem 5.1 in [17] applies directly to such case.

Corollary 5.2. *Under the assumptions in Theorem 5.1, any $\mathbf{x}$ that meets $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \epsilon$ and $\|\mathbf{x}\|_1 \leq \|\mathbf{x}^o\|_1$ satisfies that*

$$\|\mathbf{x} - \mathbf{x}^o\|_2 \leq C_0^{std} K^{-1/2} \|\mathbf{x}^o - \mathbf{x}^K\|_1 + C_1^{std} \epsilon \quad (5.18)$$

with C_0^{std}, C_1^{std} as defined in Theorem 5.1.

An illustration of Corollary 5.2 is presented in Fig. 5.1, where the shaded band area refers to the feasible domain of BPDN in (5.4) and all points in the triangular area, the intersection of the feasible domain of BPDN and the ℓ_1 ball $\{\mathbf{x} : \|\mathbf{x}\|_1 \leq \|\mathbf{x}^o\|_1\}$, are good candidates for recovery of $\mathbf{x}^o$. The reason why one seeks for the optimal solution $\mathbf{x}^*$ is to guarantee that the inequality $\|\mathbf{x}\|_1 \leq \|\mathbf{x}^o\|_1$ holds since $\|\mathbf{x}^o\|_1$ is generally unavailable *a priori*. Corollary 5.2 can explain why a satisfactory recovery can be obtained in practice using some algorithm that may not produce an optimal solution to BPDN, e.g., rONE-L1 proposed in Chapter 3. Corollaries 5.1 and 5.2 are useful for checking the effectiveness of an algorithm in the case when the output

cannot be guaranteed to be optimal.² Namely, an algorithm is called effective in solving some ℓ_1 minimization problem if it can produce a feasible solution with its ℓ_1 norm no larger than that of the original signal.

5.3 Algorithms for P-BPDN

5.3.1 Special Case: Positive Signals

This subsection studies a special case where the original signal $\mathbf{x}^o$ is positive-valued (except zero entries). Such a case has been studied in standard CS [32, 33]. By incorporating the positiveness of $\mathbf{x}^o$, P-BPDN is modified into the positive P-BPDN (PP-BPDN) problem

$$\min_{\mathbf{x}, \boldsymbol{\beta}} \mathbf{1}^T \mathbf{x}, \text{ subject to } \begin{cases} \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta}) \mathbf{x}\|_2 \leq \epsilon, \\ \mathbf{x} \succeq \mathbf{0}, \\ r\mathbf{1} \succeq \boldsymbol{\beta} \succeq -r\mathbf{1}, \end{cases}$$

where $\succeq$ operates elementwise for vectors and $\mathbf{0}$, $\mathbf{1}$ are column vectors composed of 0, 1 respectively with proper dimensions. It is noted that the robustly stable signal recovery results in this chapter apply directly to the solution to PP-BPDN in such case. This subsection shows that the nonconvex PP-BPDN problem can be transformed into a convex one and hence its optimal solution can be efficiently obtained. Denote $\mathbf{p} = \boldsymbol{\beta} \odot \mathbf{x}$. A new, convex problem (P_1) is introduced as follows.

$$(P_1) \quad \min_{\mathbf{x}, \mathbf{p}} \mathbf{1}^T \mathbf{x}, \text{ subject to } \begin{cases} \left\| \mathbf{y} - \Psi \begin{bmatrix} \mathbf{x} \\ \mathbf{p} \end{bmatrix} \right\|_2 \leq \epsilon, \\ \mathbf{x} \succeq \mathbf{0}, \\ r\mathbf{x} \succeq \mathbf{p} \succeq -r\mathbf{x}. \end{cases}$$

²It is common when the problem to be solved is nonconvex, such as P-BPDN as discussed in Section 5.3 and ℓ_p ($0 \leq p < 1$) minimization approaches [121–123] in standard CS. In addition, Corollaries 5.1 and 5.2 can be readily extended to the ℓ_p ($0 \leq p < 1$) minimization approaches.

Theorem 5.6. *Problems PP-BPDN and (P_1) are equivalent in the sense that, if $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ is an optimal solution to PP-BPDN, then there exists $\mathbf{p}^* = \boldsymbol{\beta}^* \odot \mathbf{x}^*$ such that $(\mathbf{x}^*, \mathbf{p}^*)$ is an optimal solution to (P_1) , and that, if $(\mathbf{x}^*, \mathbf{p}^*)$ is an optimal solution to (P_1) , then there exists $\boldsymbol{\beta}^*$ with $\beta_j^* = \begin{cases} p_j^*/x_j^*, & \text{if } x_j^* > 0; \\ 0, & \text{otherwise} \end{cases}$ such that $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ is an optimal solution to PP-BPDN.*

Proof. We only prove the first part of Theorem 5.6 using contradiction. The second part follows similarly. Suppose that $(\mathbf{x}^*, \mathbf{p}^*)$ with $\mathbf{p}^* = \boldsymbol{\beta}^* \odot \mathbf{x}^*$ is not an optimal solution to (P_1) . Then there exists $(\mathbf{x}', \mathbf{p}')$ in the feasible domain of (P_1) such that $\|\mathbf{x}'\|_1 < \|\mathbf{x}^*\|_1$. Define $\boldsymbol{\beta}'$ as $\beta'_j = \begin{cases} p'_j/x'_j, & \text{if } x'_j > 0; \\ 0, & \text{otherwise} \end{cases}$. It is easy to show that $(\mathbf{x}', \boldsymbol{\beta}')$ is a feasible solution to PP-BPDN. By $\|\mathbf{x}'\|_1 < \|\mathbf{x}^*\|_1$ we conclude that $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ is not an optimal solution to PP-BPDN, which leads to contradiction. ■

Theorem 5.6 states that an optimal solution to PP-BPDN can be efficiently obtained by solving the convex problem (P_1) .

5.3.2 AA-P-BPDN: Alternating Algorithm for P-BPDN

For general signals, P-BPDN in (5.11) is nonconvex. A simple method is to solve a series of BPDN problems with

$$\mathbf{x}^{(j+1)} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta}^{(j)})\mathbf{x}\|_2 \leq \epsilon, \quad (5.19)$$

$$\boldsymbol{\beta}^{(j+1)} = \arg \min_{\boldsymbol{\beta} \in [-r, r]^N} \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta})\mathbf{x}^{(j+1)}\|_2 \quad (5.20)$$

starting from $\boldsymbol{\beta}^{(0)} = \mathbf{0}$, where the superscript (j) indicates the j th iteration and $\boldsymbol{\Delta}^{(j)} = \text{diag}(\boldsymbol{\beta}^{(j)})$. Denote AA-P-BPDN the alternating algorithm defined by (5.19) and (5.20). To analyze AA-P-BPDN, we first present the following two lemmas.

Lemma 5.1. *For a matrix sequence $\{\boldsymbol{\Phi}^{(j)}\}_{j=1}^\infty$ composed of fat matrices, let $\mathcal{D}^j = \{\mathbf{v} : \|\mathbf{y} - \boldsymbol{\Phi}^{(j)}\mathbf{v}\|_2 \leq \epsilon\}$, $j = 1, 2, \dots$, and $\mathcal{D}^* = \{\mathbf{v} : \|\mathbf{y} - \boldsymbol{\Phi}^*\mathbf{v}\|_2 \leq \epsilon\}$ with $\epsilon > 0$.*

If $\Phi^{(j)} \rightarrow \Phi^*$, as $j \rightarrow +\infty$, then for any $\mathbf{v} \in \mathcal{D}^*$ there exists a sequence $\{\mathbf{v}^{(j)}\}_{j=1}^{\infty}$ with $\mathbf{v}^{(j)} \in \mathcal{D}^{(j)}$, $j = 1, 2, \dots$, such that $\mathbf{v}^{(j)} \rightarrow \mathbf{v}$, as $j \rightarrow +\infty$.

Lemma 5.1 studies the variation of feasible domains $\mathcal{D}^j$, $j = 1, 2, \dots$, of a series of BPDN problems whose sensing matrices $\Phi^{(j)}$, $j = 1, 2, \dots$, converge to Φ^* . It states that the sequence of the feasible domains also converges to $\mathcal{D}^*$ in the sense that for any point in $\mathcal{D}^*$, there exists a sequence of points, each of which belongs to one $\mathcal{D}^j$, that converges to the point. To prove Lemma 5.1, we first show that it holds for any interior point of $\mathcal{D}^*$ by constructing such a sequence. Then we show that it also holds for a boundary point of $\mathcal{D}^*$ by that for any boundary point there exists a sequence of interior points of $\mathcal{D}^*$ that converges to it. The detailed proof is given in Appendix A.3.

Lemma 5.2. *An optimal solution $\mathbf{x}^*$ to the BPDN problem in (5.4) satisfies that $\mathbf{x}^* = \mathbf{0}$, if $\|\mathbf{y}\|_2 \leq \epsilon$, or $\|\mathbf{y} - \Phi\mathbf{x}^*\|_2 = \epsilon$, otherwise.*

Proof. It is trivial for the case where $\|\mathbf{y}\|_2 \leq \epsilon$. Consider the other case where $\|\mathbf{y}\|_2 > \epsilon$. Note first that $\mathbf{x}^* \neq \mathbf{0}$. We use contradiction to show that the equality $\|\mathbf{y} - \Phi\mathbf{x}^*\|_2 = \epsilon$ holds. Suppose that $\|\mathbf{y} - \Phi\mathbf{x}^*\|_2 < \epsilon$. Introduce $f(\theta) = \|\mathbf{y} - \theta\Phi\mathbf{x}^*\|_2$. Then $f(0) > \epsilon$, and $f(1) < \epsilon$. There exists θ_0 , $0 < \theta_0 < 1$, such that $f(\theta_0) = \epsilon$ since $f(\theta)$ is continuous on the interval $[0, 1]$. Hence, $\mathbf{x}' = \theta_0\mathbf{x}^*$ is a feasible solution to BPDN in (5.4). We conclude that $\mathbf{x}^*$ is not optimal by $\|\mathbf{x}'\|_1 = \theta_0 \|\mathbf{x}^*\|_1 < \|\mathbf{x}^*\|_1$, which leads to contradiction. ■

Lemma 5.2 studies the location of an optimal solution to the BPDN problem. It states that the optimal solution locates at the origin if the origin is a feasible solution, or at the boundary of the feasible domain otherwise. This can be easily observed from Fig. 5.1. Based on Lemmas 5.1 and 5.2, we have the following results for AA-P-BPDN.

Theorem 5.7. *Any accumulation point $(\mathbf{x}^*, \beta^*)$ of the sequence $\{(\mathbf{x}^{(j)}, \beta^{(j)})\}_{j=1}^{\infty}$ is a stationary point of AA-P-BPDN in the sense that*

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^*)\mathbf{x}\|_2 \leq \epsilon, \quad (5.21)$$

$$\boldsymbol{\beta}^* = \arg \min_{\boldsymbol{\beta} \in [-r, r]^N} \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta}) \mathbf{x}^*\|_2 \quad (5.22)$$

with $\boldsymbol{\Delta}^* = \text{diag}(\boldsymbol{\beta}^*)$.

Theorem 5.8. *An optimal solution $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ to P-BPDN in (5.11) is a stationary point of AA-P-BPDN.*

Theorem 5.7 studies the property of the solution $(\mathbf{x}^{(j)}, \boldsymbol{\beta}^{(j)})$ produced by AA-P-BPDN. It shows that $(\mathbf{x}^{(j)}, \boldsymbol{\beta}^{(j)})$ is arbitrarily close to a stationary point of AA-P-BPDN as the iteration index j is large enough.³ Hence, the output of AA-P-BPDN can be considered as a stationary point provided that an appropriate termination criterion is set. Theorem 5.8 tells that an optimal solution to P-BPDN is a stationary point of AA-P-BPDN. So, it is possible for AA-P-BPDN to produce an optimal solution to P-BPDN. The proofs of Theorems 5.7 and 5.8 are provided in Appendix A.4 and Appendix A.5 respectively.

Remark 5.5. *The following formulation of P-BPDN can possibly provide an efficient approach to an optimal solution to P-BPDN. Let $\mathbf{x} = \mathbf{x}_+ - \mathbf{x}_-$ where $\mathbf{x}_+ \succeq \mathbf{0}$, $\mathbf{x}_- \succeq \mathbf{0}$ and $\mathbf{x}_+ \odot \mathbf{x}_- = \mathbf{0}$. Then we have $|\mathbf{x}| = \mathbf{x}_+ + \mathbf{x}_-$ where $|\cdot|$ applies elementwise. Denote $\mathbf{p} = \boldsymbol{\beta} \odot \mathbf{x}$. A convex problem can be cast as follows:*

$$\begin{aligned} & \min_{\mathbf{x}_+, \mathbf{x}_-, \mathbf{p}} \mathbf{1}^T (\mathbf{x}_+ + \mathbf{x}_-), \\ & \text{subject to } \begin{cases} \left\| \mathbf{y} - \begin{bmatrix} \mathbf{A} & -\mathbf{A} & \mathbf{B} \end{bmatrix} \begin{bmatrix} \mathbf{x}_+ \\ \mathbf{x}_- \\ \mathbf{p} \end{bmatrix} \right\|_2 \leq \epsilon, \\ \mathbf{x}_+ \succeq \mathbf{0}, \\ \mathbf{x}_- \succeq \mathbf{0}, \\ r(\mathbf{x}_+ + \mathbf{x}_-) \succeq \mathbf{p} \succeq -r(\mathbf{x}_+ + \mathbf{x}_-). \end{cases} \end{aligned} \quad (5.23)$$

³It is shown in the proof of Theorem 5.7 in Appendix A.4 that the sequence $\{(\mathbf{x}^{(j)}, \boldsymbol{\beta}^{(j)})\}_{j=1}^\infty$ is bounded. And it can be shown, for example, using contradiction, that for a bounded sequence $\{a_j\}_{j=1}^\infty$, there exists an accumulation point of $\{a_j\}_{j=1}^\infty$ such that a_j is arbitrarily close to it as j is large enough.

The above convex problem can be considered as a convex relaxation of P-BPDN since it can be shown (like that in Theorem 5.6) that an optimal solution to P-BPDN can be obtained based on an optimal solution to the problem in (5.23) incorporated with an additional nonconvex constraint $\mathbf{x}_+ \odot \mathbf{x}_- = \mathbf{0}$. An interesting phenomenon has been observed through numerical simulations that an optimal solution to the problem in (5.23) still satisfies the constraint $\mathbf{x}_+ \odot \mathbf{x}_- = \mathbf{0}$. Based on such an observation, an efficient approach to P-BPDN is to firstly solve (5.23), and then check whether its solution, denoted by $(\mathbf{x}_+^*, \mathbf{x}_-^*, \mathbf{p}^*)$, satisfies $\mathbf{x}_+ \odot \mathbf{x}_- = \mathbf{0}$. If it does, then $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ is an optimal solution to P-BPDN where $\mathbf{x}^* = \mathbf{x}_+^* - \mathbf{x}_-^*$ and $\beta_j^* = \begin{cases} p_j^*/x_j^*, & \text{if } x_j^* \neq 0; \\ 0, & \text{otherwise.} \end{cases}$ Otherwise, we may turn to AA-P-BPDN again. But we note that it is still an open problem whether an optimal solution to (5.23) always satisfies the constraint $\mathbf{x}_+ \odot \mathbf{x}_- = \mathbf{0}$. In addition, the convex relaxation in (5.23) does not apply to the complex signal case as in the DOA estimation studied in Chapter 8.

5.3.3 Effectiveness of AA-P-BPDN

As reported in the last subsection, it is possible for AA-P-BPDN to produce an optimal solution to P-BPDN. But it is not easy to check the optimality of the output of AA-P-BPDN because of the nonconvexity of P-BPDN. Instead, we study the effectiveness of AA-P-BPDN in solving P-BPDN in this subsection with the concept of effectiveness as defined in Subsection 5.2.5. By Corollary 5.1, a good signal recovery $\hat{\mathbf{x}}$ of $\mathbf{x}^o$ is not necessarily an optimal solution. It requires only that $(\hat{\mathbf{x}}, \hat{\boldsymbol{\beta}})$, where $\hat{\boldsymbol{\beta}}$ denotes the recovery of $\boldsymbol{\beta}^o$, be a feasible solution to P-BPDN and that $\|\hat{\mathbf{x}}\|_1 \leq \|\mathbf{x}^o\|_1$ holds. As shown in the proof of Theorem 5.7 in Appendix A.4, that $(\mathbf{x}^{(j)}, \boldsymbol{\beta}^{(j)})$ for any $j \geq 1$ is a feasible solution to P-BPDN and that the sequence $\{\|\mathbf{x}^{(j)}\|_1\}_{j=1}^\infty$ is monotone decreasing and converges. So, the effectiveness of AA-P-BPDN in solving P-BPDN can be assessed via numerical simulations by checking whether $\|\mathbf{x}^{AA}\|_1 \leq \|\mathbf{x}^o\|_1$ holds with $\mathbf{x}^{AA}$ denoting the output of AA-P-BPDN. The effectiveness of AA-P-BPDN is verified in Section 5.4 via numerical simulations,

where we observe that the inequality $\|\mathbf{x}^{AA}\|_1 \leq \|\mathbf{x}^o\|_1$ holds in all experiments (over 3700 trials).

5.4 Numerical Simulations

This section demonstrates the robustly stable signal recovery results of SP-CS, as well as the effectiveness of AA-P-BPDN in solving P-BPDN in (5.11), via numerical simulations. AA-P-BPDN is implemented in Matlab with problems in (5.19) and (5.20) being solved using CVX [124]. AA-P-BPDN is terminated as $\frac{|\|\mathbf{x}^{(j)}\|_1 - \|\mathbf{x}^{(j-1)}\|_1|}{\|\mathbf{x}^{(j-1)}\|_1} \leq 1 \times 10^{-6}$ or the maximum number of iterations, set to 200, is reached. PP-BPDN is also implemented in Matlab and solved by CVX.

We first consider general signals. The sparse signal case is mainly studied. The variation of the signal recovery error is studied with respect to the noise level, perturbation level and number of measurements respectively. Besides AA-P-BPDN for P-BPDN in SP-CS, performances of three other approaches are also studied. The first one assumes that the perturbation is known *a priori* and recovers the original signal $\mathbf{x}^o$ by solving, namely, the oracle (O-) BPDN problem

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\mathbf{\Delta}^o)\mathbf{x}\|_2 \leq \epsilon.$$

The O-BPDN approach produces the best recovery result of SP-CS within the scope of ℓ_1 minimization of CS since it exploits the exact perturbation (oracle information). The second one corresponds to the robust signal recovery of perturbed CS as described in Subsection 5.1.3 and solves N-BPDN in (5.7) where $\epsilon_{\mathbf{E}, \mathbf{x}^o} = \|\mathbf{B}\mathbf{\Delta}^o\mathbf{x}^o\|_2$ is used though it is not available in practice. The last one refers to the other approach to SP-CS that seeks for the signal recovery by solving TPS-BPDN in (5.16) as discussed in Subsection 5.2.5.

The first experiment studies the signal recovery error with respect to the noise level. We set the signal length $N = 200$, sample size $M = 80$, sparsity level $K = 10$

Figure 5.2: Signal and perturbation recovery errors with respect to the noise level ϵ with parameter settings $(N, M, K, r) = (200, 80, 10, 0.1)$. Both signal and β^o recovery errors of AA-P-BPDN for P-BPDN in SP-CS are proportional to ϵ .

and perturbation parameter $r = 0.1$. The noise level ϵ varies from 0.05 to 2 with interval 0.05. For each combination of (N, M, K, r, ϵ) , the signal recovery error, as well as β^o recovery error (on the support of $\mathbf{x}^o$), is averaged over $R = 50$ trials. In each trial, matrices $\mathbf{A}$ and $\mathbf{B}$ are generated from Gaussian distribution and each column of them has zero mean and unit norm after proper scaling. The sparse signal $\mathbf{x}^o$ is composed of unit spikes with random signs and locations. Entries of β^o are uniformly distributed in $[-r, r]$. The noise $\mathbf{e}$ is zero mean Gaussian distributed and then scaled such that $\|\mathbf{e}\|_2 = \epsilon$. Using the same data, the four approaches, including O-BPDN, N-BPDN, TPS-BPDN and AA-P-BPDN for P-BPDN, are used to recover $\mathbf{x}^o$ respectively in each trial. The simulation results are shown in Fig. 5.2. It can be seen that both signal and β^o recovery errors of AA-P-BPDN for P-BPDN in SP-CS are proportional to the noise, which is consistent with our robustly stable signal recovery result. The error of N-BPDN grows linearly with the noise but a large error still exhibits in the noise free case. Except the ideal case of O-BPDN, our proposed P-BPDN has the smallest error.

The second experiment studies the effect of the structured perturbation. Experiment settings are the same as those in the first experiment except that we set $(N, M, K, \epsilon) = (200, 80, 10, 0.5)$ and vary $r \in \{0.05, 0.1, \dots, 1\}$. Fig. 5.3 presents our simulation results. A nearly constant error is obtained using O-BPDN in standard CS since the perturbation is assumed to be known in O-BPDN. The error of

Figure 5.3: Signal and perturbation recovery errors with respect to the perturbation level in terms of r with parameter settings $(N, M, K, \epsilon) = (200, 80, 10, 0.5)$. The error of AA-P-BPDN for P-BPDN in SP-CS slowly increases with the perturbation level and is quite close to that of the ideal case of O-BPDN for a moderate perturbation.

AA-P-BPDN for P-BPDN in SP-CS slowly increases with the perturbation level and is quite close to that of O-BPDN for a moderate perturbation. Such a behavior is consistent with our analysis. Besides, it can be observed that the error of N-BPDN grows linearly with the perturbation level. Again, our proposed P-BPDN has the smallest error except O-BPDN.

The third experiment studies the variation of the recovery error with the number of measurements. We set $(N, K, r, \epsilon) = (200, 10, 0.1, 0.2)$ and vary $M \in \{30, 35, \dots, 100\}$. Simulation results are presented in Fig. 5.4. Signal recovery errors of all four approaches decrease as the number of measurements increases. Again, it is observed that O-BPDN of the ideal case achieves the best result followed by our proposed P-BPDN. For example, to obtain the signal recovery error of 0.05, about 55 measurements are needed for O-BPDN while the numbers are, respectively, 65 for AA-P-BPDN and 95 for TPS-BPDN. It is impossible for N-BPDN to achieve such a small error in our observation because of the existence of the perturbation.

We next consider a compressible signal that is generated by taking a fixed sequence $\{2.8843 \cdot j^{-1.5}\}_{j=1}^N$ with $N = 200$, randomly permuting it, and multiplying by a random sign sequence (the coefficient 2.8843 is chosen such that the compressible signal has the same ℓ_2 norm as the sparse signals in the previous experiments). It is sought to be recovered from $M = 70$ noisy measurements with $\epsilon = 0.2$ and $r = 0.1$.

Figure 5.4: Signal and perturbation recovery errors with respect to the number of measurements with parameter settings $(N, K, r, \epsilon) = (200, 10, 0.1, 0.2)$. AA-P-BPDN for P-BPDN in SP-CS has the best performance except the ideal case of O-BPDN.

Give experiment results in one instance as an example. The signal recovery error of AA-P-BPDN for P-BPDN in SP-CS is about 0.239, while errors of O-BPDN, N-BPDN and TPS-BPDN are about 0.234, 0.361 and 0.314 respectively.

For the special positive signal case, an optimal solution to PP-BPDN can be efficiently obtained. An experiment result is shown in Fig. 5.5, where a sparse signal of length $N = 200$, composed of $K = 10$ positive unit spikes, is exactly recovered from $M = 50$ noise free measurements with $r = 0.1$ by solving (P_1) .

5.5 Conclusion

This chapter studied the CS problem in the presence of measurement noise and a structured matrix perturbation. A concept named as robust stability for signal recovery was introduced. It was shown that the robust stability can be achieved for a sparse signal by solving an ℓ_1 minimization problem P-BPDN under mild conditions. In the presence of measurement noise, the recovery error is at most proportional to the noise level and the recovery is exact in the special noise free case. A general result for compressible signals was also reported. An alternating algorithm named as AA-P-BPDN was proposed to solve the nonconvex P-BPDN problem, and numerical simulations were carried out, verifying our theoretical analysis. A

Figure 5.5: Exact recovery of a positive sparse signal from noise-free measurements with $(M, N, K, r, \epsilon) = (200, 50, 10, 0.1, 0)$. PP-BPDN is solved by solving (P_1) . β^o and its recovery are shown only on the support of $\mathbf{x}^o$. Black circles: original signal and β^o ; red stars: recoveries.

practical application of the proposed perturbed CS framework and solution to the DOA estimation will be provided in Chapter 8.

Chapter 6

Bayesian Compressed Sensing with New Sparsity-Inducing Prior

Besides the ℓ_1 optimization methods that have been studied in the previous chapters, sparse Bayesian learning (SBL) [39] is another popular approach to the sparse signal recovery in CS, see a list of publications [38, 41–51, 54]. This chapter and the following two are built upon this approach. In this chapter, we introduce a new sparsity-inducing prior named as Gaussian shifted-truncated-gamma (G-STG) prior that generalizes the Laplace prior in [46] and the Gaussian-gamma prior in [54]. The extended flexibility of the new prior promotes its capability of modeling sparse signals. In fact, we show that the Gaussian-gamma prior cannot work in the main algorithm of this chapter. From the perspective of MAP estimation, the G-STG prior corresponds to a nonconvex objective function in optimization methods in general. For signal recovery we propose an iterative algorithm based on an evidence procedure and a fast greedy algorithm inspired by the algorithm in [57]. We show that similar theoretical guarantees shown in [40] for the basic SBL also hold for the new SBL method. Specifically, every local optimum of the SBL cost function is achieved at a sparse solution and the global optimum is achieved at the maximally sparse solution. Moreover, the proposed algorithm produces a sparser solution than existing

SBL methods. We provide simulation results with 1D synthetic signals and 2D images to verify our analysis and compare the proposed method with state-of-the-art ones.

6.1 G-STG Prior

6.1.1 Mathematical Formulation

We introduce the hierarchical Gaussian shifted-truncated-gamma (G-STG) prior for a sparse signal $\mathbf{x} \in \mathbb{R}^N$ as follows:

$$p(\mathbf{x}|\boldsymbol{\alpha}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{\Lambda}), \quad (6.1)$$

$$p(\boldsymbol{\alpha}; \tau, \epsilon, \eta) = \prod_{i=1}^N p(\alpha_i; \tau, \epsilon, \eta) = \prod_{i=1}^N \Gamma_{\tau}(\alpha_i|\epsilon, \eta), \quad (6.2)$$

where $\mathbf{\Lambda} = \text{diag}(\boldsymbol{\alpha})$, $\boldsymbol{\alpha} \in \mathbb{R}^N$, $p(\alpha_i; \tau, \epsilon, \eta)$ is a shifted-truncated-gamma (STG) distribution for $\alpha_i \geq 0$ with

$$\Gamma_{\tau}(\alpha_i|\epsilon, \eta) = \frac{\eta^{\epsilon}}{\Gamma_{\eta\tau}(\epsilon)} (\alpha_i + \tau)^{\epsilon-1} \exp\{-\eta(\alpha_i + \tau)\}, \quad (6.3)$$

$\epsilon \geq 0$ is the shape parameter, $\eta \geq 0$ is the rate parameter, $\tau \geq 0$ is the threshold parameter and $\Gamma_{\tau}(\epsilon) = \int_{\tau}^{\infty} t^{\epsilon-1} e^{-t} dt$ denotes an incomplete gamma function. The first layer of the prior is a commonly used Gaussian prior that leads to convenient computations as shown later. In the second layer α_i , $i = 1, \dots, N$, are assumed to be independent, and further $\alpha_i + \tau$, $i = 1, \dots, N$, are i.i.d. truncated gamma distributed (that is why we say that $p(\alpha_i; \tau, \epsilon, \eta)$ is an STG distribution). By $p(\alpha_i) \propto (\alpha_i + \tau)^{\epsilon-1} \exp\{-\eta\alpha_i\}$, $i = 1, \dots, N$, it is obvious that α_i is favored to be zero in the second layer if $\epsilon \leq 1$, resulting in that x_i is favored to be zero. Thus the hierarchical prior is a sparsity-inducing prior. In general, there is no explicit

expression for the marginal distribution

$$\begin{aligned} p(\mathbf{x}; \tau, \epsilon, \eta) &= \prod_{i=1}^N p(x_i; \tau, \epsilon, \eta) \\ &= \prod_{i=1}^N \int_0^\infty \mathcal{N}(x_i|0, \alpha_i) \Gamma_\tau(\alpha_i|\epsilon, \eta) d\alpha_i. \end{aligned} \quad (6.4)$$

In the following we study some special cases and show that the G-STG prior generalizes those in [46, 54].

1) $\epsilon = 1$: In this case, the second layer is reduced to an exponential prior independent of τ since $\Gamma_\tau(\alpha_i|1, \eta) = \eta e^{-\eta\alpha_i}$. Then, the G-STG prior coincides with the Laplace prior in [46] with $p(x_i; \tau, 1, \eta) = \sqrt{\eta/2} \exp(-\sqrt{2\eta}|x_i|)$.

2) $\tau = 0$: The second layer becomes a gamma prior for each α_i . As a result, the proposed prior becomes the Gaussian-gamma prior in [54] and $p(x_i) = \frac{2^{3/4-\epsilon/2}}{\sqrt{\pi}\Gamma(\epsilon)} \eta^{\frac{2\epsilon+1}{4}} |x_i|^{\epsilon-\frac{1}{2}} \mathcal{K}_{\epsilon-\frac{1}{2}}(\sqrt{2\eta}|x_i|)$ where $\mathcal{K}_\nu(\cdot)$ is the modified Bessel function of the second kind and order $\nu \in \mathbb{R}$. In addition, we have that $p(0) = +\infty$ if $\epsilon \leq \frac{1}{2}$ and $p(0) < +\infty$ if $\epsilon > \frac{1}{2}$. Though the G-STG prior generalizes the Gaussian-gamma prior, it should be noted that the main algorithm based on the G-STG prior proposed in this chapter works differently from that in [54].

3) $\tau \rightarrow +\infty$: By l'Hospital's rule it can be shown that $\lim_{\tau \rightarrow +\infty} \Gamma_\tau(\alpha_i|\epsilon, \eta) = \eta e^{-\eta\alpha_i}$, i.e., the prior for α_i in the second layer approaches an exponential prior in such a case. Consequently, the proposed G-STG prior coincides with the Laplace prior as in *Case 1*.

To visualize the variation of the G-STG prior with respect to the two parameters τ and ϵ , we plot the PDF $p(x_i; \tau, \epsilon, \eta)$ in Fig. 6.1 with $\eta = 1$. Fig. 6.1(a) is for the case of $\epsilon = 0.1$ and varying τ . Obviously, the G-STG prior is a sparsity-inducing prior with the PDFs highly peaked at the origin, especially when $\tau \rightarrow 0$. The main difference between the cases $\tau = 0$ and $\tau = 1 \times 10^{-8}$ is near the origin where $p(x_i)$ approaches infinity for $\tau = 0$ while it is always finite for $\tau > 0$. As τ gets larger, less density concentrates near the origin and the resulting prior gets closer to the Laplace

prior that corresponds to $\tau = +\infty$. Fig. 6.1(b) is for the case of $\tau = 1 \times 10^{-8}$ and varying ϵ . It is shown that the G-STG prior gets less sparsity-inducing as ϵ gets larger, and that it ceases to promote sparsity as $\epsilon > 1$. From the perspective of MAP estimation, the G-STG prior corresponds to a nonconvex optimization method as $\epsilon < 1$ since the term $-\log p(\mathbf{x})$ is nonconvex in such a case.

6.1.2 Intuitive Interpretation and Threshold Parameter Setting

Strictly speaking, a continuous prior is not suitable for sparse signals since any vector generated from a continuous prior is only approximately sparse (the probability of a zero-valued entry is zero). In the following, we provide an intuitive explanation about why the G-STG prior works for sparse signals by setting the threshold parameter τ according to the noise level.

We consider a Gaussian ensemble sensing matrix $\mathbf{A}$ (the entries of $\mathbf{A}$ are i.i.d. Gaussian $\mathcal{N}(0, M^{-1})$ where the variance is set to M^{-1} to make columns of $\mathbf{A}$ have expected unit norm). Then, consider a compressible signal $\mathbf{z}$ and the observation model $\mathbf{y} = \mathbf{A}\mathbf{z}$. In the Bayesian framework, we assume that $\mathbf{z}$ is distributed according to some sparsity-inducing prior. Here we adopt the Gaussian-gamma prior in [54], i.e., we assume that $p(\mathbf{z}|\boldsymbol{\beta}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \boldsymbol{\beta})$ and $p(\boldsymbol{\beta}; \epsilon, \eta) = \prod_{i=1}^N \Gamma(\beta_i|\epsilon, \eta)$ where $\boldsymbol{\beta} = \text{diag}(\boldsymbol{\beta})$ and $\Gamma(\beta_i|\epsilon, \eta)$ refers to $\Gamma_\tau(\beta_i|\epsilon, \eta)$ with $\tau = 0$. However, theoretical results [23, 58] state that only significant entries of $\mathbf{z}$ can be recovered while insignificant ones play as noises. So we write $\mathbf{z}$ into $\mathbf{z} = \mathbf{x} + \mathbf{w}$ where $\mathbf{x}$ denotes the significant, “recoverable” component and $\mathbf{w}$ refers to the insignificant, “unrecoverable” part. Then the observation model becomes $\mathbf{y} = \mathbf{A}(\mathbf{x} + \mathbf{w}) = \mathbf{A}\mathbf{x} + \mathbf{e}$ with $\mathbf{e} = \mathbf{A}\mathbf{w}$. By the structure of $\mathbf{A}$, $\mathbf{e}$ is a zero-mean AWGN with the noise variance $\sigma^2 = M^{-1} \|\mathbf{w}\|_2^2$. We see that only the power of $\mathbf{w}$ is reflected in the noise. That is, for any vector $\tilde{\mathbf{w}}$ satisfying $\|\tilde{\mathbf{w}}\|_2^2 = \|\mathbf{w}\|_2^2$, $\mathbf{A}\tilde{\mathbf{w}}$ and $\mathbf{e}$ are identically distributed. So we may replace $\mathbf{w}$ by $\tilde{\mathbf{w}}$ in the observation model (i.e., $\mathbf{x} + \tilde{\mathbf{w}}$ and $\mathbf{z}$ are indistinguishable for CS approaches). Then we may model $\tilde{\mathbf{w}}$ as an i.i.d. zero-mean Gaussian

(a) $\epsilon = 0.1$ (b) $\tau = 1 \times 10^{-8}$

Figure 6.1: PDFs of the G-STG prior in the case of (a) $\epsilon = 0.1$ and varying τ and (b) $\tau = 1 \times 10^{-8}$ and varying ϵ with $\eta = 1$.

vector with variance $\tau = N^{-1} \|\mathbf{w}\|_2^2 = (M/N) \sigma^2$. In addition, $\tilde{\mathbf{w}}$ is independent of $\mathbf{x}$. So, under the assumption of a Gaussian ensemble matrix $\mathbf{A}$, a compressible signal $\mathbf{z}$ is equivalent to the sum of its significant part $\mathbf{x}$ plus a white Gaussian noise $\tilde{\mathbf{w}}$. Applying the Gaussian-gamma prior for $\mathbf{z}$ to $\mathbf{x} + \tilde{\mathbf{w}}$, we obtain that $\beta_i = \alpha_i + \tau$, $i = 1, \dots, N$, where $\boldsymbol{\alpha}$ is as defined in (6.1). Then we get $p(\alpha_i) = p(\beta_i - \tau | \beta_i \geq \tau)$ as a conditional distribution, resulting in that $p(\boldsymbol{\alpha})$ is in the exact form of (6.2).

In this chapter, we mainly consider the observation model $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$ where $\mathbf{x}$ is a sparse signal and $\mathbf{e}$ is a zero-mean AWGN with known variance σ^2 . The same sparsity-inducing prior for $\mathbf{x}$ can be obtained by a reverse procedure and the details are omitted. So we can set the threshold parameter $\tau = (M/N) \sigma^2$ in the G-STG prior. Though this setting is only based on intuition without rigorous analysis, it indeed leads to good performance as to be reported via simulations in Section 6.4, where it is also observed that this setting applies to other sensing matrix ensembles besides the Gaussian one.

6.2 Sparse Bayesian Learning for Signal Recovery

6.2.1 Bayesian Formulation

In SBL, the signal of interest and noise are modeled as random variables. Under the common assumption of zero-mean AWGNs, i.e., $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, where σ^2 is the noise variance, the PDF of the compressive measurements is

$$p(\mathbf{y}|\mathbf{x}; \sigma^2) = \mathcal{N}(\mathbf{y}|\mathbf{A}\mathbf{x}, \sigma^2 \mathbf{I}). \quad (6.5)$$

In this chapter we assume that the noise variance σ^2 is known *a priori*. This assumption has been widely made in the CS literature, e.g., [58, 125]. Moreover, it is shown in [41] that to estimate σ^2 jointly with the signal recovery process (e.g., in [39]) can lead to very inaccurate estimate.

The G-STG prior introduced in Section 6.1 is adopted as the sparsity-inducing prior for the sparse signal $\mathbf{x}$. The hyperparameters τ and ϵ are chosen manually according to the reasoning in Section 6.1 and Subsection 6.3.2. Numerical simulations will be provided in Section 6.4 to illustrate their performance. To estimate η from the measurements, we assume a gamma hyperprior for η : $p(\eta; c, d) = \Gamma(\eta|c, d)$, where we let $c, d \rightarrow 0$ to obtain a uniform hyperprior (over a logarithmic scale). So the joint PDF of the observation model is $p(\mathbf{y}, \mathbf{x}, \boldsymbol{\alpha}, \eta; \sigma^2, \tau, \epsilon, c, d) = p(\mathbf{y}|\mathbf{x}; \sigma^2) p(\mathbf{x}|\boldsymbol{\alpha}) p(\boldsymbol{\alpha}|\eta; \tau, \epsilon) p(\eta; c, d)$, where $\mathbf{y}$ is the observation, $\mathbf{x}$ is the unknown signal of interest, $\boldsymbol{\alpha}$ and η are unknown parameters, and σ^2 , τ , ϵ , c and d are fixed. The task is to estimate $\mathbf{x}$.

6.2.2 Bayesian Inference

Note that the exact Bayesian inference is intractable since $p(\mathbf{x}|\mathbf{y})$ is computationally intractable. Some approximations have to be made. Following from [39], we decompose the posterior $p(\mathbf{x}, \boldsymbol{\alpha}, \eta|\mathbf{y})$ into two terms as

$$p(\mathbf{x}, \boldsymbol{\alpha}, \eta|\mathbf{y}) = p(\mathbf{x}|\mathbf{y}, \boldsymbol{\alpha}) p(\boldsymbol{\alpha}, \eta|\mathbf{y}). \quad (6.6)$$

The first term $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\alpha})$ is the posterior for $\mathbf{x}$ given the hyperparameter $\boldsymbol{\alpha}$, which will be later shown to be a Gaussian PDF. Then, we compute the most probable estimates of $\boldsymbol{\alpha}$ and η , say $\boldsymbol{\alpha}_{MP}$ and η_{MP} , that maximize the second term $p(\boldsymbol{\alpha}, \eta|\mathbf{y})$. We use $\boldsymbol{\alpha}_{MP}$ to obtain the posterior for $\mathbf{x}$. From the perspective of signal estimation, this is equivalent to requiring

$$p(\mathbf{x}|\mathbf{y}) = \int p(\mathbf{x}, \boldsymbol{\alpha}, \eta|\mathbf{y}) d\boldsymbol{\alpha} d\eta \approx p(\mathbf{x}|\mathbf{y}, \boldsymbol{\alpha}_{MP}), \quad (6.7)$$

where $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\alpha}_{MP})$ refers to $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\alpha})$ at $\boldsymbol{\alpha}_{MP}$. Similar procedures have been adopted in [40, 46]. We will provide theoretical evidence to show that this approach leads to desirable properties in Section 6.2.3. Simulation results presented in Section 6.4 also suggest that the signal recovery based on this approximation is very effective.

Since $p(\mathbf{y}|\mathbf{x})$ and $p(\mathbf{x}|\boldsymbol{\alpha})$ are both Gaussian, it is easy to show that the posterior for $\mathbf{x}$ and the marginal distribution for $\mathbf{y}$ are both Gaussian with $p(\mathbf{x}|\mathbf{y}, \boldsymbol{\alpha}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$, $p(\mathbf{y}|\boldsymbol{\alpha}) = \mathcal{N}(\mathbf{y}|\mathbf{0}, \mathbf{C})$, where

$$\boldsymbol{\mu} = \sigma^{-2} \boldsymbol{\Sigma} \mathbf{A}^T \mathbf{y}, \quad (6.8)$$

$$\boldsymbol{\Sigma} = (\sigma^{-2} \mathbf{A}^T \mathbf{A} + \boldsymbol{\Lambda}^{-1})^{-1}, \quad (6.9)$$

$$\mathbf{C} = \sigma^2 \mathbf{I} + \mathbf{A} \boldsymbol{\Lambda} \mathbf{A}^T. \quad (6.10)$$

The maximization of $p(\boldsymbol{\alpha}, \eta|\mathbf{y})$ is equivalent to that of $p(\mathbf{y}, \boldsymbol{\alpha}, \eta) = p(\mathbf{y}|\boldsymbol{\alpha}) p(\boldsymbol{\alpha}|\eta) p(\eta)$ by the relation $p(\boldsymbol{\alpha}, \eta|\mathbf{y}) = p(\mathbf{y}, \boldsymbol{\alpha}, \eta) / p(\mathbf{y})$. We consider $\log \eta$ as the hidden variable instead of η since the uniform hyperprior is assumed over a logarithmic scale. By $p(\log \eta) = \eta p(\eta)$ we see that the hyperprior for η leads to a noninformative prior by setting $c = d = 0$. So the log-likelihood function is

$$\begin{aligned} \mathcal{L}(\boldsymbol{\alpha}, \log \eta) &= \log p(\mathbf{y}, \boldsymbol{\alpha}, \log \eta) \\ &= -\frac{1}{2} \log |\mathbf{C}| - \frac{1}{2} \mathbf{y}^T \mathbf{C}^{-1} \mathbf{y} \\ &\quad + (\epsilon - 1) \sum_{i=1}^N \log(\alpha_i + \tau) - \eta \sum_{i=1}^N (\alpha_i + \tau) \\ &\quad + (N\epsilon + c) \log \eta - N \log \Gamma_{\eta\tau}(\epsilon) - d\eta + C_1, \end{aligned} \quad (6.11)$$

where C_1 is a constant. The maximizer of $\mathcal{L}$ will be analyzed in Subsection 6.2.3. In the following we provide an iterative procedure to maximize $\mathcal{L}$ by recognizing the identities $\log |\mathbf{C}| = \log |\boldsymbol{\Lambda}| + M \log \sigma^2 - \log |\boldsymbol{\Sigma}|$ and $\mathbf{y}^T \mathbf{C}^{-1} \mathbf{y} = \sigma^{-2} \|\mathbf{y} - \mathbf{A}\boldsymbol{\mu}\|_2^2 + \boldsymbol{\mu}^T \boldsymbol{\Lambda}^{-1} \boldsymbol{\mu}$.

6.2.2.1 Update of α

For α_i , $i = 1, \dots, N$, we have

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \alpha_i} &= -\frac{1}{2\alpha_i} + \frac{E\{x_i^2\}}{2\alpha_i^2} + \frac{\epsilon - 1}{\alpha_i + \tau} - \eta \\ &= -\frac{f(\alpha_i)}{2\alpha_i^2(\alpha_i + \tau)}, \end{aligned} \quad (6.12)$$

where $f(t) = 2\eta t^3 + (3 - 2\epsilon + 2\eta\tau)t^2 + (\tau - E\{x_i^2\})t - \tau E\{x_i^2\}$ for $t \in \mathbb{R}$ is a cubic function and $E\{x_i^2\} = \mu_i^2 + \Sigma_{ii}$, where Σ_{ii} is the i th diagonal entry of Σ . We need the following lemma.

Lemma 6.1. *For a cubic function $g(t) = \lambda_1 t^3 + \lambda_2 t^2 + \lambda_3 t + \lambda_4$, if $\lambda_1, \lambda_2 > 0$ and $\lambda_4 < 0$, then $g(t) = 0$ has a unique root on $(0, +\infty)$.*

Proof. By $g(0) = \lambda_4 < 0$ and $\lim_{t \rightarrow +\infty} g(t) = +\infty$ there exists at least one root in $(0, +\infty)$. We show that this root is unique using contradiction. Suppose there exists more than one positive root. Then there must exist three positive roots and that $g(t)$ has two positive stationary points. That is, the two solutions of $\frac{dg(t)}{dt} = 3\lambda_1 t^2 + 2\lambda_2 t + \lambda_3 = 0$ are both positive, resulting in that $\lambda_2 < 0$ (contradiction). ■

By Lemma 6.1, it is easy to show that the maximum of $\mathcal{L}$ is achieved at the unique positive root, say $\alpha_i^* > 0$, of $f(\alpha_i) = 0$. We note that explicit expressions are available for the roots of a cubic function and hence α_i , $i = 1, \dots, N$, can be efficiently updated.

6.2.2.2 Update of η

In general, there is no explicit expression for updating η . Since the first and second derivatives of $\mathcal{L}$ with respect to $\log \eta$ can be easily computed, $\mathcal{L}$ can be efficiently maximized with respect to $\log \eta$ using numerical methods, e.g., gradient ascending method or Newton's method. In addition, the computational complexity hardly depends on the CS problem dimension.

Remark 6.1. *The computation of the incomplete gamma function $\Gamma_{\eta\tau}(\epsilon)$ is involved in the update of η . This term can be efficiently computed using functions provided in Matlab if ϵ is properly bounded away from zero. But a numerical integration is needed if $\epsilon = 0$. In Section 6.4, we observe through simulations that the update of η may take considerably long time in the case of $\epsilon = 0$. But such differences are negligible in the case of a high-dimensional CS problem since the computation of η hardly depends on the problem dimension unlike other computations, such as the update of α .*

As a result, an iterative algorithm can be implemented to obtain α_{MP} and η_{MP} by iteratively updating Σ in (6.9), μ in (6.8), α and η . It is easy to show that this algorithm can be implemented using an EM algorithm [126]. So the likelihood $\mathcal{L}$ increases monotonically at each iteration and the algorithm is guaranteed to converge. After convergence, the signal $\mathbf{x}$ is estimated using its posterior mean μ . One shortcoming of the iterative algorithm is that at each iteration a high-dimensional matrix inversion has to be calculated for updating Σ though this computation can be possibly alleviated using the Woodbury matrix identity.

6.2.3 Analysis of Global and Local Maxima

We analyze the global and local maxima of the likelihood $\mathcal{L}$ in (6.11) in this subsection. Our analysis is rooted in [40] and shows that the theoretical results on the basic SBL in [40] can be extended to our case with necessary modifications. In the following, we assume that $c = d = 0$ and $\eta > 0$ is fixed (it is a similar case if η is chosen to maximize $\mathcal{L}$ as well). Thus we may write $\mathcal{L}$ (with respect to α) as

$$\begin{aligned} \mathcal{L}(\alpha) = & -\frac{1}{2} \log |\mathbf{C}| - \frac{1}{2} \mathbf{y}^T \mathbf{C}^{-1} \mathbf{y} \\ & + (\epsilon - 1) \sum_{i=1}^N \log(\alpha_i + \tau) - \eta \sum_{i=1}^N \alpha_i + C_4, \end{aligned} \quad (6.13)$$

where C_4 is a constant independent of α . We first consider the global maxima in the noise free case.

Theorem 6.1. *Let $\tau \geq 0$, $0 \leq \epsilon \leq 1$, $\mathbf{e} = 0$ and $\sigma^2 = 0$. Assume that $\mathbf{x}^0$ satisfying $\|\mathbf{x}^0\|_0 < M < N$ is the maximally sparse solution to the linear equation $\mathbf{y} = \mathbf{A}\mathbf{x}$. Then, there exists some $\boldsymbol{\alpha}^0$ with $\|\boldsymbol{\alpha}^0\|_2 < +\infty$ and $\|\boldsymbol{\alpha}^0\|_0 = \|\mathbf{x}^0\|_0$ such that at $\boldsymbol{\alpha}^0$, $\mathcal{L}$ is globally maximized and the corresponding posterior mean is $\boldsymbol{\mu} = \mathbf{x}^0$.*

Proof. Note first that $\lim_{\sigma^2 \rightarrow 0} \boldsymbol{\mu} = \boldsymbol{\Lambda}^{\frac{1}{2}} \left(\mathbf{A} \boldsymbol{\Lambda}^{\frac{1}{2}} \right)^\dagger \mathbf{y} = \boldsymbol{\Lambda}^{\frac{1}{2}} \left(\mathbf{A} \boldsymbol{\Lambda}^{\frac{1}{2}} \right)^\dagger \mathbf{A} \mathbf{x}^0$ by (6.8). Since $\|\mathbf{x}^0\|_0 < M$ there must exist $\boldsymbol{\alpha}^0 \in \mathbb{R}^N$, $\|\boldsymbol{\alpha}^0\|_2 < +\infty$ and $\|\boldsymbol{\alpha}^0\|_0 = \|\mathbf{x}^0\|_0$, such that $\boldsymbol{\mu} = \mathbf{x}^0$ at $\boldsymbol{\alpha} = \boldsymbol{\alpha}^0$. In fact, $\boldsymbol{\alpha}^0$ can be any vector satisfying that for $j = 1, \dots, N$, $\alpha_j^0 > 0$ if $x_j^0 \neq 0$, and $\alpha_j^0 = 0$ otherwise. In the following we show that the global maximum of $\mathcal{L}$ is achieved at $\boldsymbol{\alpha} = \boldsymbol{\alpha}^0$. Let $C_5 = \|\boldsymbol{\alpha}^0\|_\infty < \infty$ and denote $C_6 > 0$ the minimum nonzero α_i , $i = 1, \dots, N$. Following from the proof of [40, Theorem 1], we have $|\mathbf{C}| = |\sigma^2 \mathbf{I} + \mathbf{A} \boldsymbol{\Lambda} \mathbf{A}^T| \rightarrow 0$ and $\mathbf{y}^T \mathbf{C}^{-1} \mathbf{y} \leq \|\mathbf{x}^0\|_2^2 / C_6$ if $\boldsymbol{\alpha} \rightarrow \boldsymbol{\alpha}^0$. In addition, we have $\sum_{i=1}^N \log(\alpha_i + \tau) \leq N \log(C_5 + \tau)$ and $\sum_{i=1}^N \alpha_i \leq NC_5$. So we have $\mathcal{L} = +\infty$ at $\boldsymbol{\alpha} = \boldsymbol{\alpha}^0$, which completes the proof. ■

Theorem 6.1 shows that the global maximum of the objective function is achieved at the maximally sparse solution. Thus the proposed approach has no structural errors and the remaining question of whether the algorithm can produce this solution is a convergence issue. Further, we have the following result which is similar to that in [40, Theorem 2].

Theorem 6.2. *Let $\tau \geq 0$ and $0 \leq \epsilon \leq 1$. Every local maximum of $\mathcal{L}$ is achieved at a sparse $\boldsymbol{\alpha}$ with $\|\boldsymbol{\alpha}\|_0 \leq M$ that leads to a sparse posterior mean $\boldsymbol{\mu}$ with $\|\boldsymbol{\mu}\|_0 \leq M$, regardless of the existence of noise.*

Proof. Note that $\|\boldsymbol{\mu}\|_0 \leq \|\boldsymbol{\alpha}\|_0$ since $\mu_i \rightarrow 0$ as $\alpha_i \rightarrow 0$, $i = 1, \dots, N$. So we need only to show that every local maximum of $\mathcal{L}$ is achieved at a sparse $\boldsymbol{\alpha}$ with $\|\boldsymbol{\alpha}\|_0 \leq M$. Let $q(\boldsymbol{\alpha}) = -\frac{1}{2} \log |\mathbf{C}| + (\epsilon - 1) \sum_{i=1}^N \log(\alpha_i + \tau) - \eta \sum_{i=1}^N \alpha_i$, which is convex with respect to $\boldsymbol{\alpha} \in \mathbb{R}_+^N$ if $\epsilon \leq 1$. Suppose that $\boldsymbol{\alpha}^*$ is a local maximum point of $\mathcal{L}$. We may construct a closed, bounded convex polytope $\mathcal{P} \subset \mathbb{R}_+^N$ following from [40] (we omit the details) such that $\boldsymbol{\alpha}^* \in \mathcal{P}$ and if $\boldsymbol{\alpha} \in \mathcal{P}$, then the second

term of $\mathcal{L}$, $-\frac{1}{2}\mathbf{y}^T\mathbf{C}^{-1}\mathbf{y}$, equals a constant $C_7 = -\frac{1}{2}\mathbf{y}^T(\sigma^2\mathbf{I} + \mathbf{A}\text{diag}(\boldsymbol{\alpha}^*)\mathbf{A}^T)^{-1}\mathbf{y}$. In addition, all extreme points of $\mathcal{P}$ are sparse with support size no more than M . As a result, $\boldsymbol{\alpha}^*$ is a local maximum point of $q(\boldsymbol{\alpha})$ with respect to $\boldsymbol{\alpha} \in \mathcal{P}$. By the convexity of $q(\boldsymbol{\alpha})$, $\boldsymbol{\alpha}^*$ must be an extreme point of $\mathcal{P}$ with $\|\boldsymbol{\alpha}^*\|_0 \leq M$. ■

Theorem 6.2 states that all local maxima of $\mathcal{L}$ are achieved at sparse solutions. Since we can locally maximize $\mathcal{L}$ efficiently in practice, based on Theorem 6.2, we introduce a fast algorithm in Section 6.3 that searches for a sparse solution that locally maximizes $\mathcal{L}$.

6.3 Fast Algorithm

6.3.1 Fast Greedy Algorithm

Based on Theorem 6.2 in Section 6.2.3, the following algorithm aims to find a sparse solution that locally maximizes $\mathcal{L}$. We consider the contribution of a single basis vector $\mathbf{a}_j$ (the j th column of $\mathbf{A}$), $j = 1, \dots, N$, and determine whether it should be included in the model (or in the active set) for the maximization of $\mathcal{L}$. Denote $\mathbf{C}_{-j} = \sigma^2\mathbf{I} + \sum_{i \neq j} \alpha_i \mathbf{a}_i \mathbf{a}_i^T = \mathbf{C} - \alpha_j \mathbf{a}_j \mathbf{a}_j^T$ that is independent of the j th basis vector $\mathbf{a}_j$. Then we have $|\mathbf{C}| = |\mathbf{C}_{-j}| |1 + \alpha_j \mathbf{a}_j^T \mathbf{C}_{-j}^{-1} \mathbf{a}_j|$ and $\mathbf{C}^{-1} = \mathbf{C}_{-j}^{-1} - (\alpha_j^{-1} + \mathbf{a}_j^T \mathbf{C}_{-j}^{-1} \mathbf{a}_j)^{-1} \mathbf{C}_{-j}^{-1} \mathbf{a}_j \mathbf{a}_j^T \mathbf{C}_{-j}^{-1}$. Using the two identities above, we rewrite the log-likelihood function (with respect to $\boldsymbol{\alpha}$) into

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\alpha}) &= -\frac{1}{2} \log |\mathbf{C}_{-j}| - \frac{1}{2} \mathbf{y}^T \mathbf{C}_{-j}^{-1} \mathbf{y} \\
&\quad + (\epsilon - 1) \sum_{i=1}^N \log(\alpha_i + \tau) - \eta \sum_{i=1}^N \alpha_i \\
&\quad - \frac{1}{2} \log |1 + \alpha_j \mathbf{a}_j^T \mathbf{C}_{-j}^{-1} \mathbf{a}_j| + \frac{q_j^2}{2(\alpha_j^{-1} + s_j)} + C_2 \\
&= \mathcal{L}(\boldsymbol{\alpha}_{-j}) + \ell(\alpha_j) + C_3,
\end{aligned} \tag{6.14}$$

where $\mathcal{L}(\boldsymbol{\alpha}_{-j})$ denotes $\mathcal{L}(\boldsymbol{\alpha})$ after removing the contribution of the j th basis vector, $\ell(\alpha_j) = -\frac{1}{2} \log |1 + \alpha_j s_j| + \frac{q_j^2}{2(\alpha_j^{-1} + s_j)} + (\epsilon - 1) \log \left(\frac{\alpha_j}{\tau} + 1 \right) - \eta \alpha_j$ with $s_j = \mathbf{a}_j^T \mathbf{C}_{-j}^{-1} \mathbf{a}_j$ and $q_j = \mathbf{a}_j^T \mathbf{C}_{-j}^{-1} \mathbf{y}$, and C_2, C_3 are constants independent of $\boldsymbol{\alpha}$. Note that $\ell(\alpha_i)$ has been modified by a constant such that $\ell(0) = 0$. In the following we compute the maximum point, say α_j^* , of $\ell(\alpha_j)$ on $[0, +\infty)$. If $\alpha_j^* > 0$, then the basis vector $\mathbf{a}_j$ should be preserved in the active set since it gives positive contribution to the likelihood. Otherwise, it should be removed from the active set (by setting $\alpha_j = 0$). We consider only the general case $\tau > 0$, $0 \leq \epsilon \leq 1$ and $\eta > 0$ since it is simple for other cases. First we have

$$\begin{aligned} \frac{d\ell(\alpha_j)}{d\alpha_j} &= -\frac{s_j}{2(1 + \alpha_j s_j)} + \frac{q_j^2}{2(1 + \alpha_j s_j)^2} + \frac{\epsilon - 1}{\alpha_j + \tau} - \eta \\ &= -\frac{h(\alpha_j)}{2(\alpha_j + \tau)(1 + \alpha_j s_j)^2}, \end{aligned} \quad (6.15)$$

where $h(\alpha_j) = c_1 \alpha_j^3 + c_2 \alpha_j^2 + c_3 \alpha_j + c_4$ is a cubic function of α_j with

$$c_1 = 2\eta s_j^2, \quad (6.16)$$

$$c_2 = (3 - 2\epsilon) s_j^2 + 4\eta s_j + 2\eta \tau s_j^2, \quad (6.17)$$

$$c_3 = (5 - 4\epsilon) s_j + 2\eta - q_j^2 + \tau (4\eta s_j + s_j^2), \quad (6.18)$$

$$c_4 = 2 - 2\epsilon + \tau (s_j + 2\eta - q_j^2). \quad (6.19)$$

To compute the maximum point of $\ell(\alpha_j)$, we need the following result.

Lemma 6.2 ([127]). *For a cubic function $g(t) = \lambda_1 t^3 + \lambda_2 t^2 + \lambda_3 t + \lambda_4$, let the discriminant $\Delta = 18\lambda_1 \lambda_2 \lambda_3 \lambda_4 - 4\lambda_2^3 \lambda_4 + \lambda_2^2 \lambda_3^2 - 4\lambda_1 \lambda_3^3 - 27\lambda_1^2 \lambda_4^2$. If $\Delta > 0$, then $g(t) = 0$ has three distinct real roots. If $\Delta = 0$, then the equation has three real roots including a multiple root. If $\Delta < 0$, then the equation has one real root and two complex conjugate roots.*

Note that $s_j > 0$ for $j = 1, \dots, N$ since $\mathbf{C}_{-j}$ and $\mathbf{C}_{-j}^{-1}$ are positive definite, and then $c_1, c_2 > 0$. We divide our discussions into three scenarios.

6.3.1.1 $c_4 < 0$

This case is the same as the update of α in Subsection 6.2.2. By Lemma 6.1, $h(\alpha_j) = 0$ has a unique solution α_j^* on $(0, +\infty)$. Then it is easy to show that $\ell(\alpha_j)$ increases monotonically on $(0, \alpha_j^*]$ and decreases monotonically on $(\alpha_j^*, +\infty)$. Thus $\ell(\alpha_j)$ obtains the maximum at $\alpha_j^* > 0$.

6.3.1.2 $c_4 \geq 0, c_3 < 0$ and $\Delta > 0$

Let Δ denote the determinant of $h(\alpha_j)$. We see that $h(\alpha_j) = 0$ has three real distinct roots by Lemma 6.2 since $\Delta > 0$. The two stationary points of $h(\alpha_j)$ lie on different sides of the y -axis since $\frac{dh(\alpha_j)}{d\alpha_j} = 3c_1\alpha_j^2 + 2c_2\alpha_j + c_3$ with $c_3 < 0$. Thus the three roots include one negative root and two distinct positive roots since $h(0) = c_4 \geq 0$. Denote $\alpha'_j > 0$ the largest root. We see that $\ell(\alpha_j)$ decreases from $\alpha_j = 0$ to some point, then increases until $\alpha_j = \alpha'_j$ and decreases again. As a result, $\ell(\alpha_j)$ obtains the maximum at 0 or α'_j . So we have

- if $\ell(\alpha'_j) > 0$, then the maximum point is $\alpha_j^* = \alpha'_j > 0$;
- if $\ell(\alpha'_j) \leq 0$, then $\alpha_j^* = 0$;

6.3.1.3 $c_4 \geq 0$, and $c_3 \geq 0$ or $\Delta \leq 0$

In this scenario, the maximum of $\ell(\alpha_j)$ is obtained at $\alpha_j^* = 0$. We divide our discussions into three cases: 1) $\Delta < 0$, 2) $\Delta = 0$, and 3) $\Delta > 0$ and $c_3 \geq 0$. In *Case 1*, $h(\alpha_j) = 0$ has only one (negative) real root. In *Case 2*, the equation has three negative roots (two of them coincide), or one negative root and a multiple positive root. In *Case 3* the equation has three distinct negative roots. So in all the cases $\ell'(\alpha_j) \leq 0$ for $\alpha_j \geq 0$, resulting in that $\ell(\alpha_j)$ decreases monotonically on $[0, +\infty)$.

Based on the analysis above, we can compute efficiently α_j^* (given s_j and q_j) at which the likelihood is maximized with respect to a single basis vector $\mathbf{a}_j$, $j = 1, \dots, N$.

So, the likelihood consistently increases if we update a single α_j at one time with the basis $\mathbf{a}_j$, $j = 1, \dots, N$, properly chosen. As a result, a greedy algorithm can be implemented. At the beginning, no basis vectors are included in the model (i.e., the active set is empty or all $\alpha_j = 0$). Then choose a vector $\mathbf{a}_j$ at each iteration such that it gives the largest likelihood increment by updating α_j from its current value α_j^0 to the maximum point α_j^* . If $\alpha_j^0 = 0$ and $\alpha_j^* > 0$, then the basis vector $\mathbf{a}_j$ is added to the model with $\alpha_j = \alpha_j^*$. If $\alpha_j^0 > 0$ and $\alpha_j^* > 0$, then α_j is re-estimated in the model. If $\alpha_j^0 > 0$ and $\alpha_j^* = 0$, then $\mathbf{a}_j$ is removed from the model with $\alpha_j = 0$. After that, update η as in Subsection 6.2.2. The process is repeated until convergence (that is guaranteed since $\mathcal{L}$ increases monotonically). Based on the results of [57], we see that $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, s_j and q_j , $j = 1, \dots, N$, can be efficiently updated without matrix inversions. Consequently, the greedy algorithm is computationally efficient.

Remark 6.2. *The main differences between the proposed algorithm and those in [46, 57] are the updates of $\boldsymbol{\alpha}$ and η . Since roots of a cubic equation have explicit expressions and the update of η hardly depends on the problem dimension, the proposed greedy algorithm has the same computational complexity as those in [46, 57] at each iteration.*

6.3.2 Analysis of Basis Selection Condition

We study the basis selection condition of the fast algorithm in more details in this subsection. Based on the analysis in Subsection 6.3.1, we have the following result.

Proposition 6.1. *Suppose that the log-likelihood $\mathcal{L}$ has been locally maximized in the fast algorithm. If $\alpha_j > 0$ for some basis vector $\mathbf{a}_j$, $j = 1, \dots, N$, then $q_j^2 > \min \left\{ s_j + 2\eta + \frac{2-2\epsilon}{\tau}, (5-4\epsilon)s_j + 2\eta + \tau(4\eta s_j + s_j^2) \right\}$ where s_j , q_j are as defined in Subsection 6.3.1.*

Proof. Note that the inequalities $q_j^2 > s_j + 2\eta + \frac{2-2\epsilon}{\tau}$ and $q_j^2 > (5-4\epsilon)s_j + 2\eta + \tau(4\eta s_j + s_j^2)$ are equivalent to $c_4 < 0$ and $c_3 < 0$ respectively. Let us suppose that the conclusion does not hold. Then we have $c_4 \geq 0$ and $c_3 \geq 0$. It follows from the

analysis in Subsection 6.3.1.3 that the likelihood increases if $\mathbf{a}_j$ is removed from the model, leading to contradiction. ■

Remark 6.3. *Proposition 6.1 provides a necessary condition for the basis vectors in the final active set. Note that this condition is generally insufficient since additional requirements are needed in the case of $c_4 \geq 0$ (e.g., $\Delta > 0$) according to the analysis in Subsection 6.3.1. In a special case where $q_j^2 > s_j + 2\eta + \frac{2-2\epsilon}{\tau}$ holds, we can conclude that $\mathbf{a}_j$ is in the model according to Subsection 6.3.1.1.*

The basis selection condition concerns the sparsity level of the solution of the fast algorithm. We have illustrated that different settings of τ and ϵ lead to different sparsity-inducing priors in Section 6.1. In the following we will see how the parameters affect the sparsity of the algorithm solution. We discuss the effects of τ and ϵ separately.

Let us first fix $\tau > 0$ and vary $\epsilon \in [0, 1]$. It is easy to show that as ϵ decreases both the terms $s_j + 2\eta + \frac{2-2\epsilon}{\tau}$ and $(5 - 4\epsilon)s_j + 2\eta + \tau(4\eta s_j + s_j^2)$ increase. Thus the necessary condition in Proposition 6.1 becomes stronger. As a result, the solution of the greedy algorithm will be sparser, which is consistent with the fact that a smaller ϵ leads to a more sparsity-inducing prior as shown in Section 6.1. In the case of $\epsilon = 1$, the inequality turns to be $q_j^2 > s_j + 2\eta$ that coincides with the result in [46]. So, the proposed algorithm will produce a sparser solution than that of [46] by simply setting $\epsilon < 1$.

It is not obvious for the case of fixed $\epsilon < 1$ and varying τ (the case $\epsilon = 1$ has been discussed before) since, as τ decreases, the first term $s_j + 2\eta + \frac{2-2\epsilon}{\tau}$ increases while the second term $(5 - 4\epsilon)s_j + 2\eta + \tau(4\eta s_j + s_j^2)$ decreases. To make a correct conclusion, we observe that $\ell(\alpha_j)$ in Subsection 6.3.1 is a strictly increasing function with respect to $\tau > 0$ for any fixed $\alpha_j > 0$ and keeps equal to zero at $\alpha_j = 0$. As a result, as τ decreases, it is less likely that the maximum of $\ell(\alpha_j)$ is achieved at a positive point, resulting in that the solution of the greedy algorithm gets sparser. This is consistent with that a smaller τ leads to a more sparsity-inducing prior as

shown in Section 6.1. In fact, the greedy algorithm produces a zero solution if $\tau = 0$ and $\epsilon < 1$ since in such a case the maximum of the likelihood is always achieved at the origin with respect to every basis vector. So intuitively, a smaller τ should be used to obtain a more sparsity-inducing prior but too small τ may lead to inaccuracy for the greedy algorithm. Numerical simulations in Section 6.4 will illustrate that the recommended τ in Subsection 6.1.2 is a good choice.

Remark 6.4. *In the case of $\tau = 0$, the proposed G-STG prior coincides with the Gaussian-gamma prior in [54]. The greedy algorithm using this prior developed in [54] is claimed to follow from the same framework in [57] and maximize the likelihood sequentially. However, it should be noted that the algorithm in [54] does not really maximize the likelihood since, if it does, then it should produce a zero solution as discussed above. Specifically, the authors of [54] compute only a local maximum point of $\ell(\alpha_j)$ which cannot guarantee to increase the likelihood $\mathcal{L}$ while the global maximum of $\ell(\alpha_j)$ is always obtained at the origin. Hence the algorithm in [54] is technically incorrect because of the inappropriate basis update scheme. Moreover, the algorithm in [54] has not been shown to provide guaranteed convergence.*

6.4 Numerical Simulations

In this section, we present numerical results to illustrate the performance of the proposed method (we consider only the fast algorithm in Subsection 6.3.1). We consider both one-dimensional synthetic signals and two-dimensional images, and compare with existing methods, including ℓ_1 optimization (BP or BPDN), reweighted (RW-) ℓ_1 optimization [128], StOMP [129], the basic BCS [57] and BCS with the Laplace prior (denoted by Laplace) [46]. BCS, Laplace and the proposed method are SBL methods. ℓ_1 optimization is a convex optimization method. RW- ℓ_1 is related to non-convex optimization. StOMP is a greedy method. The Matlab codes of BP, StOMP and BCS are obtained from the SparseLab package¹, and that of BPDN is from the

¹Available at <http://sparselab.stanford.edu>.

ℓ_1 -magic package². When composing this thesis, we have also repeated all simulations by implementing BP and BPDN using a recent CS algorithm SPGL1 [130], which results in slightly different accuracy since the same problem is solved. Its computational speed will also be reported. The code of Laplace is available at <https://netfiles.uiuc.edu/dbabacan/www/links.html>. The number of iterations is set to 5 for RW- ℓ_1 (i.e., 5 ℓ_1 minimization problems are solved iteratively). To make a fair comparison, we use the same convergence criterion in the proposed method as in BCS and Laplace with the stopping tolerance set to 1×10^{-8} .

The performance metrics adopted include the relative mean squared error (RMSE, calculated by $\|\hat{\mathbf{x}} - \mathbf{x}\|_2^2 / \|\mathbf{x}\|_2^2$), the support size ($\|\hat{\mathbf{x}}\|_0$), the number of iterations (for the three BCS methods) and the CPU time, where $\hat{\mathbf{x}}$ and $\mathbf{x}$ denote the recovered and original signals, respectively.

6.4.1 One-Dimensional Synthetic Signals

6.4.1.1 Performance with respect to τ

An explicit determination of τ has been recommended in Subsection 6.1.2. In the following, we show that, indeed, this setting leads to good performance in the signal recovery. In our simulation, we set the signal length $N = 512$ and the number of nonzero entries $K = 20$, and vary the sample size M from 40 to 140 with step size of 5. The nonzero entries of the sparse signal are randomly located with the amplitudes following from a zero-mean unit-variance Gaussian distribution. We consider two matrix ensembles for the sensing matrix $\mathbf{A}$, including Gaussian ensemble and uniform spherical ensemble (with columns uniformly distributed on the sphere $\mathbb{S}^{N-1}$). To obtain the desired SNR, zero-mean AWGNs are added to the linear measurements where the noise variance is set to $\sigma^2 = (K/M) 10^{-\text{SNR}/10}$. We set $\text{SNR} = 25$ dB. The noisy measurements are used in the following signal recovery process. In the proposed algorithm, we set $\epsilon = 0.01$ which results in both fast and accurate recovery

²Available at <http://users.ece.gatech.edu/~justin/l1magic>.

(this will be illustrated in the next experiment). Denote $\tau_0 = (M/N) \sigma^2$. We set $\tau = \theta \tau_0$ and consider six values of $\theta = 10^{-4}, 10^{-2}, 10^{-1}, 1, 10$ and 100 . Thus $\theta = 1$ leads to the recommended value of τ in Subsection 6.1.2. For each M , 100 random problems are generated and solved respectively using the proposed method with different τ . The metrics are averaged results over the 100 trials.

Our simulation results are presented in Fig. 6.2. Fig. 6.2(a) and Fig. 6.2(b) plot the RMSEs and support sizes of the proposed algorithm with Gaussian sensing matrices. It is shown that the recommended τ leads to approximately the smallest error with a reasonable number of measurements while the errors are almost the same when the sample size is small for different τ 's. Fig. 6.2(b) shows that the recommended τ results in the most accurate estimation of the support size in most cases. In addition, it is shown that a sparser solution is obtained if a smaller τ is used in the algorithm as expected. Almost identical performance is shown in Fig. 6.2(c) and 6.2(d) by using the uniform spherical ensemble. Thus, we consider only the uniform spherical ensemble in the following experiments.

6.4.1.2 Performance with respect to ϵ

We study now the performance of the proposed algorithm with respect to ϵ . We repeat the simulation above using the recommended τ and consider five values of $\epsilon = 0, 0.01, 0.1, 0.5$ and 1 . Note that the case $\epsilon = 1$ corresponds to the Laplace prior. Our simulation results are presented in Fig. 6.3. It is shown in Fig. 6.3(a) that the signal recovery error decays as the sample size increases in general. As the sample size is small the estimation errors differ slightly. But with a reasonable number of measurements a smaller ϵ results in a smaller error. It is shown in Fig. 6.3(b) that a smaller ϵ leads to a sparser solution as expected and more accurate support size estimation. Another advantage of adopting a small ϵ can be observed in Fig. 6.3(c) where it is shown that a smaller ϵ leads to less number of iterations. In general, the time consumption is proportional to the number of iterations since the computational workload is approximately the same at each iteration. Fig. 6.3(d)

Figure 6.2: Performance of the proposed algorithm with respect to different settings of τ with (a) & (b) Gaussian ensemble, and (c) & (d) uniform spherical ensemble.

Figure 6.3: Performance of the proposed algorithm with respect to different settings of ϵ .

shows an exception at $\epsilon = 0$ as illustrated in Remark 6.1. In this case, the update of η takes most of the computational time in our simulation. Since it is shown in Figs. 6.3(a) – 6.3(c) that the performance at $\epsilon = 0$ and 0.01 is hardly distinguishable, we use $\epsilon = 0.01$ in the rest simulations.

6.4.1.3 Comparison with existing methods

We consider two simulation setups. In the first case we repeat the simulations above (i.e., we fix the SNR = 25 dB and vary M). In each trial, all methods share the same data. We adopt the FAR thresholding strategy in StOMP. Our simulation results are presented in Fig. 6.4. It is shown in Fig. 6.4(a) that the reconstruction errors of the three SBL methods (BCS, Laplace and our proposed method) are very close to each other and larger than those of BPDN, RW-BPDN and StOMP if the sample

size is small. With a reasonable sample size it can be seen that our proposed method has the smallest error. Fig. 6.4(b) shows the average support size of the recovered signal. The results of BPDN and RW-BPDN are omitted since they are global optimization methods and their numerical solutions have no exact zero entries. In general, the estimated support sizes of StOMP, BCS and Laplace increase with the sample size. As expected the proposed method produces sparser solutions than BCS and Laplace. It is shown that the proposed method can accurately estimate the support of the sparse signal in most cases and has the best performance. Fig. 6.4(c) plots the number of iterations of the three SBL methods, where it is shown that the proposed one uses the least number of iterations and thus is the fastest one in computational speed. On average, StOMP uses the least computation time (about 0.01s), followed by the proposed method, BCS and Laplace (from 0.06 to 0.1s), and then BPDN (about 1s) and RW-BPDN (about 2s). Note also that when implemented with SPGL1, the speed of BPDN is comparable with the proposed method.

In the next simulation we set the sample size $M = 120$ and vary the SNR from 0 to 50dB with step size of 5dB. The simulation results are presented in Fig. 6.5. Fig. 6.5(a) shows that the proposed method has consistently the smallest signal recovery error. Fig. 6.5(b) shows that the proposed method produces the sparsest solution and the most accurate support size estimation. Fig. 6.5(c) shows that among the three SBL methods the proposed one uses the least number of iterations at all SNR levels. In the low SNR regime, it can be 6 and 3 times less in comparison with Laplace and BCS respectively, leading to that the proposed method is much faster than Laplace and BCS. As in the last simulation, in general SPGL1 and the proposed method is comparable in speed. The difference is that the computational time of SPGL1 increases with respect to the SNR while it decreases for the proposed method.

In summary, the proposed method has improved performance for sparse signals in comparison with existing ones. It outperforms its SBL peers in both signal recovery accuracy and computational speed.

Figure 6.4: Performance comparison of the proposed algorithm with existing ones with $(N, K, \text{SNR}) = (512, 20, 25\text{dB})$.

Figure 6.5: Performance comparison of the proposed algorithm with existing ones with $(N, M, K) = (512, 120, 20)$.

6.4.2 Images

In this section, we revisit the widely used multiscale CS reconstruction [131] of the 512×512 Mondrian image in SparseLab. We use the same simulation setup, i.e., we choose the “symmlet8” wavelet as the sparsifying basis with a coarsest scale $j_0 = 4$ and a finest scale $j_1 = 6$. The number of wavelet samples is $N = 4096$ and the sample size of CS methods is $M = 2713$. The parameters of BP and StOMP with the FDR and FAR thresholding strategies (denoted by FDR and FAR respectively) are set as in SparseLab. Since the wavelet expansion of the Mondrian image is compressible but not exactly sparse, we set $\sigma^2 = 0.01 \text{Var}(\mathbf{y})$ in Laplace and our proposed method as in BCS, where $\text{Var}(\mathbf{y})$ denotes the variance of the entries of $\mathbf{y}$.

Table 6.1 presents the experimental results over 100 trials. Linear reconstruction from 4096 wavelet samples has a reconstruction error of 0.1333 that represents a lower bound of the error of the considered CS methods. The global optimization method BP has the smallest error among the CS methods, followed by BCS, Laplace, the proposed method and StOMP. The presented results verify again that the proposed method produces a sparser solution than BCS and Laplace. In fact, it produces the sparsest solution among all the methods. So it is reasonable that the proposed method has a slightly worse reconstruction error in comparison with BCS and Laplace since the original signal is not exactly sparse. We note that the proposed method is faster than BCS and Laplace. FDR uses the least time but has the worst accuracy. In comparison with FAR, the proposed method is slightly slower but more accurate. Finally, it can be observed that the proposed method has the most stable performance among the CS methods except BP by comparing the standard deviation of the metrics. When implemented with SPGL1, BP takes about four times computational time of the proposed method, which is consistent with an observation in the last simulation that SPGL1 takes longer time as the SNR increases. Fig. 6.6 shows a typical example of reconstructed images where faithful reconstructions of the Mondrian image can be observed.

In summary, the proposed method has satisfactory performance for compressible

Table 6.1: Averaged Relative MSEs, CPU Times and Number of Nonzero Entries (mean $\pm$ standard deviation) for Multiscale CS Reconstruction of the Mondrian Image.

	RMSE	Time (s)	# Nonzeros
Linear	0.1333	—	4096
BP	0.1393 ± 0.0008	42.2 ± 4.03	4096 ± 0
FDR	0.1999 ± 0.0487	8.84 ± 2.12	2155 ± 122
FAR	0.1499 ± 0.0033	17.0 ± 4.35	1142 ± 41
BCS	0.1423 ± 0.0023	27.2 ± 5.92	1305 ± 67
Laplace	0.1429 ± 0.0013	25.7 ± 6.11	1218 ± 65
Proposed	0.1448 ± 0.0011	21.3 ± 4.09	1049 ± 21

signals and there exists a tradeoff between its recovery accuracy and estimated signal sparsity. In particular, a smaller ϵ leads to a sparser signal estimate which may be less accurate, while a larger ϵ (as for Laplace where $\epsilon = 1$) leads to a less sparse but more accurate signal estimate since more basis vectors contribute to the estimate.

6.5 Conclusion

The sparse signal recovery problem in CS was studied in this chapter. Within the framework of Bayesian CS, a new hierarchical sparsity-inducing prior was introduced and efficient signal recovery algorithms were developed. Similar theoretical results on the global and local optimizers of the proposed method were proven as that for the basic SBL. The main algorithm was shown to produce sparser solutions than its existing SBL peers. Numerical simulations were carried out to demonstrate the performance of the proposed sparsity-inducing prior and solution. The proposed G-STG prior preserves the general structure of existing hierarchical sparsity-inducing priors and can be implemented in other SBL-based methods with ease.

Figure 6.6: The 512×512 Mondrian image (a) and its reconstructions using (b) linear reconstruction (RMSE = 0.1333) from $N = 4096$ wavelet samples and a multiscale CS scheme from $M = 2713$ linear measurements by (c) BP (RMSE = 0.1391, time = 44.7s and # nonzeros = 4096), (d) StOMP with FDR thresholding (RMSE = 0.1751, time = 10.3s and # nonzeros = 2014), (e) StOMP with FAR thresholding (RMSE = 0.1529, time = 15.7s and # nonzeros = 1088), (f) BCS (RMSE = 0.1448, time = 24.1s and # nonzeros = 1293), (g) Laplace (RMSE = 0.1427, time = 25.5s and # nonzeros = 1229) and (h) our proposed method (RMSE = 0.1440, time = 19.8s and # nonzeros = 1033).

Chapter 7

Variational Bayesian Algorithm for Quantized Compressed Sensing

As Chapter 6, this chapter is built upon SBL. However, in this chapter we are concerned about the CS problem with quantized measurements, known as the quantized CS problem. Notice that quantization is inevitable in practical considerations though most studies of CS are focused on the standard formulation where the compressive measurements have infinite bit precision. Two scenarios are mainly studied in the literature which are categorized according to the number of bits per measurement used and include multi-bit CS (≥ 2 bits) [1, 58–61, 132] and 1-bit CS (a single bit) [62–68]. Since only the sign information of the measurements is retained in the 1-bit case and signal scaling is lost, it is quite different from the multi-bit CS and has been separately studied. In this chapter, we introduce a Bayesian framework for quantized CS that unifies the multi- and 1-bit cases. The new framework deals with quantization errors and measurement noises separately, allows data saturation in multi-bit CS, and does not need the signal sparsity information. Based on the new problem formulation, we propose an algorithm within the Bayesian CS framework where the quantization errors are modeled as random variables and jointly estimated with the signal of interest. The performance of the proposed algorithm is studied by extensive numerical simulations in various scenarios. This chapter is mainly based on [9].

7.1 A New Framework for Quantized CS

Multi- and 1-bit CS problems are typically studied separately in the literature due to their big difference. In this section, we propose a Bayesian framework that unifies both the cases. The new framework is applicable to various scenarios including noiseless/noisy environment and unsaturated/saturated quantizer. Its relations with existing methods are studied through an MAP interpretation.

7.1.1 A Unified Observation Model

In quantized CS, the observed samples are noisy linear measurements of the original signal after quantization:

$$\mathbf{z} = \mathcal{Q}(\mathbf{y}), \quad \mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \quad (7.1)$$

where $\mathbf{x}$ is the signal of interest, $\mathbf{A}$ is the sensing matrix, $\mathbf{n}$ is the measurement noise vector, $\mathbf{y}$ is the pre-quantized noisy measurement vector, $\mathcal{Q}$ denotes a quantizer and $\mathbf{z}$ is the observation vector. Notice that the noise is not denoted by $\mathbf{e}$ as in the previous chapters. Later we will see that $\mathbf{e}$ is reserved for the quantization error. A quantizer $\mathcal{Q}(v)$ for a scalar $v \in \mathbb{R}$ is defined as

$$\mathcal{Q}(v) = \begin{cases} v_0, & \text{if } v \in (u_0, u_1), \\ v_1, & \text{if } v \in [u_1, u_2), \\ \cdots, & \cdots, \\ v_{L-1}, & \text{if } v \in [u_{L-1}, u_L), \end{cases} \quad (7.2)$$

where L denotes the number of the quantization levels and typically satisfies $L = 2^B$ with B denoting the bit depth (bits per quantized measurement), $u_0 < u_1 < \cdots < u_L$, and $v_i \in [u_i, u_{i+1})$ for $i = 0, \cdots, L-1$. The quantizer $\mathcal{Q}(v)$ is called unsaturated if (u_0, u_L) is a finite interval, or saturated otherwise. For a vector $\mathbf{v}$, $\mathcal{Q}(\mathbf{v})$ operates elementwise. Multi-bit CS refers to the case $B \geq 2$ while 1-bit CS corresponds to $B = 1$.

7.1.1.1 Multi-bit CS

We consider first a multi-bit quantizer where $B \geq 2$. Denote $\mathcal{D}_y$ the domain of $\mathbf{y}$. Then we have

$$\mathcal{D}_y = \mathcal{Q}^{-1}(\mathbf{z}) := \{\mathbf{y} \in \mathbb{R}^M | \mathcal{Q}(\mathbf{y}) = \mathbf{z}\}. \quad (7.3)$$

We introduce an auxiliary variable $\mathbf{e} = \mathbf{z} - \mathbf{y}$ denoting the quantization error with its domain

$$\mathcal{D}_e = \mathbf{z} - \mathcal{D}_y := \{\mathbf{z} - \mathbf{y} | \mathbf{y} \in \mathcal{D}_y\}. \quad (7.4)$$

Note that $\mathcal{D}_e$ is unbounded when data saturation occurs.

7.1.1.2 1-bit CS

In the case of 1-bit quantizer we set $u_0 = -\infty$, $u_1 = 0$ and $u_2 = +\infty$. The sign information of $\mathbf{y}$ is preserved in the quantized measurement $\mathbf{z}$. But the scaling information of $\mathbf{y}$ and $\mathbf{x}$ is lost. Without loss of generality, we let the 1-bit quantizer for a scalar $v \in \mathbb{R}$ be

$$\mathcal{Q}(v) = \varsigma \operatorname{sgn}(v),$$

where $\varsigma \rightarrow 0_+$ (ς is an arbitrarily small positive number) and $\operatorname{sgn}(\cdot)$ is the sign function. For convenience, we set $\operatorname{sgn}(0) = 1$ (the choice is arbitrary and can be replaced by $\operatorname{sgn}(0) = -1$). Then we have $\mathbf{z} \rightarrow \mathbf{0}$. To solve the signal scaling problem we impose a constraint that $\mathbf{y}$ has fixed unit norm, i.e.,

$$\|\mathbf{y}\|_s = 1 \quad (7.5)$$

with $s \geq 1$. Different from the multi-bit quantizer case we have in such a case that

$$\mathcal{D}_y = \{\mathbf{y} \in \mathbb{R}^M | \operatorname{sgn}(\mathbf{y}) = \operatorname{sgn}(\mathbf{z}), \|\mathbf{y}\|_s = 1\}, \quad (7.6)$$

$$\mathcal{D}_e = \{\mathbf{e} \in \mathbb{R}^M | \operatorname{sgn}(\mathbf{e}) = -\operatorname{sgn}(\mathbf{z}), \|\mathbf{e}\|_s = 1\}. \quad (7.7)$$

As a result, an observation model that unifies the multi- and 1-bit CS problems can

be written into

$$\mathbf{z} = \mathbf{A}\mathbf{x} + \mathbf{e} + \mathbf{n}, \quad \mathbf{e} \in \mathcal{D}_e, \quad (7.8)$$

which is the observation model to be used in this chapter to recover $\mathbf{x}$.

Remark 7.1. *In practice, one is able to know the domain of $\mathbf{e}$ but it is difficult to characterize its exact relationship with $\mathbf{x}$. Thus, the dependence of $\mathbf{e}$ on $\mathbf{x}$ is dropped when we write (7.1) into (7.8) which is the observation model we use for the signal recovery, i.e., the only information we attempt to exploit during the signal recovery process is its domain. We note that the signal recovery performance may be improved if the dependence can be properly exploited which, however, is rather difficult.*

7.1.2 Bayesian Formulation of Quantized CS

In this subsection we formulate the quantized CS problem from a Bayesian perspective based on the observation model in (7.8). According to Remark 7.1 we treat $\mathbf{e}$ as a random variable independent of $\mathbf{x}$. The joint probability density function (PDF) $p(\mathbf{z}, \mathbf{x}, \mathbf{e})$ is decomposed as

$$p(\mathbf{z}, \mathbf{x}, \mathbf{e}) = p(\mathbf{z}|\mathbf{x}, \mathbf{e}) p(\mathbf{x}) p(\mathbf{e}).$$

We define the three distributions on the right hand side as follows.

7.1.2.1 Noise model

Under an assumption of white Gaussian measurement noise, i.e., $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ where σ^2 is the noise variance and $\mathbf{I}$ denotes an identity matrix of proper dimension, we have

$$p(\mathbf{z}|\mathbf{x}, \mathbf{e}; \sigma^2) = \mathcal{N}(\mathbf{z}|\mathbf{A}\mathbf{x} + \mathbf{e}, \sigma^2 \mathbf{I}). \quad (7.9)$$

7.1.2.2 Sparse signal model

A sparse prior is needed for the sparse signal $\mathbf{x}$ of interest. Here we do not give an explicit distribution to the sparse signal $\mathbf{x}$ but denote $p(\mathbf{x})$ its PDF. Then we let $f(\mathbf{x}) = -C_1 \log p(\mathbf{x}) + C_2$ where C_1 and C_2 are proper constants. The only thing that we assume for $p(\mathbf{x})$ is that it favors entries of $\mathbf{x}$ being zero. As an example, a commonly used sparse prior for $\mathbf{x}$ is a Laplace prior [133, 134]: $p(\mathbf{x}) = \lambda^N \exp\{-\lambda \|\mathbf{x}\|_1\}$ with λ being a positive constant. In such a case, we have $f(\mathbf{x}) = \|\mathbf{x}\|_1$ that has been extensively studied in deterministic optimization methods.

7.1.2.3 Quantization error model

We assume a uniform, noninformative prior for $\mathbf{e}$:

$$\mathbf{e} \sim U(\mathcal{D}_e) \quad (7.10)$$

since the only information of $\mathbf{e}$ that we use is $\mathbf{e} \in \mathcal{D}_e$.

Remark 7.2. *The uniform prior may not characterize well the quantization error in the case of a very small bit depth B . But it is noted that a sophisticated prior needs more information besides the domain $\mathcal{D}_e$ which is difficult to obtain.*

To obtain an MAP estimator of $\mathbf{x}$ requires to integrate out $\mathbf{e}$ from $p(\mathbf{z}, \mathbf{x}, \mathbf{e})$ that is computationally intractable. We propose to estimate $\mathbf{x}$ and $\mathbf{e}$ simultaneously using their joint MAP estimator:

$$\begin{aligned} \{\hat{\mathbf{x}}, \hat{\mathbf{e}}\} &= \arg \max_{\mathbf{x}, \mathbf{e}} \log p(\mathbf{x}, \mathbf{e} | \mathbf{z}) \\ &= \arg \max_{\mathbf{x}, \mathbf{e}} \log p(\mathbf{z}, \mathbf{x}, \mathbf{e}) \\ &= \arg \max_{\mathbf{x}, \mathbf{e}} \log \{p(\mathbf{z} | \mathbf{x}, \mathbf{e}) p(\mathbf{x}) p(\mathbf{e})\} \\ &= \arg \min_{\mathbf{x}, \mathbf{e} \in \mathcal{D}_e} \left\{ f(\mathbf{x}) + \frac{C_1}{2\sigma^2} \|\mathbf{z} - \mathbf{e} - \mathbf{A}\mathbf{x}\|_2^2 \right\}. \end{aligned} \quad (7.11)$$

An equivalent form of the problem in (7.11) is

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{e}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\mathbf{z} - \tilde{\mathbf{e}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \tilde{\mathbf{e}} \in \mathcal{D}_e, \end{cases} \quad (7.12)$$

where ϵ is a proper scalar that controls the noise energy. The first constraint in (7.12) is to ensure the data consistency against the measurement noise. In multi-bit CS, the second one concerns data consistency due to quantization. In 1-bit CS, an additional signal scaling constraint is included in the second constraint that prevents an optimal solution for $\mathbf{x}$ from $\mathbf{0}$. Before proceeding to our algorithm within the framework of Bayesian CS, we study in the next subsection relations of the proposed Bayesian framework with existing methods.

7.1.3 Relations with Existing Methods in Quantized CS

We first note that problem (7.12) is equivalent to the problem

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{y}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\tilde{\mathbf{y}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \tilde{\mathbf{y}} \in \mathcal{D}_y. \end{cases} \quad (7.13)$$

In the following we show that many existing problem formulations of quantized CS are special cases of or related to (7.13).

7.1.3.1 Multi-bit CS

We consider the case of ℓ_1 optimization where $f(\mathbf{x}) = \|\mathbf{x}\|_1$. In the noise free case where $\epsilon = 0$, the problem in (7.13) can be written into

$$\min_{\tilde{\mathbf{x}}} \|\tilde{\mathbf{x}}\|_1, \text{ subject to } \mathbf{A}\tilde{\mathbf{x}} \in \mathcal{D}_y, \quad (7.14)$$

which has been studied in [132]. Further, by assuming that $\mathcal{Q}$ is a uniform unsaturated quantizer the above problem can be written into

$$\min_{\tilde{\mathbf{x}}} \|\tilde{\mathbf{x}}\|_1, \text{ subject to } \|\mathbf{z} - \mathbf{A}\tilde{\mathbf{x}}\|_\infty \leq \frac{r}{2} \quad (7.15)$$

which is studied in [1, 59] with r denoting the quantization bin width. While existing methods that account for measurement noise typically mix it up with the quantization error, e.g., in [60], problem (7.13) extends existing noise free formulations to the noisy case by dealing with the two uncertainties separately.

Remark 7.3. *Under the assumption that all quantization errors are independent and uniformly distributed in a common interval $[-\frac{r}{2}, \frac{r}{2}]$, it is shown in [59] that the ℓ_∞ norm in problem (7.15) is not the best choice for the signal recovery. But it is unclear whether the result in [59] can be extended to the case of a general quantizer where the above assumption fails. It is noted that our problem formulation does not require this assumption and applies to an arbitrary quantizer. That is, by losing some optimality, we have obtained the universality.*

7.1.3.2 1-bit CS

In 1-bit CS (7.13) becomes

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{y}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\tilde{\mathbf{y}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \text{sgn}(\tilde{\mathbf{y}}) = \text{sgn}(\mathbf{z}), \\ \|\tilde{\mathbf{y}}\|_s = 1. \end{cases} \quad (7.16)$$

In the noise free case, it can be written into

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \text{sgn}(\mathbf{A}\tilde{\mathbf{x}}) = \text{sgn}(\mathbf{z}), \\ \|\mathbf{A}\tilde{\mathbf{x}}\|_s = 1. \end{cases} \quad (7.17)$$

This problem with the settings $f(\mathbf{x}) = \|\mathbf{x}\|_1$ and $s = 1$ can be shown to be convex and has been studied in [67]. So by (7.16) we extend (7.17) to the noisy case while

the authors of [67] state in [68] that “it was unclear how to modify the above convex program to account for possible noise.”

The third constraint in (7.16) serves to prevent the optimal solution for $\mathbf{x}$ from $\mathbf{0}$. If replacing it by $\|\tilde{\mathbf{x}}\|_2 = 1$, then problem (7.16) can be shown to be equivalent to the problem

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_2 \leq \epsilon, \\ \|\tilde{\mathbf{x}}\|_2 = 1, \end{cases} \quad (7.18)$$

where $(v)_- = \max\{-v, 0\}$ for a scalar v and operates elementwise for a vector. In the noise free case (7.18) becomes

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \text{sgn}(\mathbf{A}\tilde{\mathbf{x}}) = \text{sgn}(\mathbf{z}), \\ \|\tilde{\mathbf{x}}\|_2 = 1, \end{cases} \quad (7.19)$$

which with $f(\mathbf{x}) = \|\mathbf{x}\|_1$ is the earliest formulation of the 1-bit CS problem introduced in [62] and solved using RFPI in [62] and RSS in [65]. Assume that the signal sparsity information is known instead of the noise energy, another formulation of (7.18) is to pose $\|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_2$ as the objective function and $f(\tilde{\mathbf{x}}) \leq S$ as a constraint, where the constant S refers to the sparsity information. Such kind of formulations have been studied in [62, 64, 66]. In addition, the convex program in [68] is related by observing that $\text{sgn}^T(\mathbf{z}) \mathbf{A}\tilde{\mathbf{x}} = \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_+\|_1 - \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_1$, where $(v)_+ = \max\{v, 0\}$.

Remark 7.4. By (7.18) we see that the effective noise is $(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\mathbf{x})_-$ in 1-bit CS where by “effective noise” we refer to a noise that has the minimum energy and leads to the same measurement. Since its energy is much smaller than that of the true noise $\mathbf{y} - \mathbf{A}\mathbf{x}$, from this point of view, we may say that 1-bit CS is robust to the measurement noise.

7.2 Q-VMP: Variational Message Passing for Quantized CS

7.2.1 Model Selection

We assume that the noise variance σ^2 is known. Though it can be estimated by assuming an inverse Gamma prior for it in the case where it is unknown as in [46], its estimate is inaccurate due to an “identifiability issue” as addressed in [41]. For the sparse signal $\mathbf{x}$, we adopt a three-layer, Gaussian-Gamma-Gamma, hierarchical prior introduced in [54]:

$$p(\mathbf{x}; \epsilon, c, d) = \iint p(\mathbf{x}|\boldsymbol{\alpha}) p(\boldsymbol{\alpha}|\eta; \epsilon) p(\eta; c, d) d\boldsymbol{\alpha} d\eta$$

where

$$p(\mathbf{x}|\boldsymbol{\alpha}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \boldsymbol{\Lambda}), \quad (7.20)$$

$$p(\boldsymbol{\alpha}|\eta; \epsilon) = \prod_{i=1}^N \Gamma(\alpha_i|\epsilon, \eta), \quad (7.21)$$

$$p(\eta; c, d) = \Gamma(\eta|c, d) \quad (7.22)$$

with $\boldsymbol{\Lambda} = \text{diag}(\boldsymbol{\alpha})$ and constants ϵ, c, d . For a Gamma distributed variable $u \sim \Gamma(c, d)$, its PDF is $\Gamma(u|c, d) = \frac{d^c}{\Gamma(c)} u^{c-1} \exp(-du)$ with $\Gamma(c)$ being the Gamma function. By [54] the constants ϵ, c, d satisfy that $0 \leq \epsilon \leq 1, c, d \geq 0$. In this chapter, we adopt $c = 1, d = 0$ to make the prior for η in (7.22) noninformative (flat on $\mathbb{R}_+$). Further, we choose $\epsilon = 0$ since a smaller ϵ leads to a sparser prior and an estimator that approximates a hard-thresholding rule according to [54]. Readers are referred to [54] for more properties of the Gaussian-Gamma-Gamma prior and its relations with other sparse estimation techniques.

In 1-bit CS, we let $\mathbf{y}$ have unit ℓ_2 norm in (7.5), leading to that $\|\mathbf{e}\|_2 = 1$ in (7.7).

Figure 7.1: Directed graphical model that encodes the joint PDF in (7.23) of the Bayesian model. Nodes denoted with circles correspond to random variables, while nodes denoted with squares correspond to parameters of the model. Doubly circled z is the observation while single circled nodes represent hidden variables.

As a result, we have the joint PDF of the observation model (7.8):

$$p(z, \mathbf{x}, \mathbf{e}, \boldsymbol{\alpha}, \eta) = p(z|\mathbf{x}, \mathbf{e}) p(\mathbf{x}|\boldsymbol{\alpha}) p(\boldsymbol{\alpha}|\eta) p(\eta) p(\mathbf{e}) \quad (7.23)$$

with the distributions on the right hand side as defined respectively by (7.9), (7.20), (7.21), (7.22) and (7.10). A directed graphical model that encodes the factorization of the joint PDF in (7.23) is shown in Fig. 7.1.

7.2.2 Q-VMP Algorithm

It is well known that Bayesian inference is based on the posterior distribution $p(\mathbf{x}, \mathbf{e}, \boldsymbol{\alpha}, \eta|z) = p(z, \mathbf{x}, \mathbf{e}, \boldsymbol{\alpha}, \eta) / p(z)$. However, such an exact posterior distribution is intractable since $p(z) = \int \cdots \int p(z, \mathbf{x}, \mathbf{e}, \boldsymbol{\alpha}, \eta) d\mathbf{x} d\mathbf{e} d\boldsymbol{\alpha} d\eta$ cannot be expressed explicitly.

A variational inference approach [135, 136] is adopted in this chapter. Denote $\mathbf{V} = \{\mathbf{x}, \mathbf{e}, \boldsymbol{\alpha}, \eta\}$ the set of all unknown variables to be estimated. The goal in variational inference is to find a tractable distribution $q(\mathbf{V})$ that closely approximates the true posterior distribution $p(\mathbf{V}|z)$. To do this, some family of distributions that has enough flexibility is firstly chosen to represent $q(\mathbf{V})$. Then the task is to find a member of the family that minimizes the Kullback-Leibler (KL) divergence between the true posterior $p(\mathbf{V}|z)$ and the variational approximation $q(\mathbf{V})$. A commonly used variational distribution $q(\mathbf{V})$ is such that disjoint groups of variables are independent, i.e., $q(\mathbf{V})$ has a factorized form $q(\mathbf{V}) = q(\mathbf{x}) q(\mathbf{e}) q(\boldsymbol{\alpha}) q(\eta)$. Variational message passing (VMP) is proposed in [137] for the variational inference using a

message passing procedure on a graphical model. In VMP, the variational distributions $q(\mathbf{x})$, $q(\mathbf{e})$, $q(\boldsymbol{\alpha})$, $q(\eta)$ are iteratively updated to monotonically decrease the KL divergence and thus has guaranteed convergence. Readers are referred to [137] for more details of VMP. The updates of $q(\mathbf{x})$, $q(\boldsymbol{\alpha})$, $q(\eta)$ are similar to those in [54] because of the similarity between quantized and conventional CS. $q(\mathbf{e})$ is given complete flexibility in multi-bit CS as $q(\mathbf{x})$, $q(\boldsymbol{\alpha})$ and $q(\eta)$. We constrain $q(\mathbf{e})$ in 1-bit CS such that

$$q(\mathbf{e}) = \delta(\mathbf{e} - \mathbf{e}^0) \quad (7.24)$$

due to a computational issue to be discussed in Remark 7.6, where $\delta(\cdot)$ is the delta function and $\mathbf{e}^0 \in \mathbb{R}^M$ is to be estimated. Note that (7.24) is equivalent to the complete flexibility in the noise free case to be illustrated in Subsection 7.2.3.

Remark 7.5. *In the 1-bit case the convergence of the resulting algorithm is not a direct result of [137] due to the adoption of a degenerate distribution for $\mathbf{e}$. Instead, we may consider $\mathbf{e}$ as an unknown deterministic parameter in such a case. Then the resulting algorithm can be interpreted as a variational EM algorithm [136] and during the iterations the (marginal) likelihood $p(\mathbf{z}; \mathbf{e})$ is guaranteed to monotonically increase and thus convergence is guaranteed. Readers are referred to [136] for the details.*

7.2.2.1 Updates of $q(\mathbf{x})$, $q(\boldsymbol{\alpha})$ and $q(\eta)$

According to [137] we have that

$$\begin{aligned} q(\mathbf{x}) &\propto \exp \left\{ \langle \ln p(\mathbf{z}|\mathbf{x}, \mathbf{e}) \rangle_{q(\mathbf{e})} \langle \ln p(\mathbf{x}|\boldsymbol{\alpha}) \rangle_{q(\boldsymbol{\alpha})} \right\} \\ &\propto \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\}, \end{aligned}$$

and thus $q(\mathbf{x})$ is a Gaussian distribution $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with the mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$:

$$\boldsymbol{\mu} = \langle \mathbf{x} \rangle_{q(\mathbf{x})} = \sigma^{-2} \boldsymbol{\Sigma} \mathbf{A}^T \left(\mathbf{z} - \langle \mathbf{e} \rangle_{q(\mathbf{e})} \right), \quad (7.25)$$

$$\Sigma = \left(\sigma^{-2} \mathbf{A}^T \mathbf{A} + \langle \mathbf{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})} \right)^{-1}. \quad (7.26)$$

For $\boldsymbol{\alpha}$ we have

$$\begin{aligned} q(\boldsymbol{\alpha}) &\propto \exp \left\{ \langle \ln p(\mathbf{x}|\boldsymbol{\alpha}) \rangle_{q(\mathbf{x})} \langle \ln p(\boldsymbol{\alpha}|\boldsymbol{\eta}) \rangle_{q(\boldsymbol{\eta})} \right\} \\ &\propto \prod_{n=1}^N \alpha_n^{\epsilon - \frac{3}{2}} \exp \left\{ -\frac{1}{2} \alpha_n^{-1} \langle x_n^2 \rangle_{q(\mathbf{x})} - \alpha_n \langle \eta \rangle_{q(\boldsymbol{\eta})} \right\}, \end{aligned}$$

where $\langle x_n^2 \rangle_{q(\mathbf{x})} = \mu_n^2 + \Sigma_{nn}$. The expression on the right hand side is the product of generalized inverse Gaussian (GIG) PDFs and thus we have for any $i \in \mathbb{R}$ [138]:

$$\langle \alpha_n^i \rangle_{q(\boldsymbol{\alpha})} = \left(\frac{\langle x_n^2 \rangle_{q(\mathbf{x})}}{2 \langle \eta \rangle_{q(\boldsymbol{\eta})}} \right)^{\frac{i}{2}} \frac{\mathcal{K}_{\epsilon+i-\frac{1}{2}} \left(\sqrt{2 \langle \eta \rangle_{q(\boldsymbol{\eta})} \langle x_n^2 \rangle_{q(\mathbf{x})}} \right)}{\mathcal{K}_{\epsilon-\frac{1}{2}} \left(\sqrt{2 \langle \eta \rangle_{q(\boldsymbol{\eta})} \langle x_n^2 \rangle_{q(\mathbf{x})}} \right)}, \quad (7.27)$$

where $\mathcal{K}_\nu(\cdot)$ is the modified Bessel function of the second kind and order $\nu \in \mathbb{R}$. The case of $i = -1$ in (7.27) gives the evaluation of $\langle \mathbf{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}$ used in (7.26), and the case of $i = 1$ gives the calculation of $\langle \alpha_n \rangle_{q(\boldsymbol{\alpha})}$ used in a later expression in (7.28). The update of $q(\boldsymbol{\eta})$ can be shown to be $q(\boldsymbol{\eta}) = \Gamma \left(\eta | N\epsilon + c, \sum_{n=1}^N \langle \alpha_n \rangle_{q(\boldsymbol{\alpha})} + d \right)$. The first moment of η used in (7.27) is given as

$$\langle \eta \rangle_{q(\boldsymbol{\eta})} = \frac{N\epsilon + c}{\sum_{n=1}^N \langle \alpha_n \rangle_{q(\boldsymbol{\alpha})} + d}. \quad (7.28)$$

7.2.2.2 Update of $q(\mathbf{e})$ in multi-bit CS

In multi-bit CS we have

$$\begin{aligned} q(\mathbf{e}) &\propto \exp \left\{ \langle \ln p(\mathbf{z}|\mathbf{x}, \mathbf{e}) \rangle_{q(\mathbf{x})} \right\} p(\mathbf{e}) \\ &\propto \exp \left\{ -\frac{1}{2} \sigma^{-2} \langle \|\mathbf{z} - \mathbf{e} - \mathbf{A}\mathbf{x}\|_2^2 \rangle_{q(\mathbf{x})} \right\} I_{\mathbf{e}}(\mathcal{D}_{\mathbf{e}}) \\ &\propto \exp \left\{ -\frac{1}{2} \sigma^{-2} \|\mathbf{e} - (\mathbf{z} - \mathbf{A}\boldsymbol{\mu})\|_2^2 \right\} I_{\mathbf{e}}(\mathcal{D}_{\mathbf{e}}), \end{aligned} \quad (7.29)$$

where $I_{\mathbf{e}}(\mathcal{D}_{\mathbf{e}})$ is an indicator function that equals to 1 if $\mathbf{e} \in \mathcal{D}_{\mathbf{e}}$ or 0 otherwise. Hence, $q(\mathbf{e})$ is the product of PDFs of truncated Gaussian distributions, i.e., for

each $m = 1, \dots, M$, $q(e_m)$ is the PDF of a truncated Gaussian distribution. As a result, the first moment of e_m , $m = 1, \dots, M$, used in (7.25) can be given in closed form after some derivations using the PDF $\phi(\cdot)$ and cumulative distribution function (CDF) $\Phi(\cdot)$ of a standard Gaussian distribution:

$$\langle e_m \rangle_{q(\mathbf{e})} = \sigma \frac{\phi(l_{e_m}) - \phi(u_{e_m})}{\Phi(u_{e_m}) - \Phi(l_{e_m})} + \mu_{e_m}, \quad (7.30)$$

where $\mu_{e_m} = (\mathbf{z} - \mathbf{A}\boldsymbol{\mu})_m$, l_{e_m} and u_{e_m} satisfy that $\mathcal{D}_{e_m} = [\sigma l_{e_m} + \mu_{e_m}, \sigma u_{e_m} + \mu_{e_m}]$ with $\mathcal{D}_{e_m}$ denoting the domain of e_m , $\phi(u) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{u^2}{2}\right\}$ and $\Phi(u) = \int_{-\infty}^u \phi(t) dt$ for $u \in \mathbb{R}$.

Remark 7.6. Consider the case where $q(\mathbf{e})$ is given the complete flexibility in 1-bit CS. Note that entries of a point in $\mathcal{D}_e$ are no longer independent of each other in such a case, leading to that $q(\mathbf{e})$ is the PDF of a truncated multi-variable Gaussian distribution with $\mathbf{e}$ constrained in a nonconvex set $\mathcal{D}_e$ defined in (7.7). As a result, the calculation of $\langle \mathbf{e} \rangle_{q(\mathbf{e})}$ is in general computationally intractable in our considered CS problems where the dimension of $\mathbf{e}$ is large.

7.2.2.3 Update of $q(\mathbf{e})$ in 1-bit CS

According to [137], this is equivalent to finding an MAP estimator of $\mathbf{e}$ with its posterior distribution defined in (7.29). So we have

$$\langle \mathbf{e} \rangle_{q(\mathbf{e})} = \mathbf{e}^0 \quad (7.31)$$

with

$$\begin{aligned} \mathbf{e}^0 &= \arg \max_{\mathbf{e} \in \mathbb{R}^M} \left\{ \exp \left\{ -\frac{1}{2} \sigma^{-2} \|\mathbf{e} - (\mathbf{z} - \mathbf{A}\boldsymbol{\mu})\|_2^2 \right\} I_e(\mathcal{D}_e) \right\} \\ &= \arg \min_{\mathbf{e} \in \mathcal{D}_e} \|\mathbf{e} - (\mathbf{z} - \mathbf{A}\boldsymbol{\mu})\|_2 \\ &= \mathcal{P}_{\mathcal{D}_e}(\mathbf{z} - \mathbf{A}\boldsymbol{\mu}), \end{aligned}$$

where $\mathcal{D}_e$ is defined in (7.7) and $\mathcal{P}_{\mathcal{D}}(\mathbf{v})$ denotes a projection of a point $\mathbf{v}$ onto a set $\mathcal{D}$. The calculation of $\mathcal{P}_{\mathcal{D}_e}(\cdot)$ with the nonconvex set $\mathcal{D}_e$ is provided in the following lemma.

Lemma 7.1. *For a vector $\mathbf{v} \in \mathbb{R}^M$, let $\bar{\mathbf{v}} = -\text{sgn}(\mathbf{z}) \odot \mathbf{v}$. Denote $\mathcal{I}$ the index set of all positive entries of $\bar{\mathbf{v}}$. Let $\mathcal{I}^c$ be its complementary set. If $\mathcal{I}$ is nonempty, then let $\mathbf{e}^* \in \mathbb{R}^M$ with $\mathbf{e}_{\mathcal{I}}^* = \frac{\mathbf{v}_{\mathcal{I}}}{\|\mathbf{v}_{\mathcal{I}}\|_2}$ and $\mathbf{e}_{\mathcal{I}^c}^* = \mathbf{0}$. Otherwise, let $i_0 = \arg \max_i (\bar{v}_i)$ and $\mathbf{e}^*$ such that $e_{i_0}^* = -\text{sgn}(z_{i_0})$ and $e_i^* = 0$ whenever $i \neq i_0$. Then $\mathbf{e}^* = \mathcal{P}_{\mathcal{D}_e}(\mathbf{v})$ with $\mathcal{D}_e$ as defined in (7.7).*

Proof. It is easy to show the following equivalences:

$$\begin{aligned}
\mathbf{e}^* &= \mathcal{P}_{\mathcal{D}_e}(\mathbf{v}) \\
&\Leftrightarrow \mathbf{e}^* = \arg \min_{\mathbf{e} \in \mathcal{D}_e} \|\mathbf{e} - \mathbf{v}\|_2 \\
&\Leftrightarrow \mathbf{e}^* = \arg \max_{\mathbf{e} \in \mathcal{D}_e} \mathbf{v}^T \mathbf{e} \\
&\Leftrightarrow -\text{sgn}(\mathbf{z}) \odot \mathbf{e}^* = \arg \max_{\mathbf{w}} f(\mathbf{w}) = \bar{\mathbf{v}}^T \mathbf{w}, \\
&\text{subject to } \|\mathbf{w}\|_2 = 1 \text{ and } \mathbf{w} \succeq \mathbf{0}.
\end{aligned} \tag{7.32}$$

1) $\mathcal{I}$ is nonempty. Note that $\|\mathbf{w}_{\mathcal{I}}\|_2 \leq 1$ and $\bar{\mathbf{v}}_{\mathcal{I}^c} \preceq \mathbf{0}$. By the Cauchy inequality,

$$\begin{aligned}
f(\mathbf{w}) &= \bar{\mathbf{v}}_{\mathcal{I}}^T \mathbf{w}_{\mathcal{I}} + \bar{\mathbf{v}}_{\mathcal{I}^c}^T \mathbf{w}_{\mathcal{I}^c} \\
&\leq \|\mathbf{w}_{\mathcal{I}}\|_2 \|\bar{\mathbf{v}}_{\mathcal{I}}\|_2 + \bar{\mathbf{v}}_{\mathcal{I}^c}^T \mathbf{w}_{\mathcal{I}^c} \\
&\leq \|\bar{\mathbf{v}}_{\mathcal{I}}\|_2.
\end{aligned} \tag{7.33}$$

It is readily verified that the equality holds if $\mathbf{w}$ is in the form of $-\text{sgn}(\mathbf{z}) \odot \mathbf{e}^*$.

2) $\mathcal{I}$ is empty, i.e., $\bar{\mathbf{v}} \preceq \mathbf{0}$. We prove the following result: $f(\mathbf{w}) \leq \|\mathbf{w}\|_2 \max(\bar{\mathbf{v}}) = \max(\bar{\mathbf{v}})$. It is obvious that the equality holds if $\mathbf{w}$ is in the form of $-\text{sgn}(\mathbf{z}) \odot \mathbf{e}^*$.

The case of $M = 1$ is trivial. We next prove the case of $M = 2$ and then use induction to complete the proof. When $M = 2$, substitute $w_1 = \sqrt{1 - w_2^2}$ into

Figure 7.2: An illustration of Lemma 7.1 with nonnegative entries of $\mathbf{e}$. The unit circle in the first quadrant composes of $\mathcal{D}_{\mathbf{e}}$. Projections of four possible $\mathbf{v}$'s are shown.

$f(\mathbf{w})$ and then

$$g(w_2) := f\left(\sqrt{1-w_2^2}, w_2\right) = \sqrt{1-w_2^2}\bar{v}_1 + w_2\bar{v}_2. \quad (7.34)$$

It is easy to show that $g'(w_2) \leq 0$ if $0 \leq w_2 \leq \frac{|\bar{v}_2|}{\|\bar{\mathbf{v}}\|_2}$, and $g'(w_2) \geq 0$ if $\frac{|\bar{v}_2|}{\|\bar{\mathbf{v}}\|_2} \leq w_2 < 1$. So the maximum of $g(w_2)$ can only be obtained at the boundary of the interval $[0, 1]$, i.e., $f(\mathbf{w}) \leq \max(g(0), g(1)) = \max(\bar{\mathbf{v}})$.

Suppose the lemma holds when $M = n - 1$ with $n > 3$. We next show that it holds when $M = n$. Denote $\mathbf{w}_{-1} = [w_2, \dots, w_n]^T$, $\bar{\mathbf{v}}_{-1} = [\bar{v}_2, \dots, \bar{v}_n]^T$. By $w_1^2 + \|\mathbf{w}_{-1}\|_2^2 = \|\mathbf{w}\|_2^2$ and applying the results when $M = n - 1$ and $M = 2$ consecutively,

$$\begin{aligned} f(\mathbf{w}) &= \bar{v}_1 w_1 + \bar{\mathbf{v}}_{-1}^T \mathbf{w}_{-1} \\ &\leq \bar{v}_1 w_1 + \max(\bar{\mathbf{v}}_{-1}) \|\mathbf{w}_{-1}\|_2 \\ &\leq \|\mathbf{w}\|_2 \max(\bar{\mathbf{v}}) = \max(\bar{\mathbf{v}}). \end{aligned} \quad (7.35)$$

■

Lemma 7.1 tells how to calculate the projection onto the nonconvex set $\mathcal{D}_e$ defined in (7.7). An illustration of Lemma 7.1 is presented in Fig. 7.2, where we consider the two dimensional case with both entries of $\mathbf{e}$ nonnegative. The unit circle in the first quadrant composes of $\mathcal{D}_e$. Projections of four possible $\mathbf{v}$'s are shown. The resulting algorithm is summarized in Algorithm 7.1, named as variational message passing with quantization (Q-VMP).

Algorithm 7.1: Q-VMP

Input: sensing matrix $\mathbf{A}$, quantized measurement $\mathbf{z}$, domain of quantization error $\mathcal{D}_e$, and noise variance σ^2 .

1. initialize $\langle \alpha_n^{-1} \rangle_{q(\alpha)}$, $n = 1, \dots, N$, $\langle \eta \rangle_{q(\eta)}$ and $\langle \mathbf{e} \rangle_{q(\mathbf{e})}$;
 2. **while** not converged **do**
 3. update Σ according to (7.26);
 4. update μ according to (7.25);
 5. update $\langle \alpha_n^{-1} \rangle_{q(\alpha)}$ and $\langle \alpha_n \rangle_{q(\alpha)}$, $n = 1, \dots, N$, according to (7.27);
 6. update $\langle \eta \rangle_{q(\eta)}$ according to (7.28);
 7. update $\langle \mathbf{e} \rangle_{q(\mathbf{e})}$ according to (7.30) in multi-bit CS and (7.31) in 1-bit CS, respectively;
 8. **end while**
- Output: recovered signal $\hat{\mathbf{x}} = \mu$.
-

7.2.3 The Noise Free Case

In this subsection we consider Q-VMP in the noise free case. We first consider the data consistency. A consistent recovery means that the observation can be reproduced from the recovered signal. Empirical results suggest that a consistent recovery results in less errors [59, 65]. A theoretical proof is provided in [66] on the 1-bit case. The following analysis can be applied to both multi- and 1-bit CS. Taking $\sigma^2 \rightarrow 0$ at both sides of (7.25) and (7.26) gives

$$\begin{aligned} \mu &\rightarrow \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \left(\mathbf{A} \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \right)^\dagger (\mathbf{z} - \mathbf{e}), \\ \Sigma &\rightarrow \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-1} - \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \left(\mathbf{A} \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \right)^\dagger \mathbf{A} \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-1}. \end{aligned}$$

Thus we have

$$\mathbf{A}\mu \rightarrow \mathbf{z} - \mathbf{e} \in \mathcal{D}_y,$$

i.e., $\mathcal{Q}(\mathbf{A}\boldsymbol{\mu}) \rightarrow \mathbf{z}$, which indicates that the recovered signal reproduces the observation at each iteration.

We next consider the update of $q(\mathbf{e})$ in such a case. As $\sigma^2 \rightarrow 0$ we see that $q(\mathbf{e})$ degenerates into a single-point distribution by (7.29), that coincides with the stricter assumption in (7.24) in 1-bit CS.

7.2.4 Pruning a Basis Function

The most difficult computation of Q-VMP is the calculation of $\boldsymbol{\Sigma}$ that is the inverse of an $N \times N$ matrix. Using the Woodbury matrix identity, we have

$$\boldsymbol{\Sigma} = \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1} - \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1} \mathbf{A}^T \mathbf{C}^{-1} \mathbf{A} \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1}$$

with $\mathbf{C} = \sigma^2 \mathbf{I} + \mathbf{A} \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1} \mathbf{A}^T$ being an $M \times M$ matrix. Hence, to calculate $\boldsymbol{\Sigma}$ needs $O(\min\{N^3, N^2M\})$ operations. It is noted that if Q-VMP produces some $\langle \alpha_n^{-1} \rangle_{q(\boldsymbol{\alpha})} \rightarrow +\infty$ with $n \in \{1, \dots, N\}$, then the corresponding basis $\mathbf{A}_n$ can be removed from the model. To further speed up Q-VMP, we prune a basis $\mathbf{A}_n$ from the model (to reduce N) when the corresponding parameter $\langle \alpha_n^{-1} \rangle_{q(\boldsymbol{\alpha})}$ is larger than a certain threshold $\tau_{pruning}$. Similar basis pruning approaches have been used in [39, 54].

7.3 Numerical Simulations

In this section, we study the performance of the proposed observation model and Q-VMP algorithm in comparison with existing ones by numerical simulations.

7.3.1 Experimental Setup

Model Parameters in Q-VMP: We set $\epsilon = 0$, $c = 1$ and $d = 0$ in the Gaussian-Gamma-Gamma prior as discussed in Subsection 7.1.2.

Quantizer: In multi-bit CS, a uniform unsaturated quantizer is defined in (7.2) with $L = 2^B$, equispaced $u_0, u_1, \dots, u_L$ and $v_i = (u_i + u_{i+1})/2$, $i = 1, \dots, L-1$. In addition, we let $u_L = \|\mathbf{y}\|_\infty$ and $u_0 = -\|\mathbf{y}\|_\infty$ in each trial. For a saturated quantizer, we set $u_0 = -\infty$, $u_L = +\infty$.

CS problem generation: In our experiment, we set $N = 500$, $K = 10$, and vary the bit budget (total bits of all measurements) in $\{50, 100, \dots, 1000\}$. In each trial, a K -sparse signal of length N is generated with Gaussian distributed nonzero entries and then scaled to unit norm. Entries of the sensing matrix $\mathbf{A}$ are generated independently according to a Gaussian distribution $\mathcal{N}(0, M^{-1})$. Thus the noise free measurement $\mathbf{y}^0 = \mathbf{A}\mathbf{x}$ has unit norm in expectation. To obtain a desired SNR, a white Gaussian measurement noise $\mathbf{n}$ is added with the noise variance $\sigma^2 = M^{-1}10^{-\frac{\text{SNR}}{10}}$. Then the quantized measurement $\mathbf{z} = \mathcal{Q}(\mathbf{y})$ is preserved for the following signal recovery.

Performance metrics: Three metrics are considered, including reconstruction SNR (RSNR), sparsity level of the recovered signal and computational speed. RSNR is defined as

$$\text{RSNR} = -20 \log_{10} \|\mathbf{x} - \hat{\mathbf{x}}\|_2$$

where $\hat{\mathbf{x}}$ denotes the recovered signal of $\mathbf{x}$. The sparsity level is measured by the support size of the recovered signal. The computational speed is measured by the CPU time usage. All results are averaged over 200 trials.

7.3.2 Model Efficiency

We first study the efficiency of the observation model in (7.8) introduced in this chapter for quantized CS. We consider the multi-bit CS problem with a uniform unsaturated quantizer as an example. In existing methods that account for measurement noise, e.g., in [60], the quantization error and the noise are typically coupled and treated as a Gaussian noise (only the energy information is used). Then the quantized CS problem is transformed into a conventional one. We refer to this

formulation as existing method hereafter. In this subsection we compare the signal recovery performance of the proposed formulation in (7.8) with the existing one. Naturally, we use the proposed Q-VMP algorithm for our formulation. A corresponding algorithm for the existing formulation is thus VMP introduced in [54] for conventional CS. The latter algorithm can be considered as a simplified version of Q-VMP with the quantization error $\mathbf{e}$ fixed throughout the algorithm. In addition, we consider the performance of conventional CS for comparison where the true-valued measurements are used. The conventional CS problem can be considered as one in quantized CS by using oracle information of the quantization error. So its performance acts as an upper boundary of the quantized CS problem.

In our experiment, we set SNR = 30dB and the bit depth $B = 4$ that leads to the number of quantized measurements varying from 12 to 250. Q-VMP is implemented as follows.

Q-VMP: We initialize $\langle \alpha_n^{-1} \rangle_{q(\alpha)} = 1/|\mathbf{A}_n^T \mathbf{z}|$, $n = 1, \dots, N$, $\langle \eta \rangle_{q(\eta)} = 1$ and $\langle \mathbf{e} \rangle_{q(\mathbf{e})} = \mathbf{0}$. We set $\tau_{pruning} = 10^4$. Q-VMP is terminated if $\frac{\|\tilde{\alpha}^j - \tilde{\alpha}^{j-1}\|_2}{\|\tilde{\alpha}^{j-1}\|_2} < 10^{-5}$ or the maximum number of iterations, set to 2000, is reached, where the superscript j indicates the iteration and $\tilde{\alpha} = \left[\langle \alpha_1^{-1} \rangle_{q(\alpha)}^{-1}, \dots, \langle \alpha_N^{-1} \rangle_{q(\alpha)}^{-1} \right]^T$.

The VMP algorithm for the other two cases is similarly implemented. The true noise variance is used in Q-VMP and conventional CS. For VMP with the existing formulation, we set the noise variance to $r^2/12 + \sigma^2$ where $r = u_1 - u_0$ denotes the quantization bin width. This noise variance corresponds to a Gaussian noise whose energy is comparable with that of $\mathbf{e} + \mathbf{n}$ under the assumption that $\mathbf{e}$ is uniformly distributed and independent of $\mathbf{n}$. Reconstruction SNRs of the three methods are depicted in Fig. 7.3. It can be seen that the VMP algorithm based on the proposed observation model is consistently better than that with the existing formulation though there is still a large gap between it and the oracle-aided case. So it confirms that the proposed model and framework improve the signal recovery accuracy by decoupling the quantization error from measurement noise.

Figure 7.3: Reconstruction SNRs of VMP algorithms implemented respectively based on the proposed observation model in (7.8), an existing one that couples the quantization error and measurement noise, and conventional CS (oracle-aided quantized CS) as an upper boundary.

7.3.3 Performance Comparison in Multi-bit CS

7.3.3.1 Unsaturated quantizer

In multi-bit CS, we first consider the case of a uniform unsaturated quantizer. As in the last subsection, we set $\text{SNR} = 30\text{dB}$ and $B = 4$. Besides Q-VMP, we also use BPDN [58] and L1RML [61] to recover the signal for comparison. Q-VMP is implemented as in the last subsection. The other two are implemented as follows.

BPDN: BPDN takes $\mathbf{z}$ as the measurement vector and is implemented using ℓ_1 -magic¹. We let $\epsilon = \|\mathbf{z} - \mathbf{A}\mathbf{x}\|_2$ in our implementation to achieve the best result though it is unavailable in practice.

L1RML: The problem $\min_{\tilde{\mathbf{x}}} \{\lambda \|\tilde{\mathbf{x}}\|_1 - \log f_{ML}(\mathbf{A}\tilde{\mathbf{x}})\}$ is solved, where $f_{ML}(\cdot)$ is the likelihood function of the observation with $-\log f_{ML}(\mathbf{A}\tilde{\mathbf{x}})$ being a convex function of $\tilde{\mathbf{x}}$, and $\lambda > 0$ is a regularization parameter. In general, a larger λ leads to a

¹Available at <http://users.ece.gatech.edu/~justin/l1magic>

Figure 7.4: Performance comparison of Q-VMP, BPDN and L1RML with bit depth $B = 4$. SNR = 30dB. (a) Averaged reconstruction SNR; (b) Averaged support size of recovered signal; (c) Averaged CPU time.

sparser solution. Since there are no available guidelines for the choice of λ so far, we choose λ such that the recovered signal has the optimal RSNR. Additionally, we set $\tau = \frac{\sigma^2}{\|\mathbf{A}\|_2^2}$, $\epsilon = 10^{-4}$ and $\beta = 0.5$. Readers are referred to [61] for their interpretations.

The experimental results are shown in Fig. 7.4, where red solid lines denote Q-VMP, black dashed dot lines denote BPDN, and blue dashed lines denote L1RML. Fig. 7.4(a) depicts the averaged reconstruction SNRs of the three algorithms. A significant improvement of the reconstruction SNR can be observed using the proposed Q-VMP. It is over 6dB in comparison with L1RML and about an amplitude for BPDN. Moreover, Fig. 7.4(b) shows that Q-VMP produces the sparsest solution. It is noted that L1RML can produce a sparser solution by setting a larger value of λ as in [61] but at the cost of a lower RSNR. Fig. 7.4(c) shows that the speed of Q-VMP is comparable with that of BPDN and L1RML. Implemented with the basis pruning approach, Q-VMP is faster when more measurements are acquired since it is observed in such a case that the basis pruning approach works more efficiently.

7.3.3.2 Saturated Quantizer

We next consider the case of a saturated quantizer. We adopt the same experimental setup but a saturated quantizer where a noisy measurement falls in each quantization interval with the same probability. Since both the sensing matrix and measurement noise are Gaussian in the experiment, the noisy measurements are i.i.d. Gaussian $\mathcal{N}(0, M^{-1} + \sigma^2)$. Then it is easy to get the quantizer. As a result, 12.5% of the measurements are saturated in expectation. BPDN is inappropriate in such a case. We compare Q-VMP only with L1RML. The averaged reconstruction SNRs of Q-VMP and L1RML are presented in Fig. 7.5 (red solid lines). Q-VMP obtains a RSNR of about 10dB higher than L1RML when sufficient measurements are acquired. The performance of the two algorithms on support size and speed is similar to that in the uniform quantizer case and is omitted.

The experiment above may shed light on the optimal quantizer design for Q-VMP. By comparing the performance of Q-VMP in the two quantizer scenarios, it can be

Figure 7.5: Reconstruction SNRs of Q-VMP and L1RML with a saturated quantizer, as well as those with the unsaturated quantizer in Fig. 7.4.

seen from Fig. 7.5 that the saturated quantizer outperforms the uniform unsaturated one when more measurements are taken for Q-VMP while it is not so clear for L1RML. We pose the problem of the optimal quantizer design for Q-VMP as a future work.

7.3.4 Performance Comparison in 1-bit CS

The bit-depth $B = 1$ in 1-bit CS. We set $\text{SNR} = 10\text{dB}$. In such a case, 9.75% measurements flip their signs due to the noise in expectation. We compare Q-VMP with the state-of-the-art algorithms BIHT [66] and the convex programming approach in [68], denoted by CVXP. The three algorithms are implemented as follows.

Q-VMP: We initialize $\langle \alpha_n^{-1} \rangle_{q(\alpha)} = \sqrt{M} / |\mathbf{A}_n^T \text{sgn}(\mathbf{z})|$, $n = 1, \dots, N$, $\langle \eta \rangle_{q(\eta)} = 1$ and $\langle \mathbf{e} \rangle_{q(\mathbf{e})} = -\text{sgn}(\mathbf{z}) / \sqrt{M}$. As addressed in Remark 7.4, the effective noise level in 1-bit CS is much lower than the true one. We empirically find that it is a good choice to set the noise variance in Q-VMP to $10^{-3}\sigma^2$ (Q-VMP is slow in the case of a very small noise variance, which is the reason why we consider a lower SNR

Figure 7.6: Performance comparison of Q-VMP, BIHT and CVXP in 1-bit CS. SNR = 10dB. (a) Averaged reconstruction SNR; (b) Averaged support size of recovered signal; (c) Averaged CPU time.

in 1-bit CS). We set $\tau_{pruning} = 10^4$ and terminate Q-VMP as in multi-bit CS. The recovered signal is finally scaled to unit norm for comparison with the original one.

BIHT: The oracle information of K is used in the hard thresholding operation, i.e., BIHT is certain to return a reconstruction with K nonzero entries. BIHT is terminated if the Hamming error of the current recovery is below the expected Hamming error or the maximum number of iterations, set to 1000, is reached. Readers are referred to [66] for the definition of the Hamming error.

CVXP: The oracle information of $\|\mathbf{x}\|_1$ is used to achieve the best result. We implement it using CVX [124] (its speed can be accelerated in the future using more sophisticated solvers).

Our experimental results are presented in Fig. 7.6. In all figures, red solid lines denote Q-VMP, black dashed dot lines denote BIHT, and blue dashed lines denote CVXP. It is shown in Fig. 7.6(a) that the proposed Q-VMP outperforms consistently the other two algorithms in the recovery accuracy. From Fig. 7.6(b), it can be seen that Q-VMP produces a sparser solution than CVXP while BIHT uses this oracle information. Fig. 7.6(c) shows that the computational speed is a disadvantage of Q-VMP.

7.4 Conclusion

The problem of sparse signal recovery from noisy quantized compressive measurements was studied in this chapter. A Bayesian framework was presented that unifies the multi- and 1-bit CS problems and can be applied to various scenarios including noiseless/noisy environment and unsaturated/saturated quantizer. An algorithm was proposed based on variational Bayesian inference under the proposed framework. The quantization error is decoupled from the measurement noise, modeled as a random variable and estimated jointly with the signal of interest, leading to improved performance of the algorithm. Numerical simulations were provided to demonstrate its improved signal recovery accuracy over the existing results.

Part II

Applications

Chapter 8

Off-Grid Direction of Arrival Estimation

Chapters 3–7 have been focused on the theoretical analysis and algorithm design of CS. In this chapter and the next the applications of CS to direction of arrival (DOA) estimation and magnetic resonance imaging (MRI) will be studied, respectively. The DOA estimation is a classical problem in array signal processing [84] with many practical applications. Conventional covariance-based approaches, e.g., MUSIC [86, 87], require a large number of snapshots for its success. Its research has recently been advanced owing to the development of methods based on sparse signal representation or CS [88, 89, 139, 140]. While these methods have shown advantages over conventional ones, there are still difficulties in practical situations where the true DOAs are not on the discretized sampling grid. The resulting DOA estimation problem is referred to as off-grid DOA estimation in this thesis. To deal with the problem, this chapter studies an off-grid model which takes into account effects of the off-grid DOAs and has a smaller modeling error. Based on the model we introduce two new approaches to the problem solving. One is based on the results of Chapter 5. We show that the DOA estimation problem can be transformed into the perturbed CS problem studied in Chapter 5 under the off-grid model. Then the ℓ_1 optimization approach presented in Chapter 5 can be applied straightforward. But one shortcoming of this optimization approach is the high computational

complexity. The other has reduced computational workload and adopts SBL to do the estimation in an iterative manner. It can be applied to both single snapshot and multi-snapshot cases. Numerical simulations show that both the proposed approaches have improved accuracy and can maintain high estimation accuracy even under a very coarse sampling grid. This chapter is mainly based on [7] and [11].

8.1 Problem Description and Off-Grid Model

8.1.1 Problem Description

Consider K narrowband far-field sources $s_k(t)$, $k = 1, \dots, K$, impinging on an array of M omnidirectional sensors from directions d_k , $k = 1, \dots, K$. According to [84], the time delays at different sensors can be represented by simple phase shifts, leading to the observation model:

$$\mathbf{y}(t) = \mathbf{A}(\mathbf{d})\mathbf{s}(t) + \mathbf{e}(t), \quad t = 1, \dots, T, \quad (8.1)$$

where $\mathbf{y}(t) = [y_1(t), \dots, y_M(t)]^T \in \mathbb{C}^M$, $\mathbf{d} = [d_1, \dots, d_K]^T \in [0^\circ, 180^\circ]^K$, $\mathbf{s}(t) = [s_1(t), \dots, s_K(t)]^T \in \mathbb{C}^K$, $\mathbf{e}(t) = [e_1(t), \dots, e_M(t)]^T \in \mathbb{C}^M$, and $y_m(t)$ and $e_m(t)$, $m = 1, \dots, M$, denote the output and measurement noise of the m th sensor at time t respectively. The matrix $\mathbf{A}(\mathbf{d}) = [\mathbf{a}(d_1), \dots, \mathbf{a}(d_K)] \in \mathbb{C}^{M \times K}$ is an array manifold matrix and $\mathbf{a}(d_k) \in \mathbb{C}^M$ is called steering vector of the k th source which is determined by the geometry of the sensor array and is known. The entry $\mathbf{a}_m(d_k)$ contains the delay information of the k th source to the m th sensor. Note that the number of sources K may be unknown in practice. So, the objective of the DOA estimation is to find the unknown DOA vector $\mathbf{d}$ given the sensor measurements $\mathbf{y}(t)$ and the mapping $\mathbf{d} \rightarrow \mathbf{A}(\mathbf{d})$.

8.1.2 Off-Grid DOA Estimation Model

Since the DOAs of interest lie in a continuous range and CS is applicable to recovery of discrete signals, discretization of the DOA range is inevitable for applying CS. In [88, 89] a fixed sampling grid is selected which serves as the set of all candidates of DOA estimates. By assuming that all the true (unknown) DOAs are exactly on the selected grid, a CS problem can be formulated where the DOAs of interest constitute the support of the sparse signal to be recovered.

In practical situations, the true DOAs may not be on the selected grid. Using linear approximation an off-grid model was firstly proposed in [78] to account for the general case of off-grid DOA. Let $\tilde{\mathbf{d}} = \{\tilde{d}_1, \dots, \tilde{d}_N\}$ be a fixed sampling grid in the DOA range $[0, \pi)$, where N denotes the grid number and typically satisfies $N \gg M > K$. Suppose $d_k \notin \{\tilde{d}_1, \dots, \tilde{d}_N\}$ for some $k \in \{1, \dots, K\}$ and that $\tilde{d}_{n_k}$, $n_k \in \{1, \dots, N\}$, is the nearest grid point to d_k . We approximate the steering vector $\mathbf{a}(d_k)$ using linearization:

$$\mathbf{a}(d_k) \approx \mathbf{a}(\tilde{d}_{n_k}) + \mathbf{b}(\tilde{d}_{n_k})(d_k - \tilde{d}_{n_k}) \quad (8.2)$$

with $\mathbf{b}(\tilde{d}_{n_k}) = \mathbf{a}'(\tilde{d}_{n_k})$. Denote $\mathbf{A} = [\mathbf{a}(\tilde{d}_1), \dots, \mathbf{a}(\tilde{d}_N)]$, $\mathbf{B} = [\mathbf{b}(\tilde{d}_1), \dots, \mathbf{b}(\tilde{d}_N)]$ and $\Phi(\boldsymbol{\beta}) = \mathbf{A} + \mathbf{B}\text{diag}(\boldsymbol{\beta})$, where $\boldsymbol{\beta} = [\beta_1, \dots, \beta_N]^T$, for $n = 1, \dots, N$,

$$\begin{aligned} \beta_n &= d_k - \tilde{d}_{n_k}, \quad x_n(t) = s_{n_k}(t), \quad \text{if } n = n_k \text{ for any } k \in \{1, \dots, K\}; \\ \beta_n &= 0, \quad x_n(t) = 0, \quad \text{otherwise,} \end{aligned} \quad (8.3)$$

with $n_k \in \{1, \dots, N\}$ and $\tilde{d}_{n_k}$ being the nearest grid to a source d_k , $k \in \{1, \dots, K\}$. By absorbing the approximation error into the measurement noise the observation model in (8.1) can be written into

$$\mathbf{y}(t) = \Phi(\boldsymbol{\beta}) \mathbf{x}(t) + \mathbf{e}(t), \quad t = 1, \dots, T, \quad (8.4)$$

which is the off-grid model to be used in this chapter.

It should be noted that the off-grid model in (8.4) is closely related to the on-grid one that can be obtained by setting $\boldsymbol{\beta} = \mathbf{0}$ in (8.4) ($\Phi(\mathbf{0}) = \mathbf{A}$). In fact, the off-grid model can be considered as the first order approximation of the true observation model while the on-grid one is the zeroth order approximation. As a result, the off-grid model has a much smaller modeling error than the on-grid one. Such an advantage is twofold. First, by adopting the same sampling grid the off-grid model results in higher accuracy, especially in the high SNR regime where the modeling error is the dominant modeling uncertainty. Second, a coarser sampling grid can be adopted in the off-grid model to achieve a considerably reduced computational workload with a comparable modeling accuracy.

To estimate the DOA vector $\mathbf{d}$ we need to find the support of the sparse signals $\mathbf{x}(t)$, $t = 1, \dots, T$, as well as the off-grid difference $\boldsymbol{\beta}$.

8.2 Off-Grid DOA Estimation Using P-BPDN

8.2.1 Off-Grid DOA Estimation Under SP-CS

In this section we show that the off-grid DOA estimation problem can be formulated into a CS problem subject to structured matrix perturbation which has been studied in Chapter 5 and coined as SP-CS. For simplicity, we consider the single-snapshot and noise-free case with a uniform linear array (ULA). Correspondingly, we have $a_m(d_k) = \exp\{j\pi(m - \frac{M+1}{2})\cos d_k\}$, $k = 1, \dots, K$, $m = 1, \dots, M$, where $j = \sqrt{-1}$. Denote $\boldsymbol{\theta} = [\cos d_1, \dots, \cos d_K]^T \in (-1, 1]^K$ and the corresponding steering vector $\tilde{\mathbf{a}}(\theta_k) = \mathbf{a}(d_k)$. We consider a mathematically equivalent problem of the estimation of $\boldsymbol{\theta}$ rather than that of $\mathbf{d}$ in this section following from [7]. Without loss of generality, let $\tilde{\boldsymbol{\theta}} = \{\frac{1}{N} - 1, \frac{3}{N} - 1, \dots, 1 - \frac{1}{N}\}$ be a uniform sampling grid in the continuous range $(-1, 1]$ of $\boldsymbol{\theta}$ (assume N is an even number). In existing standard CS based methods $\tilde{\boldsymbol{\theta}}$ actually serves as the set of all candidate DOA estimates. As a result, their estimation accuracy is limited by the grid density since for some θ_j ,

$j \in \{1, \dots, K\}$, the best estimate of θ_j is its nearest grid point in $\tilde{\boldsymbol{\theta}}$. It can be easily shown that a lower bound for the mean squared estimation error of each θ_j is $\text{LB} = \frac{1}{3N^2}$ by assuming that θ_j is uniformly distributed in one or more grid intervals.

In the following we transform the off-grid model into the one studied in Chapter 5. Similar to (8.2), consider the Taylor expansion

$$\tilde{\mathbf{a}}(\theta_j) = \tilde{\mathbf{a}}(\tilde{\theta}_{l_j}) + \tilde{\mathbf{b}}(\tilde{\theta}_{l_j}) (\theta_j - \tilde{\theta}_{l_j}) + \mathbf{R}_j \quad (8.5)$$

with $\tilde{\mathbf{b}}(\tilde{\theta}_{l_j}) = \tilde{\mathbf{a}}'(\tilde{\theta}_{l_j})$ and $\mathbf{R}_j$ being a remainder term with respect to θ_j . Denote $\kappa = \frac{\pi}{2} \sqrt{\frac{M^2-1}{3}}$, $\tilde{\mathbf{A}} = \frac{1}{\sqrt{M}} [\tilde{\mathbf{a}}(\tilde{\theta}_1), \dots, \tilde{\mathbf{a}}(\tilde{\theta}_N)]$, $\tilde{\mathbf{B}} = \frac{1}{\kappa\sqrt{M}} [\tilde{\mathbf{b}}(\tilde{\theta}_1), \dots, \tilde{\mathbf{b}}(\tilde{\theta}_N)]$, and for $l = 1, \dots, N$,

$$\begin{aligned} \beta_l^o &= \kappa (\theta_j - \tilde{\theta}_{l_j}), \quad x_l^o = s_j, \text{ if } l = l_j \text{ for any } j \in \{1, \dots, K\}; \\ \beta_l^o &= 0, \quad x_l^o = 0, \text{ otherwise,} \end{aligned}$$

with $l_j \in \{1, \dots, N\}$ and $\tilde{\theta}_{l_j}$ being the nearest grid to a source θ_j , $j \in \{1, \dots, K\}$. It is easy to show that $\|\mathbf{x}^o\|_0 \leq K$, each column of $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ has unit norm, and $\boldsymbol{\beta}^o \in [-r, r]^N$ with $r = \frac{\kappa}{N} = \frac{\pi}{2N} \sqrt{\frac{M^2-1}{3}}$. In addition, we let $\mathbf{e} = \mathbf{R}\mathbf{s}$ with $\mathbf{R} = [\mathbf{R}_1, \dots, \mathbf{R}_K]$. We set $\epsilon = \frac{\sqrt{K}\|\mathbf{s}\|_2\pi^2}{8N^2} \sqrt{\frac{3M^4-10M^2+7}{15}}$ such that $\|\mathbf{e}\|_2 \leq \epsilon$ (the information of K and $\|\mathbf{s}\|_2$ is used). The reason is as follows. By (8.5), we have for $l = 1, \dots, M$, $j = 1, \dots, K$,

$$R_{lj} = \frac{A''_{lj}(\xi)}{2} (\theta_j - \tilde{\theta}_{l_j})^2, \quad (8.6)$$

where ξ is between θ_j and $\tilde{\theta}_{l_j}$, $A''_{lj}(\xi) = -\frac{\pi^2}{\sqrt{M}} (l - \frac{M+1}{2})^2 \exp\{i\pi(l - \frac{M+1}{2})\xi\}$, and $|\theta_j - \tilde{\theta}_{l_j}| \leq \frac{1}{N}$. It follows that for $j = 1, \dots, K$,

$$\|\mathbf{R}_j\|_2 \leq \frac{1}{2} \max \|A''_j\|_2 \cdot \frac{1}{N^2} = \frac{\pi^2}{8N^2} \sqrt{\frac{3M^4-10M^2+7}{15}}.$$

Finally, it gives the expression of ϵ by observing that

$$\|\mathbf{e}\|_2 = \|\mathbf{R}\mathbf{s}\|_2 \leq \|\mathbf{R}\|_F \|\mathbf{s}\|_2 \leq \sqrt{K} \|\mathbf{R}_1\|_2 \|\mathbf{s}\|_2.$$

After that, the DOA estimation model can be written into the form of our studied model in (5.9) of Chapter 5. The only differences are that $\tilde{\mathbf{A}}$, $\tilde{\mathbf{B}}$, $\mathbf{x}^o$ and $\mathbf{e}$ are in the complex domain rather than the real domain and that $\mathbf{e}$ denotes a modeling error term rather than the measurement noise. It is noted that the robust stability results in SP-CS apply straightforward to such complex signal case with few modifications. The objective becomes the recovery of $\mathbf{x}^o$ (its support actually) and β^o . According to Theorem 5.4 $\mathbf{x}^o$ and β^o can be stably recovered if the D-RIP condition is satisfied. Denote $\hat{\mathbf{x}}$ the recovered $\mathbf{x}^o$, $\hat{\beta}$ the recovered β^o , and $\mathcal{I}$ the support of $\hat{\mathbf{x}}$. Then we obtain the recovered θ : $\hat{\theta} = \tilde{\theta}_{\mathcal{I}} + \kappa^{-1} \hat{\beta}_{\mathcal{I}}$, where $\mathbf{v}_{\mathcal{I}}$ keeps only entries of a vector $\mathbf{v}$ on the index set $\mathcal{I}$. The empirical results in Subsection 8.2.2 will illustrate the merits of applying the SP-CS framework to estimate DOAs.

8.2.2 Numerical Simulations

This subsection studies the empirical performance of the application of the SP-CS framework in the DOA estimation. We consider the case of $N = 90$ and $K = 2$. Numerical calculations show that the D-RIP condition $\bar{\delta}_{4K}(\Psi) < \left(\sqrt{2(1+r^2)} + 1\right)^{-1}$ in Theorem 5.4 is satisfied if $M \geq 145$. Though it ceases to be a “compressed” sensing problem when $M \geq N$, it still makes sense in SP-CS since there are $2N$ variables to be estimated and hence the P-BPDN problem is still underdetermined as $M < 2N$. As noted in Subsection 5.2.3, the D-RIP condition can be possibly relaxed using recent techniques in standard CS, which may reduce the required M value. In addition, the RIP condition is a sufficient condition for guaranteed signal recovery accuracy while its conservativeness in standard CS has been studied in [30]. We next choose a much smaller $M = 30$ ($r \approx 0.302$ in such a case) and show the empirical performance of the proposed SP-CS framework on such off-grid DOA estimation.

The experimental setup is as follows. In each trial, the complex source signal $\mathbf{s}$ is generated with both entries having unit amplitude and random phases. θ_1 and θ_2 are generated uniformly from intervals $[\frac{2}{N}, \frac{4}{N}]$ and $[\frac{12}{N}, \frac{14}{N}]$ respectively ($5.1^\circ \sim 7.7^\circ$ apart in the DOA domain). The P-BPDN problem in (5.11) is solved using AA-P-BPDN presented in Subsection 5.3.2 whose settings are the same as those in Section 5.4. Our experimental results of the estimation error $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}$ for both sources are presented in Fig. 8.1 where 1000 trials are used. It can be seen that P-BPDN performs well on the off-grid DOA estimation. All estimation errors lie in the interval $[-\frac{1}{N}, \frac{1}{N}]$ with most very close to zero. To achieve a possibly comparable mean squared estimation error, a grid of length at least $N = 360$ has to be used in standard CS based methods according to the lower bound mentioned in Subsection 8.2.1. An example of performance of SP-CS and standard CS on DOA estimation is shown in Fig. 8.2, where the two approaches share the same data set and $N = 360$ is set in standard CS. From the upper two sub-figures, it can be seen that SP-CS performs well on both source signal and β^o recoveries. From the lower left one, however, it can be seen that two nonzero entries are presented in the recovered signal around the location of each source when using standard CS. Such a phenomenon is much clearer in the last sub-figure, where it can be observed that a single peak exhibits at a place very close to the true location of source 1 using the proposed SP-CS framework while two peaks occur at places further away from the true source in standard CS.

8.3 Off-Grid DOA Estimation Using SBL

The ℓ_1 optimization-based off-grid DOA estimation approach within the framework of structured CS presented in the last section has strong theoretical motivation. However, the AA-P-BPDN algorithm is computationally inefficient since a series of BPDN problems need to be solved iteratively. In this section, we propose a computationally efficient algorithm to solve the off-grid DOA estimation problem under the framework of SBL, where the sparse source signal and the grid offset are

Figure 8.1: Histogram of θ estimation error for both sources using P-BPDN for SP-CS. Statistics including mean, variance and mean squared error (MSE) are shown.

Figure 8.2: Performance comparison of SP-CS and standard CS (SCS) on DOA estimation. Upper left: signal recovery in SP-CS; upper right: β^o recovery in SP-CS (shown only on the signal support); lower left: signal recovery in standard CS; lower right: signal amplitude versus θ (near the location of source 1) in SP-CS and standard CS.

modeled as random variables and jointly estimated in an iterative manner. This section is mainly based on the work in [11]. For convenience, we follow [11] and consider the parameter estimation in the DOA domain and hence the observation model in (8.4) is adopted. As a result, the notations adopted here may have slightly different meanings from those in the last section. In addition, we assume that the number of sources K is *a priori* known. In fact, the presented approach can still be applied when K is unknown. In the latter case, we may let $K = K_{max}$ where K_{max} denotes the maximum number of sources detectable using the sensor array adopted. For example, $K_{max} = M - 1$ if an M -element ULA is adopted.

Without loss of generality, let $\tilde{\boldsymbol{\theta}} \in [0^\circ, 180^\circ)^N$ be a uniform grid in the continuous DOA range $[0^\circ, 180^\circ)$ with a grid interval $r = \tilde{\theta}_2 - \tilde{\theta}_1 \propto N^{-1}$. We derive our algorithm in the multiple measurement vector (MMV) case. The single measurement vector (SMV) scenario is a special case by simply setting $T = 1$. Denote $\mathbf{Y} = [\mathbf{y}(1), \dots, \mathbf{y}(T)] \in \mathbb{C}^{M \times T}$, $\mathbf{X} = [\mathbf{x}(1), \dots, \mathbf{x}(T)] \in \mathbb{C}^{N \times T}$ and $\mathbf{E} = [\mathbf{e}(1), \dots, \mathbf{e}(T)] \in \mathbb{C}^{M \times T}$. The off-grid DOA estimation model in (8.4) becomes

$$\mathbf{Y} = \Phi(\boldsymbol{\beta}) \mathbf{X} + \mathbf{E} \quad (8.7)$$

with $\Phi(\boldsymbol{\beta}) = \mathbf{A} + \mathbf{B} \text{diag}(\boldsymbol{\beta}) \in \mathbb{C}^{M \times N}$ and $\boldsymbol{\beta} \in [-\frac{1}{2}r, \frac{1}{2}r]^N$. The matrix $\mathbf{X}$ of interest is jointly sparse (or row-sparse), i.e., all columns of $\mathbf{X}$ are sparse and share the same support.

8.3.1 Sparse Bayesian Formulation

8.3.1.1 Noise model

Under an assumption of white (circular symmetric) complex Gaussian [141] noises, we have

$$p(\mathbf{E}|\alpha_0) = \prod_{t=1}^T \mathcal{CN}(\mathbf{e}(t)|\mathbf{0}, \alpha_0^{-1} \mathbf{I}) \quad (8.8)$$

where $\alpha_0 = \sigma^{-2}$ denotes the noise precision with σ^2 being the noise variance, the probability density function (PDF) of a (circular symmetric) complex Gaussian distributed random variable $\mathbf{u} \sim \mathcal{CN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ is [141]

$$\mathcal{CN}(\mathbf{u}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\pi^N |\boldsymbol{\Sigma}|} \exp \left\{ -(\mathbf{u} - \boldsymbol{\mu})^H \boldsymbol{\Sigma}^{-1} (\mathbf{u} - \boldsymbol{\mu}) \right\}. \quad (8.9)$$

Then we have

$$p(\mathbf{Y}|\mathbf{X}, \alpha_0, \boldsymbol{\beta}) = \prod_{t=1}^T \mathcal{CN}(\mathbf{y}(t)|\boldsymbol{\Phi}(\boldsymbol{\beta})\mathbf{x}(t), \alpha_0^{-1}\mathbf{I}). \quad (8.10)$$

In this chapter we assume that the noise precision α_0 is unknown. A Gamma hyperprior is assumed for α_0 since it is a conjugate prior of the Gaussian distribution:

$$p(\alpha_0; c, d) = \Gamma(\alpha_0|c, d) \quad (8.11)$$

where $\Gamma(\alpha_0|c, d) = [\Gamma(c)]^{-1} d^c \alpha_0^{c-1} \exp\{-d\alpha_0\}$ with $\Gamma(\cdot)$ being the Gamma function. We set $c, d \rightarrow 0$ as in [38, 39] to obtain a broad hyperprior.

Remark 8.1. *The noise term in fact contains two components: the measurement noise and the modeling error. In the low and moderate SNR regimes which are of main interest the measurement noise is the dominant uncertainty and thus the modeling error can be neglected.*

8.3.1.2 Sparse signal model

A sparse prior is needed for the jointly sparse matrix $\mathbf{X}$ of interest. We assume that the signals among snapshots are independent and adopt the following two-stage hierarchical prior: $p(\mathbf{X}; \rho) = \int p(\mathbf{X}|\boldsymbol{\alpha}) p(\boldsymbol{\alpha}; \rho) d\boldsymbol{\alpha}$, where $\rho > 0$, $\boldsymbol{\alpha} \in \mathbb{R}^N$, $\boldsymbol{\Lambda} = \text{diag}(\boldsymbol{\alpha})$ and

$$p(\mathbf{X}|\boldsymbol{\alpha}) = \prod_{t=1}^T \mathcal{CN}(\mathbf{x}(t)|\mathbf{0}, \boldsymbol{\Lambda}), \quad (8.12)$$

$$p(\boldsymbol{\alpha}; \rho) = \prod_{n=1}^N \Gamma(\alpha_n | 1, \rho). \quad (8.13)$$

It is easy to show that all columns of $\mathbf{X}$ are independent and share the same prior. According to [46], for $t = 1, \dots, T$ both $\Re\{\mathbf{x}(t)\}$ and $\Im\{\mathbf{x}(t)\}$ are identically Laplace distributed and strongly peaked at the origin. As a result, the two-stage hierarchical prior is a sparse prior that favors most rows of $\mathbf{X}$ being zero.

8.3.1.3 Off-grid distance model

We assume a uniform prior for $\boldsymbol{\beta}$:

$$\boldsymbol{\beta} \sim U\left(\left[-\frac{1}{2}r, \frac{1}{2}r\right]^N\right). \quad (8.14)$$

The prior is noninformative in the sense that the only information of $\boldsymbol{\beta}$ we use is its boundedness.

By combining the stages of the hierarchical Bayesian model, the joint PDF is

$$p(\mathbf{X}, \mathbf{Y}, \alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta}) = p(\mathbf{Y} | \mathbf{X}, \alpha_0, \boldsymbol{\beta}) p(\mathbf{X} | \boldsymbol{\alpha}) p(\boldsymbol{\alpha}) p(\alpha_0) p(\boldsymbol{\beta}) \quad (8.15)$$

with the distributions on the right hand side as defined by (8.10), (8.12), (8.13), (8.11) and (8.14) respectively.

8.3.2 Bayesian Inference

An evidence procedure [55] is exploited to perform the Bayesian inference since the exact posterior distribution $p(\mathbf{X}, \alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{Y})$ cannot be explicitly calculated. Similar approaches have been used in standard Bayesian CS methods [38, 46]. First it is easy to show that the posterior distribution of $\mathbf{X}$ is a complex Gaussian distribution:

$$p(\mathbf{X} | \mathbf{Y}, \alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \prod_{t=1}^T \mathcal{CN}(\mathbf{x}(t) | \boldsymbol{\mu}(t), \boldsymbol{\Sigma}) \quad (8.16)$$

with

$$\boldsymbol{\mu}(t) = \alpha_0 \boldsymbol{\Sigma} \boldsymbol{\Phi}^H \mathbf{y}(t), \quad t = 1, \dots, T, \quad (8.17)$$

$$\boldsymbol{\Sigma} = (\alpha_0 \boldsymbol{\Phi}^H \boldsymbol{\Phi} + \boldsymbol{\Lambda}^{-1})^{-1}. \quad (8.18)$$

Calculations of $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}(t)$, $t = 1, \dots, T$, need estimates of the hyperparameters α_0 , $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. In an evidence procedure, they are estimated using an MAP estimate that maximizes $p(\alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{Y})$. It can be easily observed that to maximize $p(\alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{Y})$ is equivalent to maximizing the joint PDF $p(\mathbf{Y}, \alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta}) = p(\alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta} | \mathbf{Y}) p(\mathbf{Y})$ since $p(\mathbf{Y})$ is independent of the hyperparameters. An expectation-maximization (EM) algorithm is implemented that treats $\mathbf{X}$ as a hidden variable and turns to maximizing $E \{ \log p(\mathbf{X}, \mathbf{Y}, \alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta}) \}$, where $p(\mathbf{X}, \mathbf{Y}, \alpha_0, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is given in (8.15) and $E \{ \cdot \}$ denotes an expectation with respect to the posterior of $\mathbf{X}$ as given in (8.16) using the current estimates of the hyperparameters.

Denote $\boldsymbol{\mathcal{U}} = [\boldsymbol{\mu}(1), \dots, \boldsymbol{\mu}(T)] = \alpha_0 \boldsymbol{\Sigma} \boldsymbol{\Phi}^H \mathbf{Y}$, $\underline{\mathbf{X}} = \mathbf{X} / \sqrt{T}$, $\underline{\mathbf{Y}} = \mathbf{Y} / \sqrt{T}$, $\underline{\boldsymbol{\mathcal{U}}} = \boldsymbol{\mathcal{U}} / \sqrt{T}$ and $\underline{\rho} = \rho / T$. Following a similar procedure as in [39], it is easy to obtain the following updates of $\boldsymbol{\alpha}$ and α_0 :

$$\alpha_n^{new} = \frac{\sqrt{1 + 4\underline{\rho} E \{ \|\underline{\mathbf{X}}^n\|_2^2 \}} - 1}{2\underline{\rho}}, \quad n = 1, \dots, N, \quad (8.19)$$

$$\alpha_0^{new} = \frac{M + (c - 1)/T}{E \{ \|\underline{\mathbf{Y}} - \boldsymbol{\Phi} \underline{\mathbf{X}}\|_F^2 \} + d/T}, \quad (8.20)$$

where $E \{ \|\underline{\mathbf{X}}^n\|_2^2 \} = \|\underline{\boldsymbol{\mathcal{U}}}^n\|_2^2 + \Sigma_{nn}$, $E \{ \|\underline{\mathbf{Y}} - \boldsymbol{\Phi} \underline{\mathbf{X}}\|_F^2 \} = \|\underline{\mathbf{Y}} - \boldsymbol{\Phi} \underline{\boldsymbol{\mathcal{U}}}\|_F^2 + \alpha_0^{-1} \sum_{n=1}^N \gamma_n$ with $\gamma_n = 1 - \alpha_n^{-1} \Sigma_{nn}$.

For $\boldsymbol{\beta}$, its estimate maximizes $E \{ \log p(\mathbf{Y} | \mathbf{X}, \alpha_0, \boldsymbol{\beta}) p(\boldsymbol{\beta}) \}$ by (8.15) and thus min-

imizes

$$\begin{aligned}
& E \left\{ \frac{1}{T} \sum_{t=1}^T \|\mathbf{y}(t) - (\mathbf{A} + \mathbf{B} \text{diag}(\boldsymbol{\beta})) \mathbf{x}(t)\|_2^2 \right\} \\
&= \frac{1}{T} \sum_{t=1}^T \|\mathbf{y}(t) - (\mathbf{A} + \mathbf{B} \text{diag}(\boldsymbol{\beta})) \boldsymbol{\mu}(t)\|_2^2 \\
&\quad + \text{Tr} \left\{ (\mathbf{A} + \mathbf{B} \text{diag}(\boldsymbol{\beta})) \boldsymbol{\Sigma} (\mathbf{A} + \mathbf{B} \text{diag}(\boldsymbol{\beta}))^H \right\} \\
&= \boldsymbol{\beta}^T \mathbf{P} \boldsymbol{\beta} - 2\mathbf{v}^T \boldsymbol{\beta} + C
\end{aligned} \tag{8.21}$$

where C is a constant term independent of $\boldsymbol{\beta}$, $\mathbf{P}$ is a positive semi-definite matrix and

$$\mathbf{P} = \Re \left\{ \overline{\mathbf{B}^H \mathbf{B}} \odot (\underline{\mathbf{U}} \cdot \underline{\mathbf{U}}^H + \boldsymbol{\Sigma}) \right\}, \tag{8.22}$$

$$\mathbf{v} = \Re \left\{ \frac{1}{T} \sum_{t=1}^T \text{diag}(\overline{\boldsymbol{\mu}(t)}) \mathbf{B}^H [\mathbf{y}(t) - \mathbf{A} \boldsymbol{\mu}(t)] - \text{diag}(\mathbf{B}^H \mathbf{A} \boldsymbol{\Sigma}) \right\}. \tag{8.23}$$

Denote $\boldsymbol{\Delta} = \text{diag}(\boldsymbol{\beta})$. The last equality of (8.21) is based on the following two equalities:

$$\begin{aligned}
& \|\mathbf{y} - (\mathbf{A} + \mathbf{B} \boldsymbol{\Delta}) \boldsymbol{\mu}\|_2^2 \\
&= \|(\mathbf{y} - \mathbf{A} \boldsymbol{\mu}) - \mathbf{B} \cdot \text{diag}(\boldsymbol{\mu}) \cdot \boldsymbol{\beta}\|_2^2 \\
&= \boldsymbol{\beta}^T \text{diag}^H(\boldsymbol{\mu}) \mathbf{B}^H \mathbf{B} \cdot \text{diag}(\boldsymbol{\mu}) \boldsymbol{\beta} \\
&\quad - 2\Re \left\{ (\mathbf{y} - \mathbf{A} \boldsymbol{\mu})^H \mathbf{B} \cdot \text{diag}(\boldsymbol{\mu}) \cdot \boldsymbol{\beta} \right\} + C_1 \\
&= \boldsymbol{\beta}^T \left(\overline{\mathbf{B}^H \mathbf{B}} \odot \boldsymbol{\mu} \boldsymbol{\mu}^H \right) \boldsymbol{\beta} \\
&\quad - 2\Re \left\{ \text{diag}(\overline{\boldsymbol{\mu}}) \mathbf{B}^H (\mathbf{y} - \mathbf{A} \boldsymbol{\mu}) \right\}^T \boldsymbol{\beta} + C_1, \\
&\text{Tr} \left\{ (\mathbf{A} + \mathbf{B} \boldsymbol{\Delta}) \boldsymbol{\Sigma} (\mathbf{A} + \mathbf{B} \boldsymbol{\Delta})^H \right\} \\
&= 2\text{Tr} \left\{ \Re \left\{ \mathbf{A} \boldsymbol{\Sigma} \boldsymbol{\Delta} \mathbf{B}^H \right\} \right\} + \text{Tr} \left\{ \mathbf{B} \boldsymbol{\Delta} \boldsymbol{\Sigma} \boldsymbol{\Delta} \mathbf{B}^H \right\} + C_2 \\
&= 2\Re \left\{ \text{Tr} \left\{ \mathbf{B}^H \mathbf{A} \boldsymbol{\Sigma} \boldsymbol{\Delta} \right\} \right\} + \text{Tr} \left\{ \boldsymbol{\Delta} \boldsymbol{\Sigma} \boldsymbol{\Delta} \mathbf{B}^H \mathbf{B} \right\} + C_2 \\
&= 2\Re \left\{ \text{diag}(\mathbf{B}^H \mathbf{A} \boldsymbol{\Sigma}) \right\}^T \boldsymbol{\beta} + \boldsymbol{\beta}^T \left(\boldsymbol{\Sigma} \odot \overline{\mathbf{B}^H \mathbf{B}} \right) \boldsymbol{\beta} + C_2,
\end{aligned}$$

where C_1, C_2 are constants independent of β , and the equality

$$\text{Tr} \{ \text{diag}^H(\mathbf{u}) \mathbf{Q} \cdot \text{diag}(\mathbf{v}) \cdot \mathbf{R}^T \} = \mathbf{u}^H (\mathbf{Q} \odot \mathbf{R}) \mathbf{v}$$

for vectors $\mathbf{u}, \mathbf{v}$ and matrices $\mathbf{Q}, \mathbf{R}$ of proper dimensions is used. Note that $\beta^T \mathbf{S} \beta \in \mathbb{R}$ for a positive semi-definite matrix $\mathbf{S}$ of proper dimension and thus $\beta^T \mathbf{S} \beta = \Re \{ \beta^T \mathbf{S} \beta \} = \beta^T \cdot \Re \mathbf{S} \cdot \beta$ since β is real-valued. Then (8.21) is obtained by observing that both $\overline{\mathbf{B}^H \mathbf{B}} \odot \boldsymbol{\mu} \boldsymbol{\mu}^H$ and $\boldsymbol{\Sigma} \odot \overline{\mathbf{B}^H \mathbf{B}}$ are positive semi-definite.

As a result, we have

$$\beta^{new} = \arg \min_{\beta \in [-\frac{1}{2}r, \frac{1}{2}r]^N} \{ \beta^T \mathbf{P} \beta - 2\mathbf{v}^T \beta \}. \quad (8.24)$$

Remark 8.2. Though an explicit expression of β^{new} cannot be given, by recognizing that β is jointly sparse with $\mathbf{x}$, the dimension of β can be reduced to K in the computation and hence β^{new} can be efficiently calculated. We provide the details in Subsection 8.3.5.

The proposed off-grid SBL algorithm, abbreviated as OGSBL, is implemented as follows. After initializations of the hyperparameters $\boldsymbol{\alpha}$, α_0 and β , we calculate $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}(t)$, $t = 1, \dots, T$, using the current values of the hyperparameters according to (8.18) and (8.17) respectively. Then we update $\boldsymbol{\alpha}$, α_0 and β according to (8.19), (8.20) and (8.24) respectively. The process is repeated until some convergence criterion is satisfied. OGSBL has guaranteed convergence since the function $p(\alpha_0, \boldsymbol{\alpha}, \beta | \mathbf{Y})$ is guaranteed to increase at each iteration by the property of EM algorithm [126].

8.3.3 OGSBL-SVD

In this subsection we recall a subspace-based idea in [88] that uses the SVD of the measurement matrix $\mathbf{Y} = \mathbf{U} \mathbf{S} \mathbf{V}^H$ to reduce the computation of the signal reconstruction process and the sensitivity to the measurement noise. Then we incorporate

it into our OGSBL algorithm. Consider the noise-free case where $\mathbf{Y} = \Phi \mathbf{X}$ with $K \leq T$. We have $\text{Rank}(\mathbf{Y}) \leq \text{Rank}(\mathbf{X}) \leq K$. Let $\mathbf{V} = [\mathbf{V}_1 \ \mathbf{V}_2]$, where $\mathbf{V}_1$ and $\mathbf{V}_2$ are matrices that consist of the first K and the rest $T - K$ columns of $\mathbf{V}$ respectively. Then we have that $\mathbf{Y}_{SV} = \mathbf{Y} \mathbf{V}_1 \in \mathbb{C}^{M \times K}$ preserves all signal information. In a general case where noises exist, by the SVD we have $\mathbf{Y} \mathbf{V} = [\mathbf{Y}_{SV} \ \mathbf{Y} \mathbf{V}_2]$, where the first part $\mathbf{Y}_{SV}$ preserves most signal information and is to be used in the following signal recovery process while the second part is abandoned. Denote $\mathbf{X}_{SV} = \mathbf{X} \mathbf{V}_1$ and $\mathbf{E}_{SV} = \mathbf{E} \mathbf{V}_1$. Then we have

$$\mathbf{Y}_{SV} = \Phi \mathbf{X}_{SV} + \mathbf{E}_{SV}. \quad (8.25)$$

In (8.25), $\mathbf{Y}_{SV}$, $\mathbf{X}_{SV}$ and $\mathbf{E}_{SV}$ can be viewed as the new matrices of sensor measurements, source signals and measurement noises respectively. The joint sparsity still holds in $\mathbf{X}_{SV}$. We do not exploit possible correlations that exist between columns of $\mathbf{X}_{SV}$ (and in $\mathbf{E}_{SV}$), i.e., we still assume that $\mathbf{X}_{SV}$ (and $\mathbf{E}_{SV}$) have independent columns.¹ It is then straightforward to apply the proposed OGSBL algorithm to estimate $\mathbf{X}_{SV}$, $\boldsymbol{\beta}$ and then the DOAs. We use OGSBL-SVD to refer to the resulting algorithm.

Based on implementation details to be introduced in Subsection 8.3.5, it can be shown that OGSBL-SVD has a computational complexity of order $O(MN^2)$ per iteration while that for OGSBL is $O(\max(MN^2, MNT))$ per iteration. An additional computational workload of order $O(\max(M^2T, MT^2))$ is for the SVD of $\mathbf{Y}$ in OGSBL-SVD. Since it is empirically found that OGSBL-SVD converges much faster than OGSBL, the whole computational workload of OGSBL-SVD is less than that of OGSBL in general.²

¹The correlations between columns of the signal matrix ($\mathbf{X}_{SV}$ in our case) have recently been studied in [49].

²A possible exception happens in the case of $T \gg N$ where the computation for the SVD is quite heavy. A modified approach in such a case is to partition $\mathbf{Y}$ firstly into blocks with each of about N columns, then operate the SVD on each block and keep the resulting signal subspaces, and finally do another SVD on the new signal matrix composed of all signal subspaces. A model similar to (8.25) can be cast.

8.3.4 Source Power and DOA Estimation

We use the estimated source powers from different directions to form a spectrum of the proposed algorithm. In the following we derive a formula to estimate the source powers. We take OGSBL-SVD as an example. The case of OGSBL is similar with some modifications and thus is omitted. Let $\widehat{\mathbf{X}} = \mathbf{X}_{SV} \mathbf{V}_1^H$ be an estimate of the signal $\mathbf{X}$. Then consider $\widehat{\mathbf{X}}$ row by row and we have $\widehat{\mathbf{X}}^n \sim \mathcal{CN}(\widehat{\mathbf{u}}^n \mathbf{V}_1^H, \widehat{\Sigma}_{nn} \mathbf{V}_1 \mathbf{V}_1^H)$ where we use $\widehat{\mathbf{u}}$ and $\widehat{\Sigma}$ to denote the final estimates of the mean and covariance of $\mathbf{X}_{SV}$ respectively. We use the expectation as an estimate of the power $\wp_n$ from direction $\tilde{\theta}_n$ (with a modification of β_n):

$$\begin{aligned}
\widehat{\wp}_n &= E\{\wp_n\} = \frac{1}{T} E\left\{\left\|\widehat{\mathbf{X}}^n\right\|_2^2\right\} \\
&= \frac{1}{T} \left(\left\|E\left\{\widehat{\mathbf{X}}^n\right\}\right\|_2^2 + E\left\{\left\|\widehat{\mathbf{X}}^n - E\left\{\widehat{\mathbf{X}}^n\right\}\right\|_2^2\right\} \right) \\
&= \frac{1}{T} \left(\left\|\widehat{\mathbf{u}}^n \mathbf{V}_1^H\right\|_2^2 + Tr\left\{\widehat{\Sigma}_{nn} \mathbf{V}_1 \mathbf{V}_1^H\right\} \right) \\
&= \frac{\left\|\widehat{\mathbf{u}}^n\right\|_2^2}{T} + \frac{K \widehat{\Sigma}_{nn}}{T}.
\end{aligned} \tag{8.26}$$

Like other spectral-based methods, the DOAs are estimated using the locations of the highest peaks of the spectrum. Suppose that the grid indices of the highest K peaks of $\widehat{\wp}$ are $\widehat{n}_k$, $k = 1, \dots, K$. The estimated K DOAs will be $\widehat{\theta}_k = \tilde{\theta}_{\widehat{n}_k} + \widehat{\beta}_{\widehat{n}_k}$, $k = 1, \dots, K$.

8.3.5 Implementation Details

This subsection presents some details of our implementations of OGSBL and OGSBL-SVD. At each iteration of OGSBL or OGSBL-SVD, an $N \times N$ matrix inversion is required when updating Σ according to (8.18). By $M < N$ the Woodbury matrix identity is applied to give $\Sigma = \Lambda - \Lambda \Phi^H \mathbf{C}^{-1} \Phi \Lambda$ with $\mathbf{C} = \alpha_0^{-1} \mathbf{I} + \Phi \Lambda \Phi^H \in \mathbb{C}^{M \times M}$.

By the fact that β is jointly sparse with $\mathbf{x}(t)$ whose K nonzero entries correspond to the locations of the K sources, we calculate only entries of β that correspond

to locations of the maximum K entries of $\boldsymbol{\alpha}$ and set others to zero. As a result, $\boldsymbol{\beta}$, $\mathbf{P}$ and $\mathbf{v}$ can be truncated into dimension of K or $K \times K$. We still use $\boldsymbol{\beta}$, $\mathbf{P}$ and $\mathbf{v}$ hereafter to denote their truncated versions for simplicity. By (8.24) and $\frac{\partial}{\partial \boldsymbol{\beta}} \{\boldsymbol{\beta}^T \mathbf{P} \boldsymbol{\beta} - 2\mathbf{v}^T \boldsymbol{\beta}\} = 2(\mathbf{P} \boldsymbol{\beta} - \mathbf{v})$ we have $\boldsymbol{\beta}^{new} = \check{\boldsymbol{\beta}}$ if $\mathbf{P}$ is invertible and $\check{\boldsymbol{\beta}} = \mathbf{P}^{-1} \mathbf{v} \in [-\frac{1}{2}r, \frac{1}{2}r]^K$. Otherwise, we update $\boldsymbol{\beta}$ elementwise, i.e., at each step we update one β_n by fixing up the other entries of $\boldsymbol{\beta}$. For $n = 1, \dots, K$, first we let

$$\check{\beta}_n = \frac{v_n - (\mathbf{P}_n)_{-n}^T \boldsymbol{\beta}_{-n}}{P_{nn}}, \quad (8.27)$$

where $\mathbf{u}_{-n}$ is $\mathbf{u}$ without the n th entry for a vector $\mathbf{u}$. Then by constraining $\beta_n \in [-\frac{1}{2}r, \frac{1}{2}r]$ we have

$$\beta_n^{new} = \begin{cases} \check{\beta}_n, & \text{if } \check{\beta}_n \in [-\frac{1}{2}r, \frac{1}{2}r]; \\ -\frac{1}{2}r, & \text{if } \check{\beta}_n < -\frac{1}{2}r; \\ \frac{1}{2}r, & \text{otherwise.} \end{cases} \quad (8.28)$$

It is easy to show that the objective function is guaranteed to decrease at each step with β_n defined in (8.28).

We terminate OGSBL and OGSBL-SVD if $\frac{\|\boldsymbol{\alpha}^{i+1} - \boldsymbol{\alpha}^i\|_2}{\|\boldsymbol{\alpha}^i\|_2} < \tau$ or the maximum number of iterations is reached, where τ is a user-defined tolerance and the superscript i refers to the iteration.

8.3.6 Numerical Simulations

In this section, we present our numerical results for the DOA estimation. A standard ULA of $M = 8$ sensors is considered. The origin is set at the middle point of the ULA to reduce the approximation error in (8.2). So, $A_{mn} = \exp\left\{j\pi\left(m - \frac{M+1}{2}\right)\cos\tilde{\theta}_n\right\}$ and $B_{mn} = -j\pi\left(m - \frac{M+1}{2}\right)\sin\tilde{\theta}_n \cdot A_{mn}$, $m = 1, \dots, M$, $n = 1, \dots, N$, with $j = \sqrt{-1}$. A uniform sampling grid $\{0^\circ, r, 2r, \dots, 180^\circ - r\}$ is considered with r being the grid interval. The number of snapshots is set to $T = 200$ in the case of MMV. We consider only OGSBL-SVD in the MMV case since it is empirically observed

to converge faster and be more accurate in comparison with OGSBL. In OGSBL-SVD, we set $\rho = 0.01$ and $c = d = 1 \times 10^{-4}$. We initialize $\alpha_0 = \frac{100K}{\sum_{t=1}^K \text{Var}\{(\mathbf{Y}_{SV})_t\}}$, $\boldsymbol{\alpha} = \frac{1}{MK} \sum_{t=1}^K |\mathbf{A}^H(\mathbf{Y}_{SV})_t|$ and $\boldsymbol{\beta} = \mathbf{0}$, where $|\cdot|$ applies elementwise. We set $\tau = 10^{-3}$ and the maximum number of iterations to 1000. We note that the proposed algorithm is insensitive to the initializations of α_0 , $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, as well as to ρ if ρ is not too large. As reported in [41], the estimate of α_0 can be inaccurate in some cases. But we have observed minimal effects on the result of DOA estimation. All experiments are carried out in Matlab v.7.7.0 on a PC with a Windows XP system and a 3GHz CPU.

8.3.6.1 Comparison with ℓ_1 -SVD

We take ℓ_1 -SVD in [88] as a representative of on-grid model based methods and compare OGSBL-SVD with it in terms of estimation bias, mean squared error (MSE) and computational time.

We first study the case where the true DOAs are exactly on the grid, i.e., the standard CS formulation describes accurately the true observation model in (8.1). Apart from OGSBL-SVD and ℓ_1 -SVD, we also consider the on-grid form of OGSBL-SVD, denoted by SBL-SVD. SBL-SVD is based on the on-grid model and hence is a special case of OGSBL-SVD by imposing $\boldsymbol{\beta} = \mathbf{0}$ during the iterations. A Monte Carlo method is used to estimate the bias of the DOA estimation. Two sources are considered with one being held fixed at direction 50° . The other source is from 54° , 55° and so on. For each pair of DOAs, the bias is calculated as the average of 100 trials. The grid interval is set to $r = 1^\circ$ in all algorithms. Fig. 8.3 presents the biases of ℓ_1 -SVD, SBL-SVD and OGSBL-SVD. Asymptotic biases exhibit in all three methods. When the two DOAs are close to each other, ℓ_1 -SVD has the largest bias while those of SBL-SVD and OGSBL-SVD are relatively small. Compared to OGSBL-SVD, SBL-SVD utilizes the additional knowledge of on-grid DOAs and has a smaller bias.

Figure 8.3: Bias comparison of ℓ_1 -SVD, SBL-SVD and OGSBL-SVD in localizing two on-grid sources as a function of angular separation with source 1 fixed at 50° and $\text{SNR} = 10$ dB. Significant bias exhibits when using ℓ_1 -SVD for closely spaced sources while small asymptotic biases occur in SBL-SVD and OGSBL-SVD.

To provide a better illustration of the DOA estimation errors, we present in Fig. 8.4 the histograms of the estimation errors of source 2 when it comes from 60° . For ℓ_1 -SVD, the errors range from -4° to 2° though it estimates the DOA accurately in over half of the 100 trials. A remarkable result is that SBL-SVD finds the DOA without error in all trials. All errors of OGSBL-SVD are within half the grid interval while it does not exploit the information that the true DOAs are on the sampling grid. The advantages of SBL-SVD and OGSBL-SVD over ℓ_1 -SVD can also be recognized from the means and variances of the estimation errors shown in Fig. 8.4.

We then consider a general case of arbitrarily located sources. Similarly, we consider two sources with one being held fixed at direction 50.3° . Fifty positions of the other source are considered, which are respectively generated uniformly from intervals $[54^\circ, 55^\circ]$, $[55^\circ, 56^\circ]$ and so on. For each pair of DOAs, the bias is calculated as the average of 100 trials. We consider OGSBL-SVD and ℓ_1 -SVD with grid interval $r = 1^\circ$. Fig. 8.5 presents the bias for each of the DOA estimations as a function

Figure 8.4: Histograms of the DOA estimation errors of (a) ℓ_1 -SVD, (b) SBL-SVD and (c) OGSBL-SVD for source 2. Source 1: 50° ; source 2: 60° . SNR = 10 dB. Grid interval $r = 1^\circ$. The errors of ℓ_1 -SVD range from -4° to 2° while SBL-SVD finds the DOAs without error in all the trials. All errors of OGSBL-SVD are within half the grid interval without exploiting the information that the true DOAs are on the sampling grid.

of the angular separation with SNR = 10dB. A significant bias occurs all the time for both sources when using ℓ_1 -SVD since the on-grid model cannot describe the true observation model accurately. As a result, the best DOA estimate that ℓ_1 -SVD can obtain is the grid point nearest to the true DOA. It can be seen that the bias of source 1 is around -0.3° when source 2 is far away from source 1 and the bias of source 2 depends on its distance to the nearest grid point. Moreover, a large bias exhibits for closely spaced sources for ℓ_1 -SVD. For OGSBL-SVD, only a small asymptotic bias exists similar to the on-grid case due to the joint estimate of the grid offset.

Next, we compare OGSBL-SVD with ℓ_1 -SVD in terms of mean squared error (MSE) and computational time with respect to the grid interval r and SNR. In our experiment, we consider SNR = 10 and 0dB, and $r = 0.5^\circ, 1^\circ, 2^\circ$ and 4° . In each trial, $K = 2$ sources θ_1, θ_2 are uniformly generated within intervals $[58^\circ, 62^\circ]$ and $[86^\circ, 90^\circ]$ respectively. For each combination (SNR, r), the MSE is averaged over $R = 200$ trials:

$$\text{MSE} = \frac{1}{RK} \sum_{i=1}^R \sum_{k=1}^K \left(\theta_k^i - \hat{\theta}_k^i \right)^2, \quad (8.29)$$

where the superscript i refers to the i th trial. According to the results in the last section, there exists a lower bound for the MSE of ℓ_1 -SVD regardless of the SNR since

Figure 8.5: Bias comparison of OGSBL-SVD and ℓ_1 -SVD in localizing two sources as a function of angular separation with source 1 fixed at 50.3° and $\text{SNR} = 10$ dB. A small asymptotic bias occurs in OGSBL-SVD while a significant bias exhibits all the time when using ℓ_1 -SVD.

the best DOA estimate that ℓ_1 -SVD can obtain is the grid point nearest to the true DOA. By assuming that the true DOA is uniformly distributed, the lower bound is $\text{LB} = r^2/12$. Fig. 8.6 presents our experimental results. In all scenarios under consideration, OGSBL-SVD is more accurate than ℓ_1 -SVD. Moreover, OGSBL-SVD can exceed the lower bound of ℓ_1 -SVD in most scenarios. The phenomenon is significant in the high SNR regime or in the coarse sampling grid case where the on-grid model is poor in describing the true observation model while the modeling error can be overcome to a large extent by the off-grid model used in this chapter.

Table 8.1 presents the averaged CPU times of OGSBL-SVD and ℓ_1 -SVD (excluding the SVD process that takes about 0.003s in our case) with respect to SNR and r .³ For both OGSBL-SVD and ℓ_1 -SVD, their CPU times decrease as the grid gets coarser. OGSBL-SVD is faster than ℓ_1 -SVD at $r = 2^\circ$ and 4° . One drawback of the proposed method is that it is slow in the case of a dense sampling grid. In practice,

³The code of ℓ_1 -SVD is provided by the author of [88]. We note that its speed can be accelerated using state-of-the-art algorithms for CS.

Figure 8.6: MSEs of OGSBL-SVD and ℓ_1 -SVD. The lower bound is for ℓ_1 -SVD regardless of the SNR.

we recommend to use a coarser grid with $r = 2^\circ$ for the proposed algorithm since it can give an accurate yet fast DOA estimation.

8.3.6.2 Comparison with STLS

The off-grid model has recently been used in [78] for the DOA estimation. In [78], a sparse total least-squares (STLS) approach is proposed. In the SMV case, STLS

Table 8.1: Averaged Time Consumptions of ℓ_1 -SVD and OGSBL-SVD with Respect to SNR and r . Time unit: sec.

SNR = 10dB				
	$r = 0.5^\circ$	$r = 1^\circ$	$r = 2^\circ$	$r = 4^\circ$
ℓ_1 -SVD	0.601	0.413	0.324	0.291
OGSBL-SVD	10.2	0.782	0.096	0.025
SNR = 0dB				
	$r = 0.5^\circ$	$r = 1^\circ$	$r = 2^\circ$	$r = 4^\circ$
ℓ_1 -SVD	0.413	0.295	0.218	0.190
OGSBL-SVD	10.9	0.831	0.104	0.024

seeks to solve the nonconvex optimization problem

$$\min_{\mathbf{x}, \boldsymbol{\beta}} \left\{ \|\boldsymbol{\beta}\|_2^2 + \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\text{diag}(\boldsymbol{\beta})) \mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1 \right\}, \quad (8.30)$$

where $\mathbf{x}$ is the sparse source signal of interest, $\mathbf{y}$ is the noisy measurement, $\mathbf{A}$, $\mathbf{B}$ and $\boldsymbol{\beta}$ are the same as defined in the off-grid model, and $\lambda > 0$ is a regularization parameter. From the Bayesian perspective, this is equivalent to seeking for an MAP solution of $(\mathbf{x}, \boldsymbol{\beta})$ by assuming that the measurement noise is white Gaussian, $\mathbf{x}$ is Laplacian and $\boldsymbol{\beta}$ is Gaussian. It is noted that the last assumption for $\boldsymbol{\beta}$ cannot properly capture the property of $\boldsymbol{\beta}$. A local minima of the problem in (8.30) is achieved in [78] by an alternating approach, i.e., alternatively solving $\mathbf{x}$ with a fixed $\boldsymbol{\beta}$, which requires a solution to an ℓ_1 -regularized least square problem, and solving $\boldsymbol{\beta}$ with a fixed $\mathbf{x}$, which requires a solution to an N dimensional linear system. As argued in [88] the SVD used in OGSBL-SVD can alleviate the sensitivity to the measurement noise in the MMV case that is not used in [78]. To make a fair comparison, we consider only the SMV case when comparing our method with STLS though a similar problem can be cast for STLS in the MMV case. In our implementation of OGSBL, we initialize $\alpha_0 = \frac{100}{\text{Var}\{\mathbf{y}\}}$ and $\boldsymbol{\alpha} = \frac{1}{MK} |\mathbf{A}^H \mathbf{y}|$. The rest settings are the same as those for OGSBL-SVD in the MMV case.

In our experiment, we consider two DOAs from 63.2° and 90.3° with $\text{SNR} = 20\text{dB}$. We consider $r = 2^\circ$ and 4° for both OGSBL and STLS. The parameter λ in (8.30) is tuned to our best such that STLS achieves the smallest error. Table 8.2 presents the averaged MSEs and CPU times of STLS and OGSBL over $R = 200$ trials. OGSBL obtains more accurate DOA estimations than STLS in both the scenarios with remarkably less computational times. We also note that it is possible to accelerate STLS using state-of-the-art algorithms for CS.

Table 8.2: Averaged MSEs and CPU Times of STLS and OGSBL in the SMV Case with Respect to r .

	MSE (dB)		Time (sec)	
	$r = 2^\circ$	$r = 4^\circ$	$r = 2^\circ$	$r = 4^\circ$
STLS	-36.5	-36.6	5.31	1.77
OGSBL	-45.1	-43.3	0.098	0.028

8.4 Conclusion

In this chapter, we studied the off-grid DOA estimation model firstly proposed in [78] for reducing the modeling error due to discretization of a continuous range. An ℓ_1 optimization-based approach was firstly presented to solve the problem under the framework of perturbed CS we have studied in Chapter 5. To alleviate the computational burden, an SBL algorithm was then proposed in which the sparse signal and the grid offset are jointly estimated iteratively. Numerical simulations were provided to validate our approaches.

Chapter 9

Sparse MRI for Motion Correction

The application of CS to the DOA estimation has been studied in the last chapter to resolve the basis mismatch problem that exists in the literature. In this chapter, we propose a novel application of CS to motion correction in magnetic resonance imaging (MRI) while CS has been widely applied to MRI for imaging acceleration. Imaging artifacts due to patient motion remain a major challenge in many MRI applications. In this chapter, we consider MR image reconstruction from k -space data corrupted by 2D translational motion. The problem is ill-posed since there exist infinitely many image candidates and possible motions which result in the same k -space data. Existing methods solve the problem by acquiring more k -space data and/or adopting a specific sampling sequence, see [95–99].

In this chapter, we present a sparsity-based approach to rigid-body motion correction for the first time. The presented approach relies on the assumption that little motion occurs during a single readout line, which can be approximately satisfied in practice and is greatly relaxed in comparison with those in [97, 99]. We model the motion occurring during the scanning process as unknown parameters in the sensing system and seek after the motion estimate such that the compensated MR image is maximally sparse/compressible among the infinite candidates. The idea is inspired by our results in Chapter 5 where it is shown that uncertain system parameters can be accurately estimated along with the sparse signal in one particular situation. We

present efficient iterative algorithms to jointly estimate the motion and the image content. The proposed method has a lot of merits, such as no need of additional data and loose requirement for the sampling sequence. Promising results are presented to demonstrate its performance. This chapter is mainly based on [12].

9.1 Problem Formulation of Motion Correction

Consider that a translational motion $\beta_{\mathbf{k}}$ occurs when acquiring the k -space measurement at $\mathbf{k}$. So it holds at the moment $\overline{m}(\mathbf{r}) = m^o(\mathbf{r} - \beta_{\mathbf{k}})$, where m^o is the MR image of interest, $\overline{m}$ is the translated image and $\mathbf{r}$ denotes the coordinate in the image domain. According to the relationship between an MR image and its k -space measurements, it holds

$$\begin{aligned}\overline{M}(\mathbf{k}) &= \iint \overline{m}(\mathbf{r}) e^{-i2\pi\langle\mathbf{k}, \mathbf{r}\rangle} d\mathbf{r} \\ &= e^{-i2\pi\langle\mathbf{k}, \beta_{\mathbf{k}}\rangle} M^o(\mathbf{k}),\end{aligned}\tag{9.1}$$

where $M^o = \mathcal{F}m^o$ and $\overline{M}$ denote the motion-free and motion-corrupted k -space data, respectively, with $\mathcal{F}$ denoting the 2D Fourier transform operator. It is shown in (9.1) that the translational motion leads to a phase error and unaltered amplitude. To represent (9.1) more compactly, let $\beta = [\cdots, \beta_{\mathbf{k}_j}, \cdots]^T$ and then write (9.1) into

$$\overline{M} = \mathcal{T}_{\beta} \mathcal{F} m^o,\tag{9.2}$$

where $\mathcal{T}_{\beta}$ denotes a linear operator caused by the translational motion β and is referred to as the translational operator hereafter. In particular, let Λ_{β} be a matrix of the same dimension as $\overline{M}$ with its element $\Lambda_{\beta}(\mathbf{k}) = e^{-i2\pi\langle\mathbf{k}, \beta_{\mathbf{k}}\rangle}$. Then $\mathcal{T}_{\beta} M^o = \Lambda_{\beta} \odot M^o$ where $\odot$ denotes the Hadamard product. In addition, it holds that $\mathcal{T}_{\beta}^{-1} = \mathcal{T}_{-\beta}$ and $\mathcal{T}_0 = \mathcal{I}$ where $\mathcal{I}$ denotes the identity operator. For a particular motion β , the linear system of equations in (9.2) relates the MR image m^o of interest and the acquired k -space data $\overline{M}$. Our objective is to reconstruct the image m^o and possibly the motion β given the corrupted k -space data $\overline{M}$.

The challenge is that the problem of recovering $\mathbf{m}^o$ from $\overline{\mathbf{M}}$ is ill-posed since there exist infinite number of solutions to (9.2). In particular, for any β there exists a solution $\mathbf{m} = \mathcal{F}^{-1}\mathcal{T}_{-\beta}\overline{\mathbf{M}}$ such that (9.2) holds. In this regard, the full sampled k -space data can be considered as “compressive” measurements in the language of CS. To choose the correct one among the infinite candidates, as a result, additional information has to be exploited.

9.2 Sparsity-Driven Motion Correction

9.2.1 Sparsity-Based Formulation

We first consider that the MR image is directly constructed from the corrupted k -space data $\overline{\mathbf{M}}$ without considering the translational motion. Then, the constructed image is

$$\widehat{\mathbf{m}} = \mathcal{F}^{-1}\overline{\mathbf{M}} = \mathcal{F}^{-1}(\mathbf{\Lambda}_{\beta} \odot \mathbf{M}^o) = (\mathcal{F}^{-1}\mathbf{\Lambda}_{\beta}) \otimes \mathbf{m}^o, \quad (9.3)$$

where $\otimes$ denotes the circular convolution operation. So the obtained image is the true image $\mathbf{m}^o$ after a convolution with $\mathcal{F}^{-1}\mathbf{\Lambda}_{\beta}$. That explains how the imaging artifacts come from the translational motion. Due to the imaging artifacts, it is natural to conjecture that the translational motion will reduce the sparsity/compressibility of MR images, i.e., the true image is the maximally sparse/compressible solution (under an appropriate basis) to (9.2). Examples of a Shepp-Logan phantom and simulated human brain are presented in Fig. 9.1. In comparison with the motion-free images (in col 1), severe artifacts are present in the motion-corrupted ones (in col 2). Numerically, it can be shown that the ℓ_1 norms (a commonly used sparsity metric) of the motion-corrupted images (under an Haar wavelet basis) are larger than those of the motion-free ones.

Based on the conjecture above, we propose to reconstruct the motion-free MR image using the maximally sparse solution that satisfies the data consistency constraint

(9.2), i.e., by solving a BP like optimization problem:

$$\min_{\mathbf{m}, \boldsymbol{\beta}} \|\mathcal{W}\mathbf{m}\|_1, \text{ subject to } \mathcal{T}_{\boldsymbol{\beta}}\mathcal{F}\mathbf{m} = \overline{\mathbf{M}} \text{ and } \boldsymbol{\beta} \in \mathcal{D}_{\boldsymbol{\beta}}, \quad (9.4)$$

where $\mathcal{W}$ is a sparsifying operator (e.g., a wavelet transform), $\|\cdot\|_1$ denotes the ℓ_1 norm (sum of amplitude of all elements) that is commonly used to promote sparsity and $\mathcal{D}_{\boldsymbol{\beta}}$ denotes the domain of $\boldsymbol{\beta}$. Given an appropriate constant C , an equivalent (LASSO-like) formulation of (9.4) is

$$\min_{\mathbf{m}, \boldsymbol{\beta}} \|\overline{\mathbf{M}} - \mathcal{T}_{\boldsymbol{\beta}}\mathcal{F}\mathbf{m}\|_{\text{F}}, \text{ subject to } \|\mathcal{W}\mathbf{m}\|_1 \leq C \text{ and } \boldsymbol{\beta} \in \mathcal{D}_{\boldsymbol{\beta}}, \quad (9.5)$$

where $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm and $C < \|\mathcal{W}\mathcal{F}^{-1}\overline{\mathbf{M}}\|_1$ is a constant that needs to be tuned in practice and can be set to $C = \|\mathcal{W}\mathbf{m}^o\|_1$ in an ideal case. The motions $\boldsymbol{\beta}_{\mathbf{k}}$ with respect to the coordinate $\mathbf{k}$ are not independent. For example, the motion should be piecewise smooth when sorted chronologically. In this chapter, we assume that the same motion is shared among every readout line, which can be approximately satisfied in practice and greatly reduces the number of unknown parameters. Note that the assumption is greatly relaxed comparing to that in [97,99] where it is assumed that little motion exists during every echotrain (each comprising several readout lines). Moreover, $\boldsymbol{\beta}$ should be properly bounded in practice. The *a priori* knowledge narrows the selection of $\boldsymbol{\beta}$ and defines $\mathcal{D}_{\boldsymbol{\beta}}$. To the best of our knowledge, (9.4) and (9.5) present the first sparsity-based formulations for the MRI motion correction problem though CS has been widely applied to MRI for imaging acceleration.

9.2.2 An Alternating Algorithm

Due to the nonconvexity with respect to $\boldsymbol{\beta}$, problem (9.5) is nonconvex. We propose an alternating algorithm that starts with $\boldsymbol{\beta}^{(0)} = \mathbf{0}$ and iteratively carries out the following steps:

1) solving

$$\widetilde{\mathbf{m}}^{(j+1)} = \arg \min_{\mathbf{m}} \|\overline{\mathbf{M}} - \mathcal{T}_{\beta^{(j)}} \mathcal{F} \mathbf{m}\|_{\text{F}}, \text{ subject to } \|\mathcal{W} \mathbf{m}\|_1 \leq C. \quad (9.6)$$

2) solving

$$\beta^{(j+1)} = \arg \min_{\beta} \|\overline{\mathbf{M}} - \mathcal{T}_{\beta} \mathcal{F} \widetilde{\mathbf{m}}^{(j+1)}\|_{\text{F}}, \text{ subject to } \beta \in \mathcal{D}_{\beta}. \quad (9.7)$$

To interpret the algorithm above, denote two sets $S_1 = \{\mathbf{m} : \|\mathcal{W} \mathbf{m}\|_1 \leq C\}$ and $S_2 = \{\mathbf{m} : \overline{\mathbf{M}} = \mathcal{T}_{\beta} \mathcal{F} \mathbf{m}, \beta \in \mathcal{D}_{\beta}\}$. We refer to S_1 as the *sparse domain* and S_2 as the *Fourier domain*, respectively. Denote $\mathcal{P}_1$ and $\mathcal{P}_2$ projections onto S_1 and S_2 , respectively. Let $\mathbf{m}^{(j)} = \mathcal{F}^{-1} \mathcal{T}_{-\beta^{(j)}} \overline{\mathbf{M}}$. Then it holds that $\widetilde{\mathbf{m}}^{(j+1)} = \mathcal{P}_1 \mathbf{m}^{(j)}$ following from *Step 1*. At *Step 2*, the equality $\mathbf{m}^{(j+1)} = \mathcal{P}_2 \widetilde{\mathbf{m}}^{(j+1)}$ holds. So the algorithm above is equivalent to the recursion

$$\mathbf{m}^{(j+1)} = \mathcal{P}_2 \mathcal{P}_1 \mathbf{m}^{(j)}, \quad (9.8)$$

starting with $\mathbf{m}^{(0)} = \widehat{\mathbf{m}}$ where $\widehat{\mathbf{m}}$ is as defined in (9.3). Since $\mathbf{m}^{(0)} \in S_2$, the recursion attempts to find a point in S_2 nearest to S_1 , or equivalently, the maximally sparse image that is consistent with the k -space observation. The constant C is generally unavailable in advance and needs to be tuned in the algorithm. A simple method is to set C such that the obtained image has the least ℓ_1 norm in the sparsifying domain (which is consistent with our objective). The remaining tasks are to compute $\mathcal{P}_1$ and $\mathcal{P}_2$.

9.2.3 Computation of $\mathcal{P}_1$

To compute $\mathcal{P}_1$ (i.e., to solve (9.6)) is a convex problem. We consider the case where the sparsifying transform $\mathcal{W}$ is a unitary operator and has a fast operation (e.g., an Haar wavelet). A fast algorithm can be developed as follows. Denote the set $S_3 = \{\mathbf{w} : \|\mathbf{w}\|_1 \leq C\}$ and $\mathcal{P}_3$ the projection onto the set. It is easy to show that

$\mathcal{P}_1 \mathbf{m} = \mathcal{W}^{-1} \mathcal{P}_3 (\mathcal{W} \mathbf{m})$. The following lemma presents a fast approach to computing $\mathcal{P}_3$ which has a computational complexity of $O(N \log N)$, where N denotes the size of the variable.

Lemma 9.1. *It holds that $\mathcal{P}_3 \mathbf{v} = \mathbf{v}$ if $\|\mathbf{v}\|_1 \leq C$. Otherwise, denote $\mathbf{v}^k$ and $|v|_{(k)}$ the best k -term approximation and the k th largest (in amplitude) entry of $\mathbf{v}$, respectively, for an integer k . Let $k^* = \max \left\{ k : \|\mathbf{v}^k\|_1 - k |v|_{(k)} < C \right\}$ and $\lambda = \frac{\|\mathbf{v}^{k^*}\|_1 - C}{k^*}$. Then $\mathcal{P}_3 \mathbf{v} = S_\lambda(\mathbf{v})$, where $S_\lambda(\cdot)$ denotes a soft thresholding operator defined in (3.4).*

Proof. Notice that

$$\mathcal{P}_3 \mathbf{v} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2, \text{ subject to } \|\mathbf{x}\|_1 \leq C. \quad (9.9)$$

It is trivial for the case when $\|\mathbf{v}\|_1 \leq C$. For the case $\|\mathbf{v}\|_1 > C$, it is easy to show that $\|\mathbf{x}^*\|_1 = C$ for the optimal solution $\mathbf{x}^*$. We next consider the Lagrangian function

$$f(\mathbf{x}, \lambda) = \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + \lambda (\|\mathbf{x}\|_1 - C). \quad (9.10)$$

So it holds that $\mathbf{x}^* = S_\lambda(\mathbf{v})$ for some $\lambda > 0$. Suppose $\|\mathbf{x}^*\|_0 = k$, i.e., $\mathbf{x}^*$ has k nonzero entries. It follows from $C = \|\mathbf{x}^*\|_1 = \|\mathbf{v}^k\|_1 - k\lambda$ that $\lambda = \frac{\|\mathbf{v}^k\|_1 - C}{k}$. The remaining task is to determine the integer k .

By $\mathbf{x}^* = S_\lambda(\mathbf{v})$ and $\|\mathbf{x}^*\|_0 = k$, we have $|v|_{(k+1)} \leq \lambda < |v|_{(k)}$. After substituting λ into $\lambda = \frac{\|\mathbf{v}^k\|_1 - C}{k}$, we have the inequality

$$k |v|_{(k+1)} \leq \|\mathbf{v}^k\|_1 - C < k |v|_{(k)}. \quad (9.11)$$

In the following we show that such k is unique. To do this, let $g(k) = \|\mathbf{v}^k\|_1 - k |v|_{(k)} - C$. It is easy to show that $g(k)$ is monotonically increasing with respect to k . Since $g(k+1) = \|\mathbf{v}^{k+1}\|_1 - (k+1) |v|_{(k+1)} - C$, the inequality (9.11) is equivalent to $g(k) < 0 \leq g(k+1)$, which by the monotonicity of $g(k)$ determines a unique k that can be expressed as in the lemma. $\blacksquare$

9.2.4 Approximate Computation of $\mathcal{P}_2$

The computation of $\mathcal{P}_2$ is equivalent to solving the optimization problem in (9.7). Under the assumption that a common displacement is shared among each readout line and the motions among different readout lines are independent, the number of unknown variables in β is reduced to twice the number of readout lines (each displacement has 2 degrees of freedom in the 2D case) and the displacement in each line can be separately calculated. However, the exact calculation of $\mathcal{P}_2$, or equivalently, to obtain the global optimum of (9.7), is generally infeasible since the Fourier domain S_2 is severely nonconvex. The β solution is typically trapped at a local minimum if optimization methods, e.g., a gradient method, are used to solve (9.7). Inspired by a navigator-based method in [98], we propose a practically efficient algorithm as follows.

In (9.7), the term $\mathcal{F}\widetilde{\mathbf{m}}^{(j)}$, denoted by $\mathbf{M}^{(j)}$, can be considered as the current estimate of the motion-free k -space data $\mathbf{M}^o$ (without accounting for the difference between $\widetilde{\mathbf{m}}^{(j)}$ and $\mathbf{m}^{(j)}$). To solve (9.7) is in fact to estimate the translational motion between the motion-corrupted observation $\overline{\mathbf{M}}$ and the current motion-free estimate $\mathbf{M}^{(j)}$. So we may consider each readout line of $\overline{\mathbf{M}}$ as a navigator and the associated line of $\mathbf{M}^{(j)}$ as the reference. Then the motion within each readout line can be estimated using the navigator-based method in [98]. We omit the details. At last, we point out that the algorithm performance can be improved in practice by a modification of the current motion-free k -space estimate: $\mathbf{M}^{(j)} = |\overline{\mathbf{M}}| \odot \text{sgn}(\mathcal{F}\widetilde{\mathbf{m}}^{(j)})$, instead of using $\mathbf{M}^{(j)} = \mathcal{F}\widetilde{\mathbf{m}}^{(j)}$, where the absolute operator $|\cdot|$ does an element-wise operation. The underlying reason is obvious according to (9.1). The translational motion changes only the phase information of the k -space data and keeps the amplitude. So intuitively, the modified estimate is closer to the true motion-free data and leads to better performance.

9.2.5 Sparse RAAR for Motion Correction

According to (9.8), the algorithm in Subsection 9.2.1 is implemented by iteratively projecting onto the two sets S_1 and S_2 . It is related to iterative projection algorithms [142] for phase retrieval [143], which aims to find an intersection point of two sets and are built upon combining projections onto the two sets in some fashion. In fact, the recursion (9.8) is exactly the error reduction (ER) algorithm [142] without accounting for the differences of the sets. One drawback of ER is its slow convergence. In our setting, we want to solve a feasibility problem:

$$\text{find } \mathbf{m} \in S_1 \cap S_2, \quad (9.12)$$

i.e., to find a point lying in both the sparse domain S_1 and the Fourier domain S_2 . The relaxed averaged alternating reflections (RAAR) algorithm introduced in [144] has fast convergence speed and stable performance. RAAR starts with some initial point $\mathbf{m}^{(0)}$ and is defined by the recursion

$$\mathbf{m}^{(j+1)} = \left(\frac{\theta}{2} (\mathcal{R}_1 \mathcal{R}_2 + \mathcal{I}) + (1 - \theta) \mathcal{P}_2 \right) \mathbf{m}^{(j)}, \quad (9.13)$$

where $\theta \in [0, 1]$ is a constant, and reflectors $\mathcal{R}_1 = 2\mathcal{P}_1 - \mathcal{I}$ and $\mathcal{R}_2 = 2\mathcal{P}_2 - \mathcal{I}$. By defining and computing $\mathcal{P}_1$ and $\mathcal{P}_2$ as before, (9.13) defines an algorithm for the motion correction problem, named as sparse RAAR (SRAAR), where $\widehat{\mathbf{m}}$ in (9.3) is used as the initial point.

9.3 Numerical Simulations

The proposed method is validated using a Shepp-Logan phantom and simulated human brain data obtained from BrainWeb¹ with both of size 256×256 . Continuous translational motions are randomly generated among the readout lines. Motion artifacts are introduced by applying varying linear phase shifts to the motion-free

¹<http://www.bic.mni.mcgill.ca/brainweb>

k -space data according to (9.1). An Haar wavelet is selected as the sparsifying transform. SRAAR is applied to the motion-corrupted k -space data to reconstruct the MR images with the setting $\theta = 0.9$. SRAAR is terminated after a fixed number of iterations.

Simulation results are presented in Fig. 9.1. For the phantom (row 1), almost exact reconstruction is obtained using SRAAR. A small amount of motions (within 5 pixels along both the readout and phase-encode directions) are studied in rows 2 (without noise) and 3 (noise added) with the human brain. It can be seen that even a small amount of motions may cause severe imaging artifacts. After the sparsity-based motion correction with SRAAR, only few artifacts remain. In the presence of both noises and a large amount of motions (3 times as large as those in rows 2 and 3), it is shown in row 4 that SRAAR may have difficulties to produce a good result though most artifacts are removed. We note that SRAAR is computationally efficient in general. In the above, each iteration takes about 1s in Matlab v.7.7.0 on a PC with a 3GHz CPU, i.e., each reconstructed image is obtained within few minutes. Moreover, SRAAR can be greatly accelerated by estimating motions in the readout lines in parallel when computing $\mathcal{P}_2$.

9.4 Conclusion

A first sparsity-based approach to motion correction was introduced in this chapter. An efficient algorithm was proposed partially inspired by phase retrieval and existing navigator-based methods. The promising results presented in this chapter indicate good potential for practical application of the proposed method. The current work considers only translational motion and encourages further studies of more complicated motions. The results of Chapter 5 reveal that unknown system parameters can be estimated even in the undersampling case, suggesting that imaging acceleration is also possible in the presence of motions.

Figure 9.1: Simulation results of sparsity-based motion correction on a Shepp-Logan phantom (row 1, 100 iterations), simulated human brain (row 2, 200 iterations), and noisy brain with small motions (row 3, 200 iterations) and large motions (row 4, 400 iterations).

Chapter 10

Conclusion and Future Work

10.1 Conclusion

As an emerging technique which enables sub-Nyquist sampling and processing, CS has attracted attentions of many researchers in the past few years. In this thesis, we have studied several problems of CS ranging from theoretical analysis and algorithm design to practical applications. The contributions of the thesis can be summarized in the following aspects.

- We have discovered a new phase transition curve for the complex CS problem. The complex phase transition curve determines precisely the sparsity-undersampling tradeoff of ℓ_1 minimization in the complex setting and is positioned well above the known one in the real domain.
- We have analyzed the CS problem subject to structured matrix perturbations which has practical relevance to the DOA estimation problem. Sufficient conditions in terms of RIP were provided such that the formulated ℓ_1 minimization problems give accurate or even exact signal recoveries.
- We have proposed efficient algorithms for the sparse signal recovery in CS, including ONE-L1 algorithms for the convex relaxation method and novel

SBL algorithms based on a new sparsity-inducing prior. Their performances were demonstrated for 1D synthetic data and 2D images.

- We have presented a unified Bayesian framework and algorithm for CS problems with multiple-bit and 1-bit quantized measurements. The novel framework and algorithm can be applied to various scenarios including noiseless/noisy environment and unsaturated/saturated quantizer.
- We have proposed a fast and accurate approach to the DOA estimation by applying CS. The new approach resolves the problem that the true DOAs do not lie on the discretized sampling grid and can maintain high estimation accuracy under a coarse sampling grid.
- We have applied CS to MRI and presented the first sparsity-based approach to rigid-body motion correction in MRI. The new method solves the problem by exploiting the image sparsity in the wavelet domain and does not need to acquire additional k -space data.

10.2 Future Work

We list some possible directions for future studies.

- **Analysis of noisy 1-bit CS:** A convex formulation of noisy 1-bit CS problem has been studied in [68] with guaranteed signal recovery accuracy proven. The formulation requires the knowledge of the signal's ℓ_1 norm which cannot be known *a priori* in practice. In Chapter 7 we have introduced a different convex formulation (problem (7.16) with $f(\mathbf{x}) = \|\mathbf{x}\|_1$ and $s = 1$) as follows:

$$\min_{\mathbf{x}, \mathbf{y}} \|\mathbf{x}\|_1, \text{ subject to } \begin{cases} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \epsilon, \\ \text{diag}(\mathbf{b}) \mathbf{y} \geq 0, \\ \mathbf{b}^T \mathbf{y} = 1, \end{cases} \quad (10.1)$$

where $\mathbf{x}$ is the signal of interest, ϵ refers to the noise level and $\mathbf{b}$ denotes the vector of sign measurements. This formulation uses the noise information rather than the signal sparsity and can be of independent interest. One future work is to analyze its theoretical guarantees.

- **Efficient algorithm for quantized CS:** One drawback of the Q-VMP algorithm in Chapter 7 for quantized CS is its high computational complexity due to an inversion of a high dimensional matrix at each iteration though it has been greatly alleviated with a basis pruning approach adopted in that chapter. So another future work is to develop fast alternatives to the current implementation.
- **Algorithm for motion correction:** In Chapter 9 we have presented the SRAAR algorithm for the motion correction in MRI. A shortcoming of the algorithm is the need of specifying a constant which determines the *sparse* domain and is typically unavailable in practice. While the process of parameter tuning can be time-consuming, it is of great interest to design efficient algorithms which do not need the constant. A possibly useful result is presented in our recent work [145].
- **Continuous DOA estimation:** Most of the existing CS-based methods for the DOA estimation are based on the standard CS formulation which can be considered as a zeroth-order approximation of the true observation model via discretizing the continuous DOA range. Our methods in Chapter 8 rely on a first-order approximation via discretization and linear approximation. One future work is to do DOA estimation directly in the continuous range without discretization. Recently, there have been some results in the case of a single snapshot (a.k.a. line spectral estimation), where one needs to estimate the frequencies of some sinusoids from their superposition [146–148]. The so-called atomic norm [149] is used in existing results which is a generalization of the ℓ_1 norm to the continuous setting. After submission of this thesis, we have developed the first discretization-free sparse method in [150], named as the

sparse and parametric approach (SPA), for continuous DOA estimation in the presence of multiple snapshots, which can be applied to uniform/sparse linear arrays. While SPA is a statistical inference method, its connection to atomic norm-based methods has also been derived (see [151] and future publications).

- **Compressive phase retrieval:** Phase retrieval [143] refers to a classical problem of recovering a signal from its Fourier magnitude measurements. This problem arises in many imaging techniques where the phase information of the frequency data are not available to the detecting and sensing devices, e.g., X-ray crystallography [143]. It is also related to the motion correction in MRI that we have studied in Chapter 9 where the phase information, though not lost, is severely distorted. Conventional approaches to phase retrieval do not consider the image structure and solve the problem by oversampling in the frequency domain. Inspired by the success of CS which exploits the signal sparsity to reduce the number of measurements, sparsity-based phase retrieval has been studied recently, known as compressive phase retrieval (CPR) [152]. Though the CPR problem exhibits many similarities to CS the key difference is the nonlinearity of the measurements in CPR. One of the future directions is to design efficient algorithms and analyze their performance for CPR. Our preliminary results are presented in [145, 153].

Appendix A

Proofs of Results in Chapter 5

A.1 Proof of Theorem 5.3

Denote $\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\beta} \odot \mathbf{x} \end{bmatrix}$ and similarly define $\mathbf{z}^o$ and $\mathbf{z}^*$. Then the problem in (5.10) can be rewritten into

$$\min_{\mathbf{x} \in \mathbb{R}^N, \boldsymbol{\beta} \in [-r, r]^N} \|\mathbf{x}\|_0, \text{ subject to } \mathbf{y} = \boldsymbol{\Psi} \mathbf{z}. \quad (\text{A.1})$$

Let $\bar{\delta}_K = \bar{\delta}_K(\boldsymbol{\Psi})$ hereafter for brevity.

First note that $\mathbf{x}^o$ is K -sparse and $\mathbf{z}^o$ is $2K$ -D-sparse. Since $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ is a solution to the problem in (A.1), we have $\|\mathbf{x}^*\|_0 \leq \|\mathbf{x}^o\|_0 \leq K$ and, hence, $\mathbf{z}^*$ is $2K$ -D-sparse. By $\mathbf{y} = \boldsymbol{\Psi} \mathbf{z}^o = \boldsymbol{\Psi} \mathbf{z}^*$ we obtain $\boldsymbol{\Psi}(\mathbf{z}^o - \mathbf{z}^*) = \mathbf{0}$ and thus $\mathbf{z}^o - \mathbf{z}^* = \mathbf{0}$ by $\bar{\delta}_{4K} < 1$ and the fact that $\mathbf{z}^o - \mathbf{z}^*$ is $4K$ -D-sparse. We complete the proof by observing that

$$\mathbf{z}^o - \mathbf{z}^* = \begin{bmatrix} \mathbf{x}^o - \mathbf{x}^* \\ \boldsymbol{\beta}^o \odot \mathbf{x}^o - \boldsymbol{\beta}^* \odot \mathbf{x}^* \end{bmatrix} = \mathbf{0}.$$

A.2 Proof of Theorems 5.4 and 5.5

We only present the proof of Theorem 5.5 since Theorem 5.4 is a special case of Theorem 5.5. We first show the following lemma.

Lemma A.1. *We have*

$$|\langle \Psi \mathbf{v}, \Psi \mathbf{v}' \rangle| \leq \bar{\delta}_{2(K+K')} \|\mathbf{v}\|_2 \|\mathbf{v}'\|_2$$

for all $2K$ -D-sparse $\mathbf{v}$ and $2K'$ -D-sparse $\mathbf{v}'$ supported on disjoint subsets.

Proof. Without loss of generality, assume that $\mathbf{v}$ and $\mathbf{v}'$ are unit vectors with disjoint supports as above. Then by the definition of D-RIP and $\|\mathbf{v} \pm \mathbf{v}'\|_2^2 = \|\mathbf{v}\|_2^2 + \|\mathbf{v}'\|_2^2 = 2$ we have $2(1 - \bar{\delta}_{2(K+K')}) \leq \|\Psi \mathbf{v} \pm \Psi \mathbf{v}'\|_2^2 \leq 2(1 + \bar{\delta}_{2(K+K')})$. Consequently, $|\langle \Psi \mathbf{v}, \Psi \mathbf{v}' \rangle| \leq \frac{1}{4} \left| \|\Psi \mathbf{v} + \Psi \mathbf{v}'\|_2^2 - \|\Psi \mathbf{v} - \Psi \mathbf{v}'\|_2^2 \right| \leq \bar{\delta}_{2(K+K')}$, which completes the proof. ■

Using the notations $\mathbf{z}$, $\mathbf{z}^o$, $\mathbf{z}^*$ and $\bar{\delta}_K$ in Appendix A.1, P-BPDN in (5.11) can be rewritten into

$$\min_{\mathbf{x} \in \mathbb{R}^N, \beta \in [-r, r]^N} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - \Psi \mathbf{z}\|_2 \leq \epsilon. \quad (\text{A.2})$$

Let $\mathbf{h} = \mathbf{x}^* - \mathbf{x}^o$ and decompose $\mathbf{h}$ into a sum of K -sparse vectors $\mathbf{h}_{T_0}, \mathbf{h}_{T_1}, \mathbf{h}_{T_2}, \dots$, where T_0 denotes the set of indices of the K largest entries (in absolute value) of $\mathbf{x}^o$, T_1 the set of the K largest entries of $\mathbf{h}_{T_0^c}$ with T_0^c being the complementary set of T_0 , T_2 the set of the next K largest entries of $\mathbf{h}_{T_0^c}$ and so on. We abuse notations $\mathbf{z}_{T_j}^* = \begin{bmatrix} \mathbf{x}_{T_j}^* \\ \beta_{T_j}^* \odot \mathbf{x}_{T_j}^* \end{bmatrix}$, $j = 0, 1, 2, \dots$, and similarly define $\mathbf{z}_{T_j}^o$. Let $\mathbf{f} = \mathbf{z}^* - \mathbf{z}^o$ and $\mathbf{f}_{T_j} = \mathbf{z}_{T_j}^* - \mathbf{z}_{T_j}^o$ for $j = 0, 1, 2, \dots$. For brevity we write $T_{01} = T_0 \cup T_1$. To bound $\|\mathbf{h}\|_2$, in the first step we show that $\|\mathbf{h}_{T_{01}^c}\|_2$ is essentially bounded by $\|\mathbf{h}_{T_{01}}\|_2$, and then in the second step we show that $\|\mathbf{h}_{T_{01}}\|_2$ is sufficiently small.

The first step follows directly from the proof of Theorem 1.3 in [17]. So we have

$$\|\mathbf{h}_{T_{01}^c}\|_2 \leq \|\mathbf{h}_{T_0}\|_2 + 2K^{-1/2}e_0 \quad (\text{A.3})$$

with $e_0 \equiv \|\mathbf{x}^o - \mathbf{x}^K\|_1$.

In the second step, we bound $\|\mathbf{h}_{T_{01}}\|_2$ by utilizing its relationship with $\|\mathbf{f}_{T_{01}}\|_2$. Note that $\mathbf{f}_{T_j}$ for each $j = 0, 1, \dots$ is $2K$ -D-sparse. By $\Psi \mathbf{f}_{T_{01}} = \Psi \mathbf{f} - \sum_{j \geq 2} \Psi \mathbf{f}_{T_j}$ we have

$$\begin{aligned} \|\Psi \mathbf{f}_{T_{01}}\|_2^2 &= \langle \Psi \mathbf{f}_{T_{01}}, \Psi \mathbf{f} \rangle - \sum_{j \geq 2} \langle \Psi \mathbf{f}_{T_{01}}, \Psi \mathbf{f}_{T_j} \rangle \\ &\leq \|\Psi \mathbf{f}_{T_{01}}\|_2 \cdot \|\Psi \mathbf{f}\|_2 + \bar{\delta}_{4K} \|\mathbf{f}_{T_0}\|_2 \sum_{j \geq 2} \|\mathbf{f}_{T_j}\|_2 \\ &\quad + \bar{\delta}_{4K} \|\mathbf{f}_{T_1}\|_2 \sum_{j \geq 2} \|\mathbf{f}_{T_j}\|_2 \end{aligned} \quad (\text{A.4})$$

$$\leq \|\mathbf{f}_{T_{01}}\|_2 \left(2\epsilon \sqrt{1 + \bar{\delta}_{4K}} + \sqrt{2\bar{\delta}_{4K}} \sum_{j \geq 2} \|\mathbf{f}_{T_j}\|_2 \right). \quad (\text{A.5})$$

We used Lemma A.1 in (A.4). In (A.5), we used the D-RIP, and inequalities $\|\mathbf{f}_{T_0}\|_2 + \|\mathbf{f}_{T_1}\|_2 \leq \sqrt{2} \|\mathbf{f}_{T_{01}}\|_2$ and

$$\|\Psi \mathbf{f}\|_2 \leq \|\mathbf{y} - \Psi \mathbf{z}^*\|_2 + \|\mathbf{y} - \Psi \mathbf{z}^o\|_2 \leq 2\epsilon. \quad (\text{A.6})$$

By noting that $\beta^o, \beta^* \in [-r, r]^N$ and

$$\mathbf{f} = \begin{bmatrix} \mathbf{h} \\ \beta^* \odot \mathbf{h} + (\beta^* - \beta^o) \odot \mathbf{x}^o \end{bmatrix} \quad (\text{A.7})$$

we have

$$\|\mathbf{f}_{T_j}\|_2 \leq \sqrt{1 + r^2} \|\mathbf{h}_{T_j}\|_2 + 2r \|\mathbf{x}_{T_j}^o\|_2, \quad j = 0, 1, \dots \quad (\text{A.8})$$

Meanwhile,

$$\sum_{j \geq 2} \|\mathbf{x}_{T_j}^o\|_2 \leq \sum_{j \geq 2} \|\mathbf{x}_{T_j}^o\|_1 = \|\mathbf{x}_{T_{01}^c}^o\|_1 \leq e_0. \quad (\text{A.9})$$

Applying the D-RIP, (A.5), (A.8) and then (A.3) and (A.9) it gives

$$\begin{aligned} (1 - \bar{\delta}_{4K}) \|\mathbf{f}_{T_{01}}\|_2^2 &\leq \|\Psi \mathbf{f}_{T_{01}}\|_2^2 \\ &\leq \|\mathbf{f}_{T_{01}}\|_2 \left\{ 2\sqrt{1 + \bar{\delta}_{4K}}\epsilon + \sqrt{2(1 + r^2)}\bar{\delta}_{4K} \|\mathbf{h}_{T_0}\|_2 \right. \\ &\quad \left. + 2\sqrt{2}\bar{\delta}_{4K} \left[\sqrt{1 + r^2}K^{-1/2} + r \right] e_0 \right\}, \end{aligned}$$

and thus $\|\mathbf{h}_{T_{01}}\|_2 \leq \|\mathbf{f}_{T_{01}}\|_2 \leq c_1\epsilon + c_0\|\mathbf{h}_{T_{01}}\|_2 + (c_2K^{-1/2} + c_3)e_0$ with $c_0 \equiv \frac{\sqrt{2(1+r^2)\bar{\delta}_{4K}}}{1-\bar{\delta}_{4K}}$, $c_1 \equiv \frac{2\sqrt{1+\bar{\delta}_{4K}}}{1-\bar{\delta}_{4K}}$, $c_2 \equiv 2c_0$ and $c_3 \equiv \frac{2\sqrt{2}\bar{\delta}_{4K}r}{1-\bar{\delta}_{4K}}$. Hence, we get a bound $\|\mathbf{h}_{T_{01}}\|_2 \leq (1-c_0)^{-1} [c_1\epsilon + (c_2K^{-1/2} + c_3)e_0]$, which together with (A.3) gives

$$\begin{aligned} \|\mathbf{h}\|_2 &\leq \|\mathbf{h}_{T_{01}}\|_2 + \|\mathbf{h}_{T_{01}^c}\|_2 \leq 2\|\mathbf{h}_{T_{01}}\|_2 + 2K^{-1/2}e_0 \\ &\leq \frac{2c_1}{1-c_0}\epsilon + \left[\left(\frac{2c_2}{1-c_0} + 2 \right) K^{-1/2} + \frac{2c_3}{1-c_0} \right] e_0, \end{aligned} \quad (\text{A.10})$$

which concludes (5.14).

By (A.6), (A.7), (A.10) and the RIP we have

$$\begin{aligned} &(1-\bar{\delta}_{4K})^{1/2} \|(\boldsymbol{\beta}_{T_0}^* - \boldsymbol{\beta}_{T_0}^o) \odot \mathbf{x}_{T_0}^o\|_2 \\ &\leq \left\| \boldsymbol{\Psi} \begin{bmatrix} \mathbf{0} \\ (\boldsymbol{\beta}_{T_0}^* - \boldsymbol{\beta}_{T_0}^o) \odot \mathbf{x}_{T_0}^o \end{bmatrix} \right\|_2 \\ &= \left\| \boldsymbol{\Psi} \left(\mathbf{f} - \begin{bmatrix} \mathbf{h} \\ \boldsymbol{\beta}^* \odot \mathbf{h} \end{bmatrix} - \begin{bmatrix} \mathbf{0} \\ (\boldsymbol{\beta}_{T_0}^* - \boldsymbol{\beta}_{T_0}^o) \odot \mathbf{x}_{T_0}^o \end{bmatrix} \right) \right\|_2 \\ &\leq 2\epsilon + \sqrt{1+r^2} \|\boldsymbol{\Psi}\|_2 \|\mathbf{h}\|_2 + 2r \|\boldsymbol{\Phi}\|_2 e_0 \\ &\leq c_4\epsilon + (c_5K^{-1/2} + c_6)e_0 \end{aligned}$$

with $c_4 \equiv 2 + \frac{2\sqrt{1+r^2}\|\boldsymbol{\Psi}\|_2 c_1}{1-c_0}$, $c_5 \equiv \sqrt{1+r^2} \left(\frac{2c_2}{1-c_0} + 2 \right) \|\boldsymbol{\Psi}\|_2$ and $c_6 \equiv \left(\frac{2\sqrt{1+r^2}c_3}{1-c_0} + 2r \right) \|\boldsymbol{\Psi}\|_2$, and thus

$$\|(\boldsymbol{\beta}_{T_0}^* - \boldsymbol{\beta}_{T_0}^o) \odot \mathbf{x}_{T_0}^o\|_2 \leq [c_4\epsilon + (c_5K^{-1/2} + c_6)e_0] / \sqrt{1-\bar{\delta}_{4K}}, \quad (\text{A.11})$$

which concludes (5.15). We complete the proof by noting that the above results make sense if $c_0 < 1$, i.e., $\bar{\delta}_{4K} < \left(\sqrt{2(1+r^2)} + 1 \right)^{-1}$.

A.3 Proof of Lemma 5.1

We first consider the case where $\mathbf{v}$ is an interior point of $\mathcal{D}^*$, i.e., it holds that $\|\mathbf{y} - \boldsymbol{\Phi}^* \mathbf{v}\|_2 = \epsilon_0 < \epsilon$. Let $\eta = \epsilon - \epsilon_0$. Construct a sequence $\{\mathbf{v}^{(j)}\}_{j=1}^\infty$ such that

$\|\mathbf{v}^{(j)} - \mathbf{v}\|_2 \leq 1/j$. It is obvious that $\mathbf{v}^{(j)} \rightarrow \mathbf{v}$. We next show that $\mathbf{v}^{(j)} \in \mathcal{D}^j$ as j is large enough. By $\Phi^{(j)} \rightarrow \Phi^*$, $\mathbf{v}^{(j)} \rightarrow \mathbf{v}$ and that the sequence $\{\mathbf{v}^{(j)}\}_{j=1}^\infty$ is bounded, there exists a positive integer j_0 such that, as $j \geq j_0$, $\|\Phi^* - \Phi^{(j)}\|_2 \|\mathbf{v}^{(j)}\|_2 \leq \eta/2$ and $\|\Phi^*\|_2 \|\mathbf{v} - \mathbf{v}^{(j)}\|_2 \leq \eta/2$. Hence, as $j \geq j_0$,

$$\begin{aligned} & \|\mathbf{y} - \Phi^{(j)} \mathbf{v}^{(j)}\|_2 \\ &= \|(\mathbf{y} - \Phi^* \mathbf{v}) + (\Phi^* - \Phi^{(j)}) \mathbf{v}^{(j)} + \Phi^* (\mathbf{v} - \mathbf{v}^{(j)})\|_2 \\ &\leq \epsilon_0 + \eta/2 + \eta/2 = \epsilon, \end{aligned}$$

from which we have $\mathbf{v}^{(j)} \in \mathcal{D}^j$ for $j \geq j_0$. By re-selecting arbitrary $\mathbf{v}^{(j)} \in \mathcal{D}^j$ for $j < j_0$ we obtain the conclusion.

For the other case where $\mathbf{v}$ is a boundary point of $\mathcal{D}^*$, there exists a sequence $\{\mathbf{v}_{(l)}\}_{l=1}^\infty \subset \mathcal{D}^*$ with all $\mathbf{v}_{(l)}$ being interior points of $\mathcal{D}^*$ such that $\mathbf{v}_{(l)} \rightarrow \mathbf{v}$, as $l \rightarrow +\infty$. According to the first part of the proof, for each $l = 1, 2, \dots$, there exists a sequence $\{\mathbf{v}_{(l)}^{(j)}\}_{j=1}^\infty$ with $\mathbf{v}_{(l)}^{(j)} \in \mathcal{D}^j$, $j = 1, 2, \dots$, such that $\mathbf{v}_{(l)}^{(j)} \rightarrow \mathbf{v}_{(l)}$, as $j \rightarrow +\infty$. The sequence $\{\mathbf{v}_{(j)}^{(j)}\}_{j=1}^\infty$ is what we expected since

$$\|\mathbf{v}_{(j)}^{(j)} - \mathbf{v}\|_2 \leq \|\mathbf{v}_{(j)}^{(j)} - \mathbf{v}_{(j)}\|_2 + \|\mathbf{v}_{(j)} - \mathbf{v}\|_2 \rightarrow 0,$$

as $j \rightarrow +\infty$.

A.4 Proof of Theorem 5.7

We first show the existence of an accumulation point. It follows from the inequality

$$\|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^{(j)}) \mathbf{x}^{(j)}\|_2 \leq \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^{(j-1)}) \mathbf{x}^{(j)}\|_2 \leq \epsilon \quad (\text{A.12})$$

that $\mathbf{x}^{(j)}$ is a feasible solution to the problem in (5.19), and thus $\|\mathbf{x}^{(j+1)}\|_1 \leq \|\mathbf{x}^{(j)}\|_1$ for $j = 1, 2, \dots$. Then we have $\|\mathbf{x}^{(j)}\|_1 \leq \|\mathbf{x}^{(1)}\|_1 \leq \|\mathbf{A}^\dagger \mathbf{y}\|_1$ for $j = 1, 2, \dots$, since $\mathbf{A}^\dagger \mathbf{y}$ is a feasible solution to the problem in (5.19) at the first iteration with

the superscript † denoting the pseudo-inverse operator. This together with $\beta^{(j)} \in [-r, r]^N$, $j = 1, 2, \dots$, leads to that the sequence $\{(\mathbf{x}^{(j)}, \beta^{(j)})\}_{j=1}^\infty$ is bounded. Thus, there exists an accumulation point $(\mathbf{x}^*, \beta^*)$ of $\{(\mathbf{x}^{(j)}, \beta^{(j)})\}_{j=1}^\infty$.

For the accumulation point $(\mathbf{x}^*, \beta^*)$ there exists a subsequence $\{(\mathbf{x}^{(j_l)}, \beta^{(j_l)})\}_{l=1}^\infty$ of $\{(\mathbf{x}^{(j)}, \beta^{(j)})\}_{j=1}^\infty$ such that $(\mathbf{x}^{(j_l)}, \beta^{(j_l)}) \rightarrow (\mathbf{x}^*, \beta^*)$, as $l \rightarrow +\infty$. By (5.20), we have, for all $\beta \in [-r, r]^N$,

$$\|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^{(j_l)})\mathbf{x}^{(j_l)}\|_2 \leq \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta)\mathbf{x}^{(j_l)}\|_2, \quad (\text{A.13})$$

at both sides of which by taking $l \rightarrow +\infty$, we have, for all $\beta \in [-r, r]^N$,

$$\|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^*)\mathbf{x}^*\|_2 \leq \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta)\mathbf{x}^*\|_2, \quad (\text{A.14})$$

which concludes (5.22).

For (5.21), we first point out that $\|\mathbf{x}^{(j)}\|_1 \rightarrow \|\mathbf{x}^*\|_1$, as $j \rightarrow +\infty$, since $\{\|\mathbf{x}^{(j)}\|_1\}_{j=1}^\infty$ is decreasing and $\mathbf{x}^*$ is one of its accumulation points. As in Lemma 5.1, let $\mathcal{D}^j = \{\mathbf{x} : \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^{(j)})\mathbf{x}\|_2 \leq \epsilon\}$ and $\mathcal{D}^* = \{\mathbf{x} : \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^*)\mathbf{x}\|_2 \leq \epsilon\}$. By $\mathbf{A} + \mathbf{B}\Delta^{(j_l)} \rightarrow \mathbf{A} + \mathbf{B}\Delta^*$, as $l \rightarrow +\infty$, and Lemma 5.1, for any $\mathbf{x} \in \mathcal{D}^*$ there exists a sequence $\{\mathbf{x}_{(l)}\}_{l=1}^\infty$ with $\mathbf{x}_{(l)} \in \mathcal{D}^{j_l}$, $l = 1, 2, \dots$, such that $\mathbf{x}_{(l)} \rightarrow \mathbf{x}$, as $l \rightarrow +\infty$. By (5.19), we have, for $l = 1, 2, \dots$, $\|\mathbf{x}^{(j_l+1)}\|_1 \leq \|\mathbf{x}_{(l)}\|_1$, at both sides of which by taking $l \rightarrow +\infty$, we have $\|\mathbf{x}^*\|_1 \leq \|\mathbf{x}\|_1$. Finally, (5.21) is concluded since the inequality holds for arbitrary $\mathbf{x} \in \mathcal{D}^*$.

A.5 Proof of Theorem 5.8

We need to show that an optimal solution $(\mathbf{x}^*, \beta^*)$ satisfies (5.21) and (5.22). It is obvious for (5.21). For (5.22), we discuss two cases based on Lemma 5.2. If $\|\mathbf{y}\|_2 \leq \epsilon$, then $\mathbf{x}^* = \mathbf{0}$ and, hence, (5.22) holds for any $\beta^* \in [-r, r]^N$. If $\|\mathbf{y}\|_2 > \epsilon$, $\|\mathbf{y} - (\mathbf{A} + \mathbf{B}\Delta^*)\mathbf{x}^*\|_2 = \epsilon$ holds by (5.21) and Lemma 5.2. Next we use contradiction to show that (5.22) holds in such case.

Suppose that (5.22) does not hold as $\|\mathbf{y}\|_2 > \epsilon$. That is, there exists $\boldsymbol{\beta}' \in [-r, r]^N$ such that

$$\|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta}') \mathbf{x}^*\|_2 < \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta}^*) \mathbf{x}^*\|_2 = \epsilon \quad (\text{A.15})$$

holds with $\boldsymbol{\Delta}' = \text{diag}(\boldsymbol{\beta}')$. Then by Lemma 5.2 we see that $\mathbf{x}^*$ is a feasible but not optimal solution to the problem

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ subject to } \|\mathbf{y} - (\mathbf{A} + \mathbf{B}\boldsymbol{\Delta}') \mathbf{x}\|_2 \leq \epsilon. \quad (\text{A.16})$$

Hence, $\|\mathbf{x}'\|_1 < \|\mathbf{x}^*\|_1$ holds for an optimal solution $\mathbf{x}'$ to the problem in (A.16). Meanwhile, $(\mathbf{x}', \boldsymbol{\beta}')$ is a feasible solution to the P-BPDN problem in (5.11). Thus $(\mathbf{x}^*, \boldsymbol{\beta}^*)$ is not an optimal solution to the P-BPDN problem in (5.11) by $\|\mathbf{x}'\|_1 < \|\mathbf{x}^*\|_1$, which leads to contradiction.

Author's Publications

Preprints:

1. Z. Yang and L. Xie, "On sparse parametric methods for line spectral estimation from complete and incomplete data," submitted to *IEEE Transactions on Signal Processing*, November 2013.
2. Z. Yang, L. Xie, and C. Zhang, "A discretization-free sparse and parametric approach for linear array signal processing," submitted to *IEEE Transactions on Signal Processing*, July 2013.
3. Z. Yang, L. Xie, and C. Zhang, "Bayesian compressed sensing with new sparsity-inducing prior," Available online at <http://arxiv.org/abs/1208.6464>, 2013.
4. Z. Yang, C. Zhang, and L. Xie, "Robust compressive phase retrieval via l1 minimization with application to image reconstruction," Available online at <http://arxiv.org/abs/1302.0081>, 2013.

Journal papers:

1. Z. Yang, L. Xie, and C. Zhang, "Variational Bayesian algorithm for quantized compressed sensing," *IEEE Transactions on Signal Processing*, vol. 61, no. 11, pp. 2815–2824, 2013.

2. A. Maleki, L. Anitori, Z. Yang, and R. Baraniuk, “Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP),” *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4290–4308, 2013.
3. Z. Yang, L. Xie, and C. Zhang, “Off-grid direction of arrival estimation using sparse Bayesian inference,” *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 38–43, 2013.
4. Z. Yang, C. Zhang, and L. Xie, “Robustly stable signal recovery in compressed sensing with structured matrix perturbation,” *IEEE Transactions on Signal Processing*, vol. 60, no. 9, pp. 4658–4671, 2012.
5. Z. Yang, C. Zhang, and L. Xie, “On phase transition of compressed sensing in the complex domain,” *IEEE Signal Processing Letters*, vol. 19, no. 1, pp. 47–50, 2012.
6. Z. Yang, C. Zhang, J. Deng, and W. Lu, “Orthonormal expansion ℓ_1 -minimization algorithms for compressed sensing,” *IEEE Transactions on Signal Processing*, vol. 59, no. 12, pp. 6285–6290, 2011.

Conference papers:

1. C. Zhang, Z. Yang, and L. Xie, “The l_1 measure of image reconstruction subject to motion corrupted k -space data,” in *10th IEEE International Conference on Control and Automation (ICCA)*, 2013, pp. 1133–1137.
2. Z. Yang, C. Zhang, and L. Xie, “Sparse MRI for motion correction,” in *10th IEEE International Symposium on Biomedical Imaging (ISBI)*, 2013, pp. 962–965.
3. Z. Yang, L. Xie, and C. Zhang, “Accurate signal recovery in quantized compressed sensing,” in *15th IEEE International Conference on Information Fusion (FUSION)*, 2012, pp. 2531–2536.

4. Z. Yang, C. Zhang, and L. Xie, “Stable signal recovery in compressed sensing with a structured matrix perturbation,” in *IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, 2012, pp. 2737–2740.
5. J. Deng, Z. Yang, C. Zhang, and W. Lu, “Orthonormal expansion ℓ_1 -minimization for compressed sensing in MRI,” in *18th IEEE International Conference on Image Processing (ICIP)*, 2011, pp. 2297–2300.
6. Z. Yang and C. Zhang, “Sparsity-undersampling tradeoff of compressed sensing in the complex domain,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 3668–3671.

Bibliography

- [1] E. Candès and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?” *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [2] D. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [3] Z. Yang, C. Zhang, J. Deng, and W. Lu, “Orthonormal expansion ℓ_1 -minimization algorithms for compressed sensing,” *IEEE Transactions on Signal Processing*, vol. 59, no. 12, pp. 6285–6290, 2011.
- [4] Z. Yang, C. Zhang, and L. Xie, “On phase transition of compressed sensing in the complex domain,” *IEEE Signal Processing Letters*, vol. 19, no. 1, pp. 47–50, 2012.
- [5] Z. Yang and C. Zhang, “Sparsity-undersampling tradeoff of compressed sensing in the complex domain,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 3668–3671.
- [6] A. Maleki, L. Anitori, Z. Yang, and R. Baraniuk, “Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP),” *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4290–4308, 2013.
- [7] Z. Yang, C. Zhang, and L. Xie, “Robustly stable signal recovery in compressed sensing with structured matrix perturbation,” *IEEE Transactions on Signal Processing*, vol. 60, no. 9, pp. 4658–4671, 2012.

- [8] —, “Stable signal recovery in compressed sensing with a structured matrix perturbation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012, pp. 2737–2740.
- [9] Z. Yang, L. Xie, and C. Zhang, “Variational Bayesian algorithm for quantized compressed sensing,” *IEEE Transactions on Signal Processing*, vol. 61, no. 11, pp. 2815–2824, 2013.
- [10] —, “Accurate signal recovery in quantized compressed sensing,” in *15th IEEE International Conference on Information Fusion (FUSION)*, 2012, pp. 2531–2536.
- [11] —, “Off-grid direction of arrival estimation using sparse Bayesian inference,” *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 38–43, 2013.
- [12] Z. Yang, C. Zhang, and L. Xie, “Sparse MRI for motion correction,” in *10th IEEE International Symposium on Biomedical Imaging (ISBI)*, 2013, pp. 962–965.
- [13] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM journal on scientific computing*, vol. 20, no. 1, pp. 33–61, 1998.
- [14] J. F. Claerbout and F. Muir, “Robust modeling with erratic data,” *Geophysics*, vol. 38, no. 5, pp. 826–844, 1973.
- [15] D. L. Donoho and M. Elad, “Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization,” *Proceedings of the National Academy of Sciences*, vol. 100, no. 5, pp. 2197–2202, 2003.
- [16] E. Candès, “Compressive sampling,” in *Proceedings of the International Congress of Mathematicians*, vol. 3, 2006, pp. 1433–1452.
- [17] —, “The restricted isometry property and its implications for compressed sensing,” *Comptes Rendus Mathematique*, vol. 346, no. 9-10, pp. 589–592, 2008.

- [18] S. Foucart, “Sparse recovery algorithms: Sufficient conditions in terms of restricted isometry constants,” in *Approximation Theory XIII: San Antonio 2010*. Springer, 2012, pp. 65–77.
- [19] T. Cai, L. Wang, and G. Xu, “New bounds for restricted isometry constants,” *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4388–4394, 2010.
- [20] J. Tropp and A. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [21] D. Needell and J. Tropp, “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples,” *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [22] W. Dai and O. Milenkovic, “Subspace pursuit for compressive sensing signal reconstruction,” *IEEE Transactions on Information Theory*, vol. 55, no. 5, pp. 2230–2249, 2009.
- [23] T. Blumensath and M. Davies, “Iterative hard thresholding for compressed sensing,” *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, 2009.
- [24] J. A. Tropp, “Greed is good: Algorithmic results for sparse approximation,” *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [25] M. A. Davenport and M. B. Wakin, “Analysis of orthogonal matching pursuit using the restricted isometry property,” *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4395–4401, 2010.
- [26] Q. Mo and Y. Shen, “A remark on the restricted isometry property in orthogonal matching pursuit,” *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3654–3656, 2012.

- [27] T. Zhang, “Sparse recovery with orthogonal matching pursuit under RIP,” *IEEE Transactions on Information Theory*, vol. 57, no. 9, pp. 6215–6221, 2011.
- [28] M. Maleki, “Approximate message passing algorithms for compressed sensing,” Ph.D. dissertation, 2011.
- [29] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” *Available at <http://arxiv.org/abs/1011.3027>*, 2010.
- [30] J. D. Blanchard, C. Cartis, and J. Tanner, “Compressed sensing: How sharp is the restricted isometry property?” *SIAM review*, vol. 53, no. 1, pp. 105–125, 2011.
- [31] D. Donoho and J. Tanner, “Counting the faces of randomly-projected hypercubes and orthants, with applications,” *Discrete and Computational Geometry*, vol. 43, no. 3, pp. 522–541, 2010.
- [32] M. Stojnic, “Various thresholds for ℓ_1 -optimization in compressed sensing,” *Available online at <http://arxiv.org/abs/0907.3666>*, 2009.
- [33] D. Donoho, A. Maleki, and A. Montanari, “Message-passing algorithms for compressed sensing,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18 914–18 919, 2009.
- [34] D. Donoho and J. Tanner, “Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing,” *Philosophical Transactions of the Royal Society A*, vol. 367, no. 1906, pp. 4273–4293, 2009.
- [35] D. L. Donoho and J. Tanner, “Precise undersampling theorems,” *Proceedings of the IEEE*, vol. 98, no. 6, pp. 913–924, 2010.
- [36] M. Lustig, D. Donoho, and J. Pauly, “Sparse MRI: The application of compressed sensing for rapid MR imaging,” *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

- [37] M. Herman and T. Strohmer, “High-resolution radar via compressed sensing,” *IEEE Transactions on Signal Processing*, vol. 57, no. 6, pp. 2275–2284, 2009.
- [38] S. Ji, Y. Xue, and L. Carin, “Bayesian compressive sensing,” *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2346–2356, 2008.
- [39] M. Tipping, “Sparse Bayesian learning and the relevance vector machine,” *The Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [40] D. P. Wipf and B. D. Rao, “Sparse Bayesian learning for basis selection,” *IEEE Transactions on Signal Processing*, vol. 52, no. 8, pp. 2153–2164, 2004.
- [41] —, “An empirical Bayesian strategy for solving the simultaneous sparse approximation problem,” *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3704–3716, 2007.
- [42] D. Wipf and S. Nagarajan, “A unified Bayesian framework for MEG/EEG source imaging,” *Neuroimage*, vol. 44, no. 3, pp. 947–966, 2009.
- [43] S. Ji, D. Dunson, and L. Carin, “Multitask compressive sensing,” *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 92–106, 2009.
- [44] L. He and L. Carin, “Exploiting structure in wavelet-based Bayesian compressive sensing,” *IEEE Transactions on Signal Processing*, vol. 57, no. 9, pp. 3488–3497, 2009.
- [45] S. D. Babacan, L. Mancera, R. Molina, and A. K. Katsaggelos, “Non-convex priors in Bayesian compressed sensing,” in *17th European Signal Processing Conference (EUSIPCO)*, 2009, pp. 110–114.
- [46] S. Babacan, R. Molina, and A. Katsaggelos, “Bayesian compressive sensing using Laplace priors,” *IEEE Transactions on Image Processing*, vol. 19, no. 1, pp. 53–63, 2010.

- [47] D. Baron, S. Sarvotham, and R. Baraniuk, "Bayesian compressive sensing via belief propagation," *IEEE Transactions on Signal Processing*, vol. 58, no. 1, pp. 269–280, 2010.
- [48] D. Shutin and B. H. Fleury, "Sparse variational Bayesian SAGE algorithm with application to the estimation of multipath wireless channels," *IEEE Transactions on Signal Processing*, vol. 59, no. 8, pp. 3609–3623, 2011.
- [49] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 5, pp. 912–926, 2011.
- [50] L. Yu, H. Sun, J.-P. Barbot, and G. Zheng, "Bayesian compressive sensing for cluster structured sparse signals," *Signal Processing*, vol. 92, no. 1, pp. 259–269, 2012.
- [51] K. E. Themelis, A. A. Rontogiannis, and K. D. Koutroumbas, "A novel hierarchical Bayesian approach for sparse semisupervised hyperspectral unmixing," *IEEE Transactions on Signal Processing*, vol. 60, no. 2, pp. 585–599, 2012.
- [52] H. Ishwaran and J. S. Rao, "Spike and slab variable selection: Frequentist and Bayesian strategies," *Annals of Statistics*, pp. 730–773, 2005.
- [53] C. P. Robert and G. Casella, *Monte Carlo statistical methods*. Citeseer, 2004, vol. 319.
- [54] N. L. Pedersen, D. Shutin, C. N. Manchón, and B. H. Fleury, "Sparse estimation using Bayesian hierarchical prior modeling for real and complex models," Available at <http://arxiv.org/pdf/1108.4324v1>, 2011.
- [55] D. J. MacKay, "Bayesian interpolation," *Neural computation*, vol. 4, no. 3, pp. 415–447, 1992.
- [56] M. J. Beal, "Variational algorithms for approximate Bayesian inference," Ph.D. dissertation, 2003.

- [57] M. Tipping and A. Faul, “Fast marginal likelihood maximisation for sparse Bayesian models,” in *9th International Workshop on Artificial Intelligence and Statistics*, vol. 1, no. 3, 2003.
- [58] E. Candès, J. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on pure and applied mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [59] L. Jacques, D. K. Hammond, and J. M. Fadili, “Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine,” *IEEE Transactions on Information Theory*, vol. 57, no. 1, pp. 559–571, 2011.
- [60] J. N. Laska, P. T. Boufounos, M. A. Davenport, and R. G. Baraniuk, “Democracy in action: Quantization, saturation, and compressive sensing,” *Applied and Computational Harmonic Analysis*, vol. 31, no. 3, pp. 429–443, 2011.
- [61] A. Zymnis, S. Boyd, and E. Candes, “Compressed sensing with quantized measurements,” *IEEE Signal Processing Letters*, vol. 17, no. 2, pp. 149–152, 2010.
- [62] P. T. Boufounos and R. G. Baraniuk, “1-bit compressive sensing,” in *42nd IEEE Annual Conference on Information Sciences and Systems (CISS)*, 2008, pp. 16–21.
- [63] J. N. Laska and R. G. Baraniuk, “Regime change: Bit-depth versus measurement-rate in compressive sensing,” *IEEE Transactions on Signal Processing*, vol. 60, no. 7, pp. 3496–3505, 2012.
- [64] P. T. Boufounos, “Greedy sparse signal reconstruction from sign measurements,” in *IEEE Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers*, 2009, pp. 1305–1309.
- [65] J. N. Laska, Z. Wen, W. Yin, and R. G. Baraniuk, “Trust, but verify: Fast and accurate signal recovery from 1-bit compressive measurements,” *IEEE Transactions on Signal Processing*, vol. 59, no. 11, pp. 5289–5301, 2011.

- [66] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, “Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors,” *IEEE Transactions on Information Theory*, vol. 59, no. 4, pp. 2082–2102, 2013.
- [67] Y. Plan and R. Vershynin, “One-bit compressed sensing by linear programming,” *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1275–1297, 2013.
- [68] Y. PLAN and R. VERSHYNIN, “Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach,” *IEEE Transactions on Information Theory*, vol. 59, no. 1, pp. 482–494, 2013.
- [69] T. Blumensath and M. Davies, “Compressed sensing and source separation,” in *International Conference on Independent Component Analysis and Signal Separation*, 2007, pp. 341–348.
- [70] T. Xu and W. Wang, “A compressed sensing approach for underdetermined blind audio source separation with sparse representation,” in *15th IEEE Workshop on Statistical Signal Processing*, 2009, pp. 493–496.
- [71] J. Zheng and M. Kaveh, “Directions-of-arrival estimation using a sparse spatial spectrum model with uncertainty,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 2848–2551.
- [72] L. Zhang, M. Xing, C. Qiu, J. Li, and Z. Bao, “Achieving higher resolution isar imaging with limited pulses via compressed sampling,” *Geoscience and Remote Sensing Letters*, vol. 6, no. 3, pp. 567–571, 2009.
- [73] S. Gleichman and Y. C. Eldar, “Blind compressed sensing,” *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6958–6975, 2011.
- [74] M. Herman and T. Strohmer, “General deviants: An analysis of perturbations in compressed sensing,” *IEEE J. Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 342–349, 2010.

- [75] M. Herman and D. Needell, "Mixed operators in compressed sensing," in *44th IEEE Annual Conference on Information Sciences and Systems (CISS)*, 2010, pp. 1–6.
- [76] Y. Chi, L. Scharf, A. Pezeshki, and A. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2182–2195, 2011.
- [77] D. Chae, P. Sadeghi, and R. Kennedy, "Effects of basis-mismatch in compressive sampling of continuous sinusoidal signals," in *2nd IEEE International Conference on Future Computer and Communication (ICFCC)*, vol. 2, 2010, pp. 739–743.
- [78] H. Zhu, G. Leus, and G. Giannakis, "Sparsity-cognizant total least-squares for perturbed compressive sampling," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2002–2016, 2011.
- [79] M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, and R. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 83–91, 2008.
- [80] J. Provost and F. Lesage, "The application of compressed sensing for photoacoustic tomography," *IEEE Transactions on Medical Imaging*, vol. 28, no. 4, pp. 585–594, 2009.
- [81] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressed channel sensing: A new approach to estimating sparse multipath channels," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1058–1076, 2010.
- [82] M. Mishali and Y. C. Eldar, "Wideband spectrum sensing at sub-Nyquist rates [applications corner]," *IEEE Signal Processing Magazine*, vol. 28, no. 4, pp. 102–135, 2011.
- [83] F. Parvaresh, H. Vikalo, S. Misra, and B. Hassibi, "Recovering sparse signals using sparse measurement matrices in compressed DNA microarrays," *IEEE*

- Journal of Selected Topics in Signal Processing*, vol. 2, no. 3, pp. 275–285, 2008.
- [84] H. Krim and M. Viberg, “Two decades of array signal processing research: The parametric approach,” *IEEE Signal Processing Magazine*, vol. 13, no. 4, pp. 67–94, 1996.
- [85] R. R. Edelman and S. Warach, “Magnetic resonance imaging,” *New England Journal of Medicine*, vol. 328, no. 10, pp. 708–716, 1993.
- [86] R. Schmidt, “A signal subspace approach to multiple emitter location spectral estimation,” Ph.D. dissertation, 1981.
- [87] P. Stoica and N. Arye, “Music, maximum likelihood, and Cramer-Rao bound,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, no. 5, pp. 720–741, 1989.
- [88] D. Malioutov, M. Cetin, and A. Willsky, “A sparse signal reconstruction perspective for source localization with sensor arrays,” *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 3010–3022, 2005.
- [89] M. Hyder and K. Mahata, “Direction-of-arrival estimation using a mixed $\ell_{2,0}$ norm approximation,” *IEEE Transactions on Signal Processing*, vol. 58, no. 9, pp. 4646–4655, 2010.
- [90] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, “Compressed sensing MRI,” *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, 2008.
- [91] H. Jung, K. Sung, K. S. Nayak, E. Y. Kim, and J. C. Ye, “k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI,” *Magnetic Resonance in Medicine*, vol. 61, no. 1, pp. 103–116, 2009.
- [92] R. Otazo, D. Kim, L. Axel, and D. K. Sodickson, “Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI,” *Magnetic Resonance in Medicine*, vol. 64, no. 3, pp. 767–776, 2010.

- [93] J. P. Haldar, D. Hernando, and Z.-P. Liang, "Compressed-sensing MRI with random encoding," *IEEE Transactions on Medical Imaging*, vol. 30, no. 4, pp. 893–903, 2011.
- [94] V. M. Patel, R. Maleh, A. C. Gilbert, and R. Chellappa, "Gradient-based image recovery methods from incomplete Fourier measurements," *IEEE Transactions on Image Processing*, vol. 21, no. 1, pp. 94–105, 2012.
- [95] M. Medley, H. Yan, and D. Rosenfeld, "An improved algorithm for 2-D translational motion artifact correction," *IEEE Transactions on Medical Imaging*, vol. 10, no. 4, pp. 548–553, 1991.
- [96] R. A. Zoroofi, Y. Sato, S. Tamura, H. Naito, and L. Tang, "An improved method for MRI artifact correction due to translational motion in the imaging plane," *IEEE Transactions on Medical Imaging*, vol. 14, no. 3, pp. 471–479, 1995.
- [97] Y. M. Kadah, A. A. Abaza, A. S. Fahmy, A.-B. M. Youssef, K. Heberlein, and X. P. Hu, "Floating navigator echo (FNAV) for in-plane 2D translational motion estimation," *Magnetic resonance in medicine*, vol. 51, no. 2, pp. 403–407, 2004.
- [98] W. Lin, F. Huang, P. Börnert, Y. Li, and A. Reykowski, "Motion correction using an enhanced floating navigator and GRAPPA operations," *Magnetic Resonance in Medicine*, vol. 63, no. 2, pp. 339–348, 2010.
- [99] J. Mendes, E. Kholmovski, and D. L. Parker, "Rigid-body motion correction with self-navigation MRI," *Magnetic Resonance in Medicine*, vol. 61, no. 3, pp. 739–747, 2009.
- [100] E. Candès and J. Romberg, " ℓ_1 -magic: Recovery of sparse signals via convex programming," *Avaiable online at <http://users.ece.gatech.edu/justin/l1magic/downloads/l1magic.pdf>*, 2005.

- [101] S. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky, “An interior-point method for large-scale ℓ_1 -regularized least squares,” *IEEE J. Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 606–617, 2008.
- [102] E. Hale, W. Yin, and Y. Zhang, “A fixed-point continuation method for ℓ_1 -regularized minimization with applications to compressed sensing,” *Available online at <http://www.caam.rice.edu/zhang/reports/tr0707.pdf>*, 2007.
- [103] S. Becker, J. Bobin, and E. Candes, “NESTA: A fast and accurate first-order method for sparse recovery,” *SIAM J. Imaging Sciences*, vol. 4, no. 1, pp. 1–39, 2011.
- [104] Y. Nesterov, “Smooth minimization of non-smooth functions,” *Mathematical Programming*, vol. 103, no. 1, pp. 127–152, 2005.
- [105] J. Nocedal and S. Wright, *Numerical Optimization*. New York: Springer Verlag, 2006.
- [106] D. Bertsekas, *Constrained optimization and Lagrange multiplier methods*. Boston: Academic Press, 1982.
- [107] Z. Wen, W. Yin, D. Goldfarb, and Y. Zhang, “A fast algorithm for sparse reconstruction based on shrinkage, subspace optimization and continuation,” *SIAM Journal on Scientific Computing*, vol. 32, no. 4, pp. 1832–1857, 2010.
- [108] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [109] K. Bredies and D. Lorenz, “Linear convergence of iterative soft-thresholding,” *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 813–837, 2008.
- [110] A. Maleki and D. L. Donoho, “Optimally tuned iterative reconstruction algorithms for compressed sensing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 330–341, 2010.

- [111] M. Stojnic, "Block-length dependent thresholds in block-sparse compressed sensing," *Available online at <http://arxiv.org/abs/0907.3679>*, 2009.
- [112] M. Stojnic, F. Parvaresh, and B. Hassibi, "On the reconstruction of block-sparse signals with an optimal number of measurements," *IEEE Transactions on Signal Processing*, vol. 57, no. 8, pp. 3075–3085, 2009.
- [113] Y. Eldar and H. Rauhut, "Average case analysis of multichannel sparse recovery using convex relaxation," *IEEE Transactions on Information Theory*, vol. 56, no. 1, pp. 505–519, 2010.
- [114] S. Foucart and R. Gribonval, "Real versus complex null space properties for sparse vector recovery," *Comptes Rendus Mathematique*, vol. 348, no. 15, pp. 863–865, 2010.
- [115] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
- [116] S. Foucart and M. Lai, "Sparsest solutions of underdetermined linear systems via ℓ_1 -minimization for $0 < q < 1$," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 395–407, 2009.
- [117] T. Cai, L. Wang, and G. Xu, "Shifting inequality and recovery of sparse signals," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1300–1308, 2010.
- [118] D. Donoho, M. Elad, and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Transactions on Information Theory*, vol. 52, no. 1, pp. 6–18, 2006.
- [119] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1982–2001, 2010.

- [120] T. Blumensath and M. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Transactions on Information Theory*, vol. 55, no. 4, pp. 1872–1882, 2009.
- [121] R. Chartrand, "Exact reconstruction of sparse signals via nonconvex minimization," *IEEE Signal Processing Letters*, vol. 14, no. 10, pp. 707–710, 2007.
- [122] R. Chartrand and V. Staneva, "Restricted isometry properties and nonconvex compressive sensing," *Inverse Problems*, vol. 24, p. 035020, 2008.
- [123] R. Saab, R. Chartrand, and O. Yilmaz, "Stable sparse approximations via non-convex optimization," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2008, pp. 3885–3888.
- [124] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming," *Available online at <http://cvxr.com/cvx>*, 2008.
- [125] E. J. Candes and Y. Plan, "A probabilistic and RIPless theory of compressed sensing," *IEEE Transactions on Information Theory*, vol. 57, no. 11, pp. 7235–7254, 2011.
- [126] G. McLachlan and T. Krishnan, *The EM algorithm and extensions*. John Wiley & Sons, 1997.
- [127] R. S. Irving, *Integers, polynomials, and rings: A course in algebra*. Springer, 2004.
- [128] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5-6, pp. 877–905, 2008.
- [129] D. Donoho, Y. Tsaig, I. Drori, and J. Starck, "Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit," *Available online at <http://www.cs.tau.ac.il/~idrori/StOMP.pdf>*, 2006.

- [130] E. Van Den Berg and M. P. Friedlander, “Probing the Pareto frontier for basis pursuit solutions,” *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.
- [131] Y. Tsaig and D. L. Donoho, “Extensions of compressed sensing,” *Signal Processing*, vol. 86, no. 3, pp. 549–571, 2006.
- [132] W. Dai and O. Milenkovic, “Information theoretical and algorithmic approaches to quantized compressive sensing,” *IEEE Transactions on Communications*, vol. 59, no. 7, pp. 1857–1866, 2011.
- [133] R. Tibshirani, “Regression shrinkage and selection via the Lasso,” *Journal of the Royal Statistical Society. Series B*, vol. 58, no. 1, pp. 267–288, 1996.
- [134] T. Park and G. Casella, “The Bayesian lasso,” *Journal of the American Statistical Association*, vol. 103, no. 482, pp. 681–686, 2008.
- [135] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, “An introduction to variational methods for graphical models,” *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [136] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, “The variational approximation for Bayesian inference,” *IEEE Signal Processing Magazine*, vol. 25, no. 6, pp. 131–146, 2008.
- [137] J. Winn and C. Bishop, “Variational message passing,” *Journal of Machine Learning Research*, vol. 6, pp. 661–694, 2005.
- [138] B. Jørgensen, *Statistical properties of the generalized inverse Gaussian distribution*. Springer New York, 1982, vol. 21.
- [139] J. M. Kim, O. K. Lee, and J. C. Ye, “Compressive MUSIC: Revisiting the link between compressive sensing and array signal processing,” *IEEE Transactions on Information Theory*, vol. 58, no. 1, pp. 278–301, 2012.

- [140] K. Lee, Y. Bresler, and M. Junge, “Subspace methods for joint sparse recovery,” *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3613–3641, 2012.
- [141] N. Goodman, “Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction),” *The Annals of Mathematical Statistics*, vol. 34, no. 1, pp. 152–177, 1963.
- [142] S. Marchesini, “A unified evaluation of iterative projection algorithms for phase retrieval,” *Review of Scientific Instruments*, vol. 78, no. 1, pp. 011 301–011 301, 2007.
- [143] R. Millane, “Phase retrieval in crystallography and optics,” *JOSA A*, vol. 7, no. 3, pp. 394–411, 1990.
- [144] D. R. Luke, “Relaxed averaged alternating reflections for diffraction imaging,” *Inverse Problems*, vol. 21, no. 1, p. 37, 2005.
- [145] Z. Yang, C. Zhang, and L. Xie, “Robust compressive phase retrieval via ℓ_1 minimization with application to image reconstruction,” *Available online at <http://arxiv.org/abs/1302.0081>*, 2013.
- [146] B. N. Bhaskar, G. Tang, and B. Recht, “Atomic norm denoising with applications to line spectral estimation,” *IEEE Transactions on Signal Processing*, DOI: 10.1109/TSP.2013.2273443, 2013.
- [147] E. J. Candès and C. Fernandez-Granda, “Towards a mathematical theory of super-resolution,” *Communications on Pure and Applied Mathematics*, DOI: 10.1002/cpa.21455, 2013.
- [148] G. Tang, B. N. Bhaskar, and B. Recht, “Near minimax line spectral estimation,” *Available online at <http://arxiv.org/abs/1303.4348>*, 2013.
- [149] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, “The convex geometry of linear inverse problems,” *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, 2012.

- [150] Z. Yang, L. Xie, and C. Zhang, “A discretization-free sparse and parametric approach for linear array signal processing,” *Submitted to IEEE Transactions on Signal Processing*, July 2013.
- [151] Z. Yang and L. Xie, “On sparse parametric methods for line spectral estimation from complete and incomplete data,” *Submitted to IEEE Transactions on Signal Processing*, November 2013.
- [152] M. Moravec, J. Romberg, and R. Baraniuk, “Compressive phase retrieval,” in *Proceedings of SPIE, the International Society for Optical Engineering*, 2007, pp. 670 120–1.
- [153] C. Zhang, Z. Yang, and L. Xie, “The l_1 measure of image reconstruction subject to motion corrupted k -space data,” in *10th IEEE International Conference on Control and Automation (ICCA)*, 2013, pp. 1133–1137.