

Fall 2013

Novel adaptive reconstruction schemes for accelerated myocardial perfusion magnetic resonance imaging

Sajan Goud Lingala
University of Iowa

Copyright 2013 Sajan Goud Lingala

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/5016>

Recommended Citation

Lingala, Sajan Goud. "Novel adaptive reconstruction schemes for accelerated myocardial perfusion magnetic resonance imaging." PhD (Doctor of Philosophy) thesis, University of Iowa, 2013.
<https://doi.org/10.17077/etd.7pl7tyno>

Follow this and additional works at: <https://ir.uiowa.edu/etd>

Part of the [Biomedical Engineering and Bioengineering Commons](#)

NOVEL ADAPTIVE RECONSTRUCTION SCHEMES FOR ACCELERATED
MYOCARDIAL PERFUSION MAGNETIC RESONANCE IMAGING

by

Sajan Goud Lingala

A thesis submitted in partial fulfillment of the
requirements for the Doctor of Philosophy degree
in Biomedical Engineering
in the Graduate College of
The University of Iowa

December 2013

Thesis Supervisor: Assistant Professor Mathews Jacob

Copyright by
SAJAN GOUD LINGALA
2013
All Rights Reserved

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Sajan Goud Lingala

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Biomedical Engineering at the December 2013 graduation.

Thesis Committee: _____

Mathews Jacob, Thesis Supervisor

Vincent Magnotta

Daniel Thedens

Joseph Reinhardt

Weiyu Xu

Dedicated to my parents

ACKNOWLEDGEMENTS

I would like to thank my advisor Dr. Mathews Jacob for his excellent guidance and providing me opportunities to work on challenging and exciting projects. I sincerely thank his support all through these years. His intuition on approaching research problems has instigated in me tremendous liking towards research, and I now hope to continue pursuing new open challenging problems in my professional research career.

I was also fortunate to collaborate with Dr. Edward DiBella and Dr. Ganesh Adluru from the University of Utah all through my thesis. Their support on myocardial perfusion MRI sequences, and insightful advices on many topics were critical towards translating many of our algorithmic ideas towards clinical settings.

I greatly thank Drs. John Newell, Alan Stolpen, and Chris McGann for their clinical collaborations.

I would also like to thank the members of my doctoral committee - Professors Joe Reinhardt, Vincent Magnotta, Daniel Thedens, Weiyu Xu - for their time, effort, and help, and for their influence on my training and research.

I am also thankful to the past and present lab mates in the Computational Biomedical Imaging Group - Dr. Ramin Eslami, Satyananda Kashyap, Zhili Yang, Yue Hu, Xuan Zhou, Greg Ongie, Merry Mani, Sampada Bhave, Chen Cui, Ipshita Bhattacharya, Arvind Balachandrasekaran, Sunrita Poddar, Yasir Baqqal for the friendly and stimulating work environment.

I am also very appreciative of Marla Kleingartner, and Autumn Craig, the MR research technologists at the University of Iowa Magnetic Resonance Research Facility for their patient help with the scans.

I would like to acknowledge financial support of my graduate studies from American Heart Association, National Institute of Health, and National Science Foundation.

Finally, nothing can express about the unconditional love and support that I have received from my family. My mother always put the children ahead of herself, she

ensured that we obtained the finest of education in difficult times. I also express my love to my sister, Rachita and my twin brother, Sujan; they are the best siblings one could hope for. There are always there to encourage and support me by all means.

ABSTRACT

Coronary artery disease (CAD) is one of the leading causes of death in the world. In the United States alone, it is estimated that approximately every 25 seconds, a new CAD event will occur, and approximately every minute, someone will die of one. The detection of CAD in its early stages is very critical to reduce the mortality rates. Magnetic resonance imaging of myocardial perfusion (MR-MPI) has been receiving significant attention over the last decade due to its ability to provide a unique view of the microcirculation blood flow in the myocardial tissue through the coronary vascular network. The ability of MR-MPI to detect changes in microcirculation during early stages of ischemic events makes it a useful tool in identifying myocardial tissues that are alive but at the risk of dying. However this technique is not yet fully established in the clinic due to fundamental limitations imposed by the MRI device physics in terms of slow imaging speed. The limitations of current MRI schemes often make it challenging to simultaneously achieve high spatio-temporal resolution, sufficient spatial coverage, and good image quality in myocardial perfusion MRI. Furthermore, the acquisitions are typically set up to acquire images during breath holding. This often results in motion artifacts due to improper breath hold patterns. This also limits its applicability to a large domain of patient populations such as those with impaired respiratory function, arrhythmias, pediatrics.

The overall objective of this thesis is to develop a novel dynamic imaging framework that can enable free breathing myocardial perfusion imaging with high spatio-temporal resolutions and close to whole heart volume coverage. To achieve this, this thesis deals with developing novel image reconstruction methods for the reconstruction of dynamic MRI data from highly accelerated / under-sampled Fourier measurements. It specifically focuses on novel blind or adaptive image models that represent the dynamic image data set using adaptive temporal bases (bases derived from the data at hand). This is in sharp contrast to classical models that rely on predetermined temporal bases (such as Fourier bases), which require assumptions such as temporal

periodicity and breath holding that are clearly not met in free breathing MR-MPI.

Three novel adaptive reconstruction methods are developed and validated: (a) low rank and sparsity based modeling (k-t SLR), (b) blind compressed sensing (BCS), and (c) motion compensated compressed sensing (MC-CS). The developed methods are applicable to a wide range of dynamic imaging problems. In the context of MR-MPI, this dissertation show feasibilities that the developed methods can enable free breathing myocardial perfusion MRI acquisitions with high spatio-temporal resolutions ($< 2\text{mm} \times 2\text{mm}^2$, 1 heart beat) and high slice coverage.

TABLE OF CONTENTS

LIST OF TABLES	x
LIST OF FIGURES	xi
CHAPTER	
1 INTRODUCTION	1
1.1 Motivation	2
1.2 Challenges with existing acceleration schemes	5
1.3 Overall objective	9
1.4 Main contributions of the thesis	9
1.5 Outline of the thesis	13
2 ACCELERATED DYNAMIC MRI BY EXPLOITING LOW RANK AND SPARSE STRUCTURE (K-T SLR)	15
2.1 Introduction	15
2.2 Dynamic MRI using KLT	18
2.2.1 Matrix recovery using nuclear norm minimization	20
2.3 k - t SLR: Formulation	21
2.3.1 Regularized matrix recovery using spectral priors	22
2.3.2 Regularized matrix recovery using spectral and sparsity priors	23
2.4 Optimization Algorithm	25
2.4.1 Implementation	29
2.5 Materials and Methods	30
2.5.1 Datasets	30
2.5.2 Comparisons against different methods	31
2.6 Results	34
2.6.1 Convergence of the algorithm	34
2.6.2 Comparisons on the PINCAT phantom	35
2.6.3 Comparisons on the invivo data	40
2.7 Discussion	41
3 ACCELERATING FREE BREATHING MYOCARDIAL PERFUSION MRI USING MULTI COIL RADIAL K-T SLR	44
3.1 Introduction	44
3.2 Theory	44
3.3 Low rank model representation	44
3.3.1 Radial $k - t$ SLR with parallel imaging	47
3.3.2 Fast augmented Lagrangian (AL) algorithm	49
3.4 Materials and Methods	51

3.4.1	Multi-coil radial acquisition of free breathing myocardial perfusion data	51
3.4.2	Undersampled reconstruction with different algorithms	52
3.4.3	Simulations to determine an optimal radial sampling trajectory	53
3.4.4	Metrics used for quantitative comparison	54
3.4.5	Qualitative evaluation: clinical scoring	55
3.5	RESULTS	56
3.5.1	Convergence analysis	56
3.5.2	Simulations to determine an optimal radial trajectory	59
3.5.3	$k - t$ SLR compared to other MPI acceleration schemes	59
3.5.4	Qualitative evaluation by a cardiologist	63
3.5.5	Prospectively undersampled radial data:	64
3.6	Discussion	65
3.7	Appendix	69
3.7.1	Augmented Lagrangian (AL) algorithm steps	69
3.7.2	Choosing the regularization parameters	70
4	BLIND COMPRESSIVE SENSING DYNAMIC MRI	73
4.1	Introduction	73
4.2	Dynamic MRI reconstruction using the BCS model	77
4.2.1	Dynamic image acquisition	77
4.2.2	The BCS representation	78
4.2.3	The objective function	79
4.2.4	The optimization algorithm	80
4.3	Experimental Evaluation	84
4.3.1	Comparison of different BCS schemes	85
4.3.2	Choice of parameters	92
4.3.3	Numerical simulations	94
4.3.4	Experiments on invivo datasets	96
4.4	Towards high spatio-temporal resolution 3D DMRI of lung using Blind CS	102
4.4.1	Preliminaries on 3D dynamic lung MRI	102
4.4.2	Retrospective undersampling of a 2D acquisition	103
4.4.3	Prospective 3D undersampling with stack of spokes GA radial acquisition	103
4.4.4	Results and discussion on dynamic lung MRI	104
4.5	Discussion on Blind compressed sensing	105
5	A NOVEL FRAMEWORK FOR MOTION COMPENSATED COMPRESSED SENSING MYOCARDIAL PERFUSION MRI	110
5.1	Introduction	110
5.2	Motion compensated compressed sensing (MC-CS)	113

5.2.1	Dynamic image acquisition	113
5.2.2	MC-CS model	113
5.3	MC-CS: Optimization algorithm	114
5.3.1	g sub-problem (Temporal denoising/dealiasing)	115
5.3.2	θ sub-problem (Motion estimation)	116
5.3.3	f sub-problem (reconstruction update)	118
5.3.4	Continuation strategy to alternate between the sub-problems	118
5.3.5	Convergence analysis	120
5.4	Materials and Methods	122
5.4.1	Datasets and $k - t$ sampling	122
5.4.2	Comparison of MC-CS with its CS variants	126
5.5	Results	127
5.6	Discussion	131
6	SUMMARY AND FUTURE WORK	134
6.1	Summary	134
6.2	Future work	136
	REFERENCES	138

LIST OF TABLES

Table

3.1 Quality scores from a cardiologist on three subjects 65

LIST OF FIGURES

Figure		
1.1	The ischemic cascade	2
1.2	Dynamic imaging acquisition set up of a myocardial perfusion MRI experiment	3
1.3	Manifestation of Gibbs ringing as dark rim artifacts	4
1.4	Dynamic data representation in the temporal Fourier transform (x-f) domain	7
1.5	Example reconstruction of the k-t SPARSE method (compressed sensing based on x-f sparsity) on free breathing MR-MPI data	8
2.1	Utility of KLT in compactly representing the dynamic image time series:	16
2.2	The PINCAT phantom used to validate the proposed scheme.	30
2.3	Convergence of the proposed continuation scheme	35
2.4	Decay of SER as a function of acceleration on the PINCAT data	36
2.5	Comparison of the two-step KLT schemes (two top rows) with the spectrally regularized reconstruction scheme	37
2.6	Performance evaluation of k-t SLR in comparison with different schemes on the PINCAT phantom:	38
2.7	Quantitative comparison of the different schemes at a range of accelerations on in-vivo data.	39
2.8	Comparisons on in-vivo data	43
3.1	The low rank Casorati matrix representation of dynamic data:	45
3.2	$k - t$ SLR with parallel MRI for accelerated imaging:	48
3.3	Convergence analysis:	56
3.4	Performance of different radial sampling schemes:	57
3.5	Comparisons of different MPI algorithms on a rest dataset with breathing motion	58
3.6	Example comparisons of different MPI acceleration algorithms using rest (i) and stress (ii) perfusion data from a patient with myocardial ischemia	60

3.7	Quantitative comparison of reconstructions from undersampled radial data (21 rays/frame) using different algorithms.	61
3.8	Peak myocardial wall enhancement images from (a) reference (72 rays/frame), and (b) $k - t$ SLR (24 rays/frame) reconstructions.	62
3.9	Example $k-t$ SLR reconstructed multi slice 2D first-pass myocardial perfusion images on a healthy subject.	66
3.10	The augmented Lagrangian frame work with the different sub problems.	71
3.11	Tuning of the regularization parameters λ_1 and λ_2 . The SER was evaluated in a field of view containing the regions of the heart. The optimal parameters were chosen corresponding to the region where the SER between the reconstruction and the reference data set was maximum.	72
4.1	Comparison of blind compressed sensing (BCS) and blind linear model (BLM) representations of dynamic imaging data:	74
4.2	Comparison of different BCS schemes:	81
4.3	Model coefficients and dictionary bases.	88
4.4	Blind CS model dependence on the regularization parameter and the dictionary size:	89
4.5	The numerical phantoms \mathbf{F}^j , which are used in the simulation study	90
4.6	Phase transition behavior of various reconstruction schemes:	91
4.7	Comparison of the proposed scheme with different methods on a retrospectively downsampled Cartesian myocardial perfusion data set with motion at 7.5 fold acceleration:	98
4.8	Comparisons of the different reconstructions schemes on a brain perfusion MRI dataset. The fully sampled data in (a) is retrospectively undersampled at a high acceleration of 10.66.	99
4.9	Comparisons of different reconstruction schemes on a stress myocardial perfusion MRI dataset with breathing motion:	100
4.10	Comparisons on 2D dynamic lung data with retrospective undersampling:	108
4.11	Example 3D dynamic images obtained with BCS using 16 spokes/frame with 16 slices, at spatial resolutions of $2.37 \times 2.37 \times 4 \text{mm}^3$, and time resolution of 0.72sec.	109

5.1	Free breathing myocardial perfusion MRI data representation in transform domains with and without motion compensation:	112
5.2	Continuation over the penalty parameter β :	124
5.3	Continuation over α : It can be seen that the continuation over the deformable force strength parameter α results in a faster convergence, and also avoids the cost to diverge in the last iterations.	125
5.4	Evolution of the region of interest signal to error ratio as the iterations proceed. The algorithm was found to be more or less robust to the choice of the initial guess of the reconstruction.	125
5.5	Qualitative comparison on the PINCAT phantom data using 20 rays/frame:	128
5.6	Quantitative comparisons of different schemes on PINCAT data	129
5.7	The fully sampled myocardial perfusion MRI dataset used as a ground-ruth during the retrospective undersampling experiments. Note the ripples in the image time profile correspond to large inter-frame motion.	129
5.8	Performance evaluation using retrospectively sampling on Cartesian data using 16 radial rays/frame	130
5.9	Quantitative comparisons of different schemes on invivo data	131
5.10	Comparison on prospectively undersampled radial data using 21 rays/frame. Here single coil reconstructions are shown on two subjects. It can be seen that MC-CS resulted in lesser motion artifacts than CS schemes in both the subjects (see the arrows that depict motion blur in CS).	132

CHAPTER 1 INTRODUCTION

Coronary artery disease (CAD) is one of the leading causes of death in the world, and is the leading cause of death in the United States. According to the American Heart Association, CAD was responsible for approximately 1 of every 6 deaths in the United States in 2008. It is estimated that in each year, approximately 7.85 million Americans will experience a new CAD attack; approximately every 25 seconds, a CAD event will occur, and approximately every minute, someone will die of one [1]. There are many available treatments for CAD including lifestyle changes (diet, exercise, stress, smoking), surgical interventions, and pharmaceuticals. However, detecting CAD in its early stages is very critical to reduce the mortality.

CAD or ischemic heart disease is the narrowing of the coronary arteries, the vessels that supply blood to the myocardium. The narrowing of the arteries results in a series of biochemical reactions/events in which an event triggers the next until the patient experiences an angina. The series of events termed as ischemic cascade starts with altered vascular dysfunction, causing reduced blood flow to the myocardium, followed by altered metabolism, followed by wall motion abnormalities, and impairment of systolic and diastolic function, followed by ECG changes, and finally resulting in angina (figure. 1.1). Changes in myocardial perfusion which occur early in the ischemic cascade serve as sensitive indicators of ischemic conditions.

Myocardial perfusion imaging is a non-invasive way to assess blood flow. It is a promising alternative to the invasive X-ray coronary angiography procedure which is the current gold standard for the diagnosis of ischemic heart disease. Presently the modality of single photon emission computed tomography (SPECT) is the most frequently used in the clinic. However SPECT involves ionizing radiation and suffers from poor spatial resolution, and attenuation artifacts [4].

Magnetic resonance imaging of myocardial perfusion (MRI-MPI) is an attractive

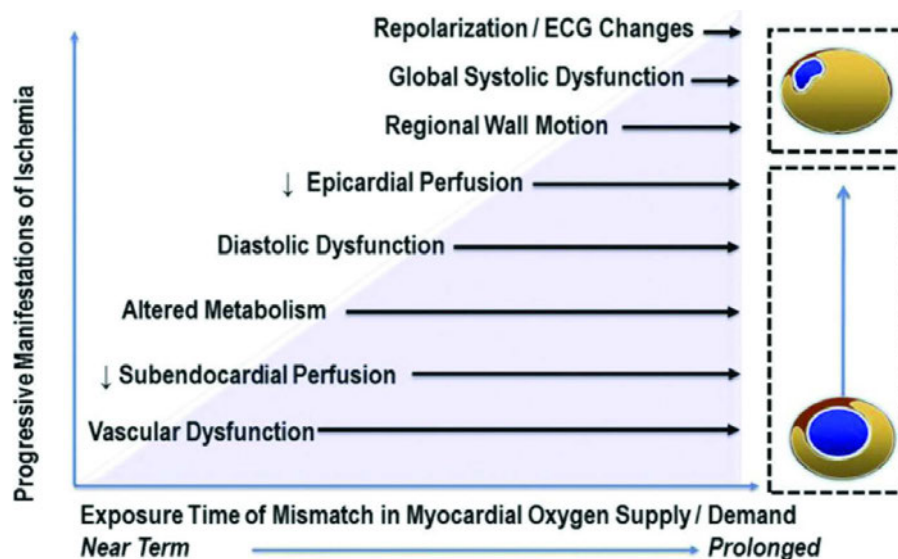


Figure 1.1: The ischemic cascade: The narrowing of coronary arteries causes changes in the vascular dysfunction leading to reduced myocardial perfusion which leads to altered metabolism, wall motion abnormalities, systolic dysfunction, ECG changes, and ultimately angina. This thesis focuses on developing non-invasive MRI methods to improve the detection of myocardial perfusion changes that aid in the early detection of coronary artery disease. (Schematic adapted from [2]).

imaging modality over SPECT due to its non-invasive nature, usefulness in longitudinal studies, and capability of providing superior spatial resolutions. The basic idea of MR-MPI is to track the dynamic variations of a paramagnetic contrast agent (typically Gadolinium) as it traverses through the different regions of the heart. It could be used to identify myocardial regions of reduced uptake / or reduced perfusion, and hence be able to identify the myocardial tissues that are alive but at the risk of dying.

1.1 Motivation

Despite its advantages, MR imaging of myocardial perfusion has not yet been fully established in the clinic due to fundamental limitations imposed by the MRI device physics. In MRI, data is acquired in the Fourier space (k-space), and the Fourier encoding process is inherently slow due to limitations on the gradient amplitudes and switching rates that could risk in resulting peripheral nerve stimulation [5], [6]. The

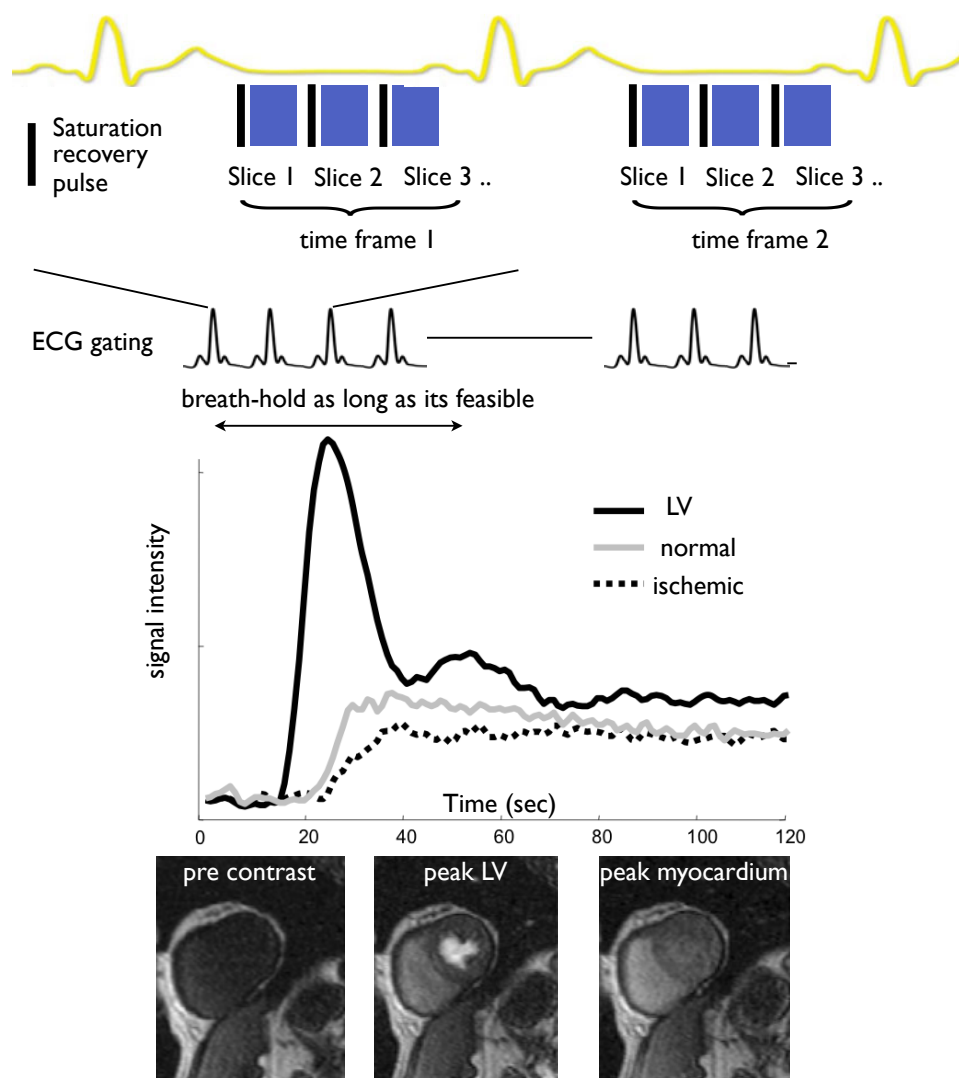


Figure 1.2: Dynamic imaging acquisition set up of a myocardial perfusion MRI experiment: An ECG gated saturation recovery sequence is used to acquire multiple slices within every heart-beat while the subject is breath-holding. After the injection of Gadolinium, the contrast first passes through the right ventricle, then enters the left ventricle, and finally the myocardium. The perfusion uptake curves in the different regions of the myocardium give an indication of any underlying ischemic conditions.




Figure 1.3: Manifestation of Gibbs ringing as dark rim artifacts: To demonstrate the relation between Gibbs ringing and the formation of the dark rim artifact (DRA), an experimental depiction from DiBella et al. (2005) [3] is shown. MR-MPI images are shown from an ex-vivo dog heart scanned at different resolutions along the phase encode direction. It is clearly seen that at low resolutions, the sharp intensity profile distribution across the myocardium and the blood pool has a poor depiction due to Gibbs ringing. The ringing manifests as a dark rim artifact along the myocardium, which can mimic the behavior of ischemic patterns. Accurate identification of the false positive DRAs are challenging at low resolutions.

classical approach of performing MR-MPI is to freeze the cardiac and respiratory motions by collecting k-space data from multiple 2-dimensional (2D) slices during each beat within the quiescent diastolic phase ($\approx 250\text{-}300$ ms), and with breath holding constraints [7] (see figure 1.2). Since the number of samples that can be acquired in this finite window is limited, this approach often restricts the image quality required for accurate depiction of myocardial perfusion changes. Specifically, the need to reliably detect subtle lesions imposes the following demands in MR-MPI:

- Good in-plane spatial resolution ($< 2\text{mm}\times 2\text{mm}$) to detect subendocardial ischemia, assess transmural extent of defects, and reduce dark rim artifacts - artifacts which could mimic subendocardial ischemia primarily caused due to Gibbs ringing at low spatial resolutions [3]. (also see figure 1.3).
- Contiguous spatial coverage (> 8 slices) to image all regions of the heart, which

makes the scheme less likely to miss ischemic areas, and allows for better sizing of ischemia.

- High temporal resolution (1-2 frames/sec) and long breath-hold duration (atleast 30 seconds) to accurately fit the kinetic perfusion model.
- Good contrast to noise ratio.

These contradictory goals are often difficult to realize with the current clinically available multi slice 2D MR-MPI schemes. Often, clinicians are forced to compromise on the spatio-temporal resolution and the coverage. Additionally, the occurrence of motion artifacts with the current perfusion protocols are common, and are majorly attributed due to improper breath hold patterns. This particularly limits the utility of perfusion imaging to its full potential in large patient domain population such as patients with respiratory insufficiencies, pediatric subjects, arrhythmias, and atrial fibrillation. Due to the above limitations, there is a need for a myocardial perfusion MRI set up that can relax the long breathhold constraints, can produce dynamic images with high spatio-temporal resolutions, and extended volume coverage.

1.2 Challenges with existing acceleration schemes

To address the challenges associated with slow MRI imaging speed, recent interest has been on accelerated schemes that recover the spatio-temporal signal from sub-sampled/accelerated k-t measurements. Current used accelerations schemes include parallel imaging [8,9], and spatio-temporal model based schemes. The spatio-temporal model based image reconstruction schemes were originally introduced to improve the spatio-temporal resolution and to minimize the acquisition time in breath-held cardiac MRI (eg: UNFOLD, k-t BLAST, k-t FOCUSS, k-t SPARSE) [10–15]. These methods exploit the banded structure or sparsity of the data in $x - f$ space to recover the dynamic images from under-sampled measurements. These schemes

model the dynamic data as a linear combination of a finite number of Fourier exponential basis functions. These bases are estimated from low spatial resolution training data. To make the recovery well posed, these methods rely on the design of specialized $k - t$ sampling patterns such that there are few signal overlaps in the temporal Fourier space. Compressed sensing (CS) schemes have demonstrated potential to improve breath held MR-MPI [16]. The key difference between CS schemes with the $k - t$ type model based methods is that they do not require any training data. They instead rely on the sparse representation of the data in the transform domain and utilize incoherent sampling conditions coupled with a non-linear reconstruction to recover the data from undersampled $k-t$ data. In the context of myocardial perfusion MRI, CS recovery based on exploiting sparse representations in transform domains such as temporal Fourier domain [16], temporal total variation [17].

However all the above methods have been observed to perform poorly in the presence of respiratory motion [17], which is often difficult to avoid in several cardiac imaging applications. For example, the dynamic contrast variations in the myocardium are typically imaged for 40-60 seconds in cardiac perfusion imaging; most patients cannot maintain a breath-hold for such long durations, especially during hyperemia. The respiratory motion and contrast variations due to bolus passage severely degrade the structure and sparsity in $x-f$ space, which makes the above model-based schemes ineffective (see figure 1.4), (also (see figure 1.5) for an example free breathing CS reconstruction using $x-f$ sparsity that resulted in motion artifacts due to reduced sparsity in the transform domain). The current clinical practice of using small image matrices and restricting the temporal resolution and spatial coverage (typically three slices are acquired) present several challenges in the interpretation of cardiac perfusion MRI data. It is also not straight-forward to extend the current model-based schemes to general dynamic imaging applications.




Figure 1.4: Dynamic data representation in the temporal Fourier transform (x-f) domain: A numerical cardiac phantom is considered to mimic breath held cine data in the top row, breath held perfusion data in the middle row, and free breathing perfusion data in the bottom row. The first column shows the spatial image (x-y) for a specific time frame; the second column depicts the image time profile (x-t) through the arrows in the first column. The x-f representations are shown in the third column. Note that the x-f space is highly structured and sparse for breathheld cine applications, while the sparsity is disturbed significantly in perfusion and breathing applications. Model based schemes such as UNFOLD, k-t BLAST/SENSE, k-t FOCUSS utilize the compactness/sparse property in the x-f space which is clearly disturbed significantly in free breathing MR-MPI.

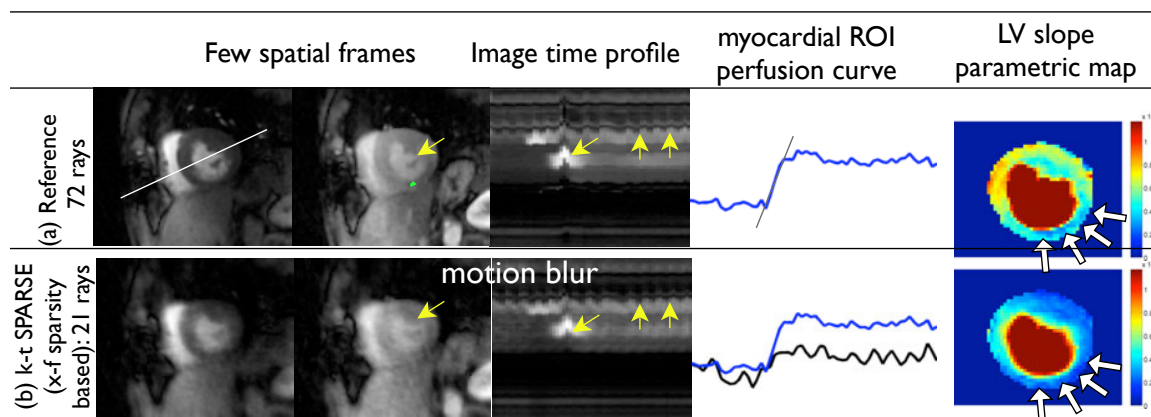


Figure 1.5: Example reconstruction of the k-t SPARSE method (compressed sensing based on x-f sparsity) on free breathing MR-MPI data. The first row corresponds to the close to fully sampled reference data with 72 radial rays per frame. Retrospective sampling of choosing 21 radial rays/frame from the acquired data was considered with k-t SPARSE. The x-y, x-t, region of interest (marked by green in the top left image) myocardial perfusion curve and the left ventricle parametric map of the slopes of the perfusion curves are shown in the columns. In the perfusion maps in (a), regions of reduced perfusion uptake are depicted in the inferior myocardium wall (as pointed by the white arrows). It is observed that k-t SPARSE was sensitive to motion artifacts and resulted in spatio-temporal blurring. This is mainly due to the reduced sparsity of the free breathing data in the x-f space; in this example, the motion blurring resulted in loss of depiction of the ischemic regions.

1.3 Overall objective

The overall objective of this thesis is to develop a novel dynamic imaging framework that can enable free breathing myocardial perfusion imaging with high spatio-temporal resolutions of the whole heart. To achieve this, this thesis deals with developing novel image reconstruction methods for the reconstruction of dynamic MRI data from highly accelerated / under-sampled Fourier measurements. It specifically focuses on novel blind or adaptive image models that represent the dynamic image data set using adaptive temporal bases (bases derived from the data at hand). This is in sharp contrast to current models that rely on predetermined temporal bases (such as Fourier bases), which require assumptions such as temporal periodicity and breath holding that are clearly not met in free breathing MR-MPI.

This dissertation deals with developing novel image reconstruction methods in conjunction with non-Cartesian sampling for the reconstruction of dynamic MRI data from highly accelerated / under-sampled Fourier measurements. The developed methods are applicable to a wide range of dynamic imaging problems. In the context of MR-MPI, this dissertation show feasibilities that the developed methods can enable free breathing myocardial perfusion MRI acquisitions with high spatio-temporal resolutions ($< 2\text{mm} \times 2\text{mm}^2$, 1 heart beat) and slice coverage (upto 7 slices).

1.4 Main contributions of the thesis

1. We have introduced a novel algorithm that exploits low rank and sparse structure of dynamic data to enable reconstruction from under-sampled $k - t$ space data (k-t SLR). The main novelties and benefits of k-t SLR over existing methods are:
 - Utilization of the Karhunen-Louve Transform (KLT): In contrast to classical model based cine MRI schemes that rely on the sparsity or banded structure in Fourier space, our scheme utilizes the compact representation

of the data in the KLT domain to exploit the correlations in the dataset. This signal adaptive strategy gives good performance even when the data is not sparse in x-f space. This makes our scheme attractive to free breathing perfusion MRI.

- Realization of a single step KLT scheme: In comparison to current KLT-based methods that rely on a two-step approach to first estimate the basis functions and then use it for reconstruction, we pose the problem as a spectrally regularized matrix recovery problem. This allows for a simultaneous estimation of the temporal bases and the KLT model coefficients from the measured data. This approach addresses the trade-offs associated with current KLT schemes, and offers significantly enhanced performance. Moreover, it can account for arbitrary non-Cartesian sampling patterns, which are more efficient in acquiring multi-dimensional data.
- The use of non convex spectral penalty: In contrast to the classical matrix recovery schemes that rely on the nuclear norm penalty, we use the non convex ($p < 1$) Schatten p-norm matrix penalty to improve the recovery rate. We use homotopy like continuation to minimize the local minima problems. Our results show that choosing $p < 1$ considerably improves the fidelity of the data for a specified number of measurements.
- Use of additional sparsity properties along with the KLT scheme: In dynamic MRI, the individual frames typically have sparse gradients. We propose to exploit this, in addition to the compact KLT representation, by penalizing the sum of total variation norms of the images. Our results demonstrate the reconstructions of the KLT model can be considerably improved when additional sparsity properties such as sparse gradients are exploited.

We have performed validations of k - t SLR using numerical phantoms and in-vivo cardiac perfusion MRI data to demonstrate the improvement in performance offered over existing model based methods. Related publications include Refs. [18–25].

2. Extensions of $k - t$ SLR to account for radial sampling and parallel imaging and acceleration of free breathing MR-MPI: We have extended $k - t$ SLR in terms of benefiting from non-Cartesian radial sampling, and parallel imaging. We have introduced a novel augmented Lagrangian framework to considerably improve the algorithm's convergence rate. Using the improved k - t SLR framework, we have successfully demonstrated the feasibility of accelerating of free breathing stress and rest myocardial perfusion MRI data. Our results demonstrate k - t SLR to provide faithful reconstructions with minimal artifacts compared to existing MR-MPI acceleration methods. Related publications include Refs. [26, 27].

3. A novel blind compressed sensing (BCS) dynamic MRI framework:

We have proposed a novel BCS framework to recover dynamic MRI from undersampled measurements. This scheme models the dynamic signal as a sparse linear combination of temporal basis functions, chosen from a large dictionary. In contrast to classical compressed sensing, the BCS scheme simultaneously estimates the dictionary and the sparse coefficients from the undersampled measurements. Apart from the sparsity of the coefficients, the key difference of the BCS scheme with current low rank methods is the non-orthogonal nature of the dictionary basis functions. Since the number of degrees of freedom of the BCS model is smaller than that of the low-rank methods, it provides improved reconstructions at high acceleration rates. We formulate the reconstruction as a constrained optimization problem; the objective function is the linear combination

of a data consistency term and sparsity promoting ℓ_1 prior of the coefficients. The Frobenius norm dictionary constraint is used to avoid scale ambiguity. We have introduced an efficient majorize minimize optimization algorithm to solve the resulting BCS problem. This algorithm is seen to be considerably faster than approaches that alternates between sparse coding and dictionary estimation, as well as the extension of K-SVD dictionary learning scheme. We have also shown that the proposed scheme is more robust to local minima compared to K-SVD method, which relies on greedy sparse coding. Our phase transition experiments demonstrate that the BCS scheme provides much better recovery rates than classical Fourier-based CS schemes, while being only marginally worse than the dictionary aware setting. Since the overhead in additionally estimating the dictionary is low, this method can be very useful in dynamic MRI applications, where the signal is not sparse in known dictionaries. We have demonstrated the utility of the BCS scheme in accelerating contrast enhanced dynamic data. Related publications include Refs. [28–31].

4. To address the problem of motion artifacts in typical compressed sensing reconstructions, we have proposed a novel motion estimation and compensated compressed sensing reconstruction scheme. The proposed scheme jointly estimates the motion and the dynamic images in first pass cardiac perfusion MR imaging. We formulate the recovery as a minimization scheme using a unified objective function that combines data consistency, and compressed sensing constraints on the motion compensated dataset. We have introduced an efficient variable splitting framework with continuation to decouple the problem into simpler sub-problems. The novelties enabled by this optimization are (a) a generalized formulation capable of handling any temporal sparsifying transform (such as temporal Fourier, temporal gradient, temporal PCA), (b) derivation of a reference dataset that is free of motion from the measurements themselves

(c) efficient decoupling of the motion estimation problem from the reconstruction problem. Unlike existing MC-CS schemes, the proposed scheme does not require fully sampled prescans or navigators for motion estimation. Validations on numerical phantoms and invivo free breathing MR-MPI data demonstrate the utility of the proposed scheme in significantly improving compressed sensing reconstructions. Related publications include [32–35].

1.5 Outline of the thesis

The structure of the dissertation is as follows:

- **Chapter 2: Accelerated dynamic MRI by exploiting low rank and sparse structure: k-t SLR:** This chapter presents a novel method to reconstruct dynamic data from undersampled Fourier data. An framework that exploits the low rank and sparse structure of dynamic data is presented. The formulation and the optimization algorithm is highlighted. Results on numerical phantom and invivo cardiac perfusion MRI data are presented to validate the method.
- **Chapter 3: Accelerating free breathing myocardial perfusion MRI using multi coil radial k-t SLR:** This chapter presents the development of k-t SLR to include non-Cartesian radial sampling and parallel imaging. The chapter discusses considerations on the sampling pattern, and an implementation of a fast augmented Lagrangian optimization algorithm. Using the improved k-t SLR frame work, this chapter demonstrates the feasibility of accelerating of free breathing stress and rest myocardial perfusion MRI data. Extensive comparisons against existing accelerated methods in the context of MR-MPI are presented.
- **Chapter 4: Blind compressed sensing dynamic MRI:** This chapter presents a novel blind compressed sensing approach to reconstruct dynamic MRI data

from undersampled measurements. It discusses the details on different formulations, and the corresponding algorithmic implementations. Empirical simulations are presented to determine the phase transition behavior of the blind compressed sensing (BCS) scheme. Validations are performed on contrast enhanced data by comparing the BCS scheme with existing compressed sensing and low rank models.

- **Chapter 5: A novel motion compensated compressed sensing framework for myocardial perfusion MRI:** This chapter presents a new joint reconstruction and motion estimation framework in the context of compressed sensing. The chapter discusses in detail the new formulation, and its novelties to account for different sparsifying transforms, independence of training data for motion estimation. Validations are performed on numerical phantom and in vivo data to study the resulting formulation in terms of convergence and its performance against existing compressed sensing methods.
- **Chapter 6: Summary and Future work:** Contributions of this thesis are summarized and recommendations on future work are suggested.

CHAPTER 2 ACCELERATED DYNAMIC MRI BY EXPLOITING LOW RANK AND SPARSE STRUCTURE (K-T SLR)

2.1 Introduction

The imaging of dynamically evolving phenomena is central to several magnetic resonance imaging (MRI) applications, including cardiac, perfusion, functional, and gastro-intestinal imaging. Achieving high spatio-temporal resolutions is challenging in dynamic MRI due to the hardware limitations and the risk of peripheral nerve stimulation. As discussed in chapter I, model based schemes that exploit the banded structure/sparsity of dynamic data in the x-f space have demonstrated success in improving breath held cardiac MRI [10–15]. However, these methods perform poorly in the presence of contrast dynamics and motion. To address this, recent interest has been on methods that exploit the compact signal representation in the Karhunen Louve transform (KLT) domain as an alternative to x-f space sparsity/structure [36–39]. Since KLT is a data-derived transform, the resulting adaptive scheme is capable of exploiting the correlations in the data, even when the temporal profiles of the voxels are not periodic. This property makes these methods applicable to a range of dynamic imaging problems. Current KLT-based algorithms rely on a two-step approach to recover the data [37–40]. Specifically, they estimate the temporal basis functions using the singular value decomposition (SVD) of a training dataset; the training dataset is an image time series with low spatial resolution and Nyquist temporal sampling rate. The training dataset is obtained as the IFFT of the central phase encodes, which is collected along with higher k-space samples at sub-Nyquist temporal sampling rates. The estimated temporal basis functions are then used to reconstruct the data with high spatio-temporal resolution from sub-Nyquist sampled k-space data. These schemes rely on the implicit assumption that the temporal basis functions estimated from the training data closely approximate the principal components of the entire data. Clearly, this approximation is heavily dependent on the number




Figure 2.1: Utility of KLT in compactly representing the dynamic image time series: The numerical simulation of breath held cine data (top row) and ungated free breathing data (bottom row), along with their corresponding representations in the $x-f$ and $x-KLT$ spaces are shown. The $x-f$ space coefficients are highly sparse/structured in the context of breath-held acquisitions due to the pseudo-periodic nature of heartbeats. The structure and sparsity of the $x-f$ space is disturbed in the presence of breathing motion. In contrast, the free breathing data is compact in the $x-KLT$ space. The few significant singular values implies that the dataset can be efficiently approximated as a low rank matrix, described by (2.1).

of phase encodes in the training data. For example, if only a single phase encode is used in the training stage, the estimated temporal functions will fail to capture the dynamics due to intermediate vertical shifts resulting from respiratory motion. Moreover, this may also result in the scheme failing to capture small details such as perfusion defects. These problems can be minimized by acquiring more phase-encodes in the training data. However, this comes at the expense of the number of higher k-space encodes that can be acquired at a specified acceleration factor, resulting in significant spatial aliasing artifacts.

In this chapter, we propose a novel algorithm to significantly accelerate dynamic MRI by exploiting the correlations between the temporal profiles of the voxels. In

contrast to the classical KLT-based schemes that use the above two-step approach [37–40], we propose to simultaneously estimate the temporal basis functions and its spatial weights directly from the entire $k - t$ space data. This approach is enabled by the re-interpretation of the KLT based reconstruction as a spectrally regularized matrix recovery scheme. Specifically, we pose the joint estimation of the bases and the signal as the recovery of a low-rank matrix, obtained by stacking the temporal dynamics of the voxels, from the measured data. This approach provides more accurate estimates of the temporal basis functions and hence result in reconstructions with better quality at a specified acceleration.

The recovery of a low-rank matrix using nuclear norm minimization has been rigorously studied by several researchers [41–45]. Motivated by the recent results in the use of non-convex penalties in compressed sensing [46,47], we introduce novel non-convex spectral penalties to minimize the number of measurements required to recover a low-rank matrix. By suppressing the singular vectors that correspond to aliasing artifacts, this approach can considerably improve the reconstructions. Moreover, the images in dynamic time series themselves can be assumed to have sparse wavelet coefficients or gradients. We propose to additionally exploit the sparsity of the matrix in pre-determined domains to further improve the recovery rate. Since the degrees of freedom in representing sparse and low-rank matrices are significantly lower than the class of arbitrary low-rank matrices, this approach enables us to improve the recovery rate. We do not promote joint sparsity as done in [48]. In our work, the temporal basis functions themselves are not constrained to be sparse in any bases; enforcing the sparsity in a specified space (eg. Fourier) may introduce significant bias in the presence of motion (and/or perfusion) (see Fig. 1). Moreover, we observe that different temporal basis functions play dominant roles in different spatial regions. Since the sparsity properties of these functions may be very different, we expect the use of joint sparsity penalty to smooth subtle motion/perfusion induced variations.

The preliminary version of this work was reported in our conference paper [20]. The work of Haldar et. al. [49], which was also published in the same proceedings, is conceptually similar to the proposed scheme. However, they do not use sparsity priors and their optimization scheme is drastically different from the proposed scheme.

Most of the existing convex matrix recovery algorithms are based on iterative singular value thresholding [42, 50, 51]. Since it is not straightforward to extend these schemes to our problem with both sparsity and low-rank penalties, we introduce a novel variable splitting algorithm for the fast minimization of the optimization criterion. This approach is the generalization of similar algorithms used for total variation minimization [52, 53] to matrix recovery. We demonstrate the utility of the proposed scheme in the context of clinical cardiac perfusion MRI. Validations using numerical phantoms and in-vivo data demonstrate the significant improvement in performance over state of the art methods. Although we focus on cardiac perfusion imaging in this chapter, the algorithm is readily applicable to most dynamic MRI applications.

2.2 Dynamic MRI using KLT

We denote the spatio-temporal signal as $\gamma(\mathbf{x}, t)$, where \mathbf{x} is the spatial location and t denotes time. The dynamic MRI measurements correspond to the samples of the signal in $k - t$ space, corrupted by noise:

$$\mathbf{b}_i = \int_{\mathbf{x}} \gamma(\mathbf{x}, t_i) \exp(-j\mathbf{k}_i^T \mathbf{x}) d\mathbf{x} + \mathbf{n}_i; \quad i = 0, \dots, s - 1.$$

Here, (\mathbf{k}_i, t_i) indicates the i^{th} sampling location. We denote the set of sampling locations as $\Xi = \{(\mathbf{k}_i, t_i), i = 0, \dots, s - 1\}$. The above expression can be rewritten in the vector form as $\mathbf{b} = \mathcal{A}(\gamma) + \mathbf{n}$, where, \mathcal{A} is the Fourier sampling operator. The goal is to recover the signal $\gamma(\mathbf{x}, t)$ from the measured k-t space samples.

In dynamic imaging applications, the temporal profiles of the voxels, indicated by

the n -dimensional vectors

$$\mathbf{q}_i = [\gamma(\mathbf{x}_i, t_0), \gamma(\mathbf{x}_i, t_1), \dots, \gamma(\mathbf{x}_i, t_{n-1})]^T; i = 0, \dots, m - 1,$$

are highly correlated/ linearly dependent. Here, m is the number of voxels. Liang et. al., proposed to re-arrange the spatio-temporal signal $\gamma(\mathbf{x}, t)$ in a matrix form to exploit the correlations [36, 37]:

$$\mathbf{\Gamma} = \begin{bmatrix} \gamma(\mathbf{x}_0, t_0) & \dots & \gamma(\mathbf{x}_0, t_{n-1}) \\ \vdots & & \\ \gamma(\mathbf{x}_{m-1}, t_0) & \dots & \gamma(\mathbf{x}_{m-1}, t_{n-1}) \end{bmatrix} \quad (2.1)$$

The rows of $\mathbf{\Gamma}$ correspond to the voxels, while the columns represent the temporal samples. Since the rows of this $m \times n$ matrix are linearly dependent, the rank of $\mathbf{\Gamma}$, is given by $r < \min(m, n)$. An arbitrary $m \times n$ matrix of rank r can be decomposed as

$$\mathbf{\Gamma} = \underbrace{\mathbf{U}}_{m \times r} \underbrace{\mathbf{\Sigma}}_{r \times r} \underbrace{\mathbf{V}^H}_{r \times n}. \quad (2.2)$$

This decomposition implies that the spatio-temporal signal $\gamma(\mathbf{x}, t)$ can be expressed as a weighted linear combination of r temporal basis functions [36, 37]:

$$\gamma(\mathbf{x}, t) = \sum_{i=0}^{r-1} \rho_i(\mathbf{x}) v_i(t). \quad (2.3)$$

The temporal basis functions $v_i(t)$ are the columns of the matrix \mathbf{V} in (2.2) while the spatial weights $\rho_i(\mathbf{x})$ are the row vectors of $\mathbf{U}\mathbf{\Sigma}$ (often termed as spatial weights). The utility of this scheme in compactly representing the dynamic time series data is illustrated in Fig. 2.1. Most of the KLT-based algorithms use the below-mentioned two-step strategy to reconstruct the spatio-temporal signal [36–40].

1. Estimate the temporal basis functions $v_i(t); i = 0, \dots, r - 1$ using SVD of the

training image time-series. The training data consists of dynamic image data, acquired with low-spatial resolution and high temporal sampling rate; it is obtained as the IFFT of the central phase encodes, acquired at the Nyquist temporal sampling rate.

2. Use the linear model specified by (2.3) to recover the cardiac data from sub-Nyquist sampled measurements, using the pre-determined temporal basis functions $v_i(t)$. This involves the estimation of the spatial weight images $\rho_i(\mathbf{x}); i = 0, \dots, r - 1$ from the under-sampled measurements. Since $r \ll n$, this approach provides a significant reduction in the number of unknowns and hence the number of measurements.

These schemes implicitly assume that the principal basis functions estimated from the low-resolution data to closely approximate the original KLT basis functions. As discussed previously, this assumption is violated when the number of phase encodes in the training data are too few, resulting in the loss of subtle details and reconstructions with inaccurate temporal dynamics. While the acquisition of more training data can minimize these problems, this comes at the expense of the number of high-frequency encodes that can be acquired at a specified acceleration rate; this can often result in aliasing artifacts. In summary, the performance of the two-step schemes requires a fine balance between the amount of training data and the number of high-frequency encodes. To overcome these problems, we introduce the single-step spectrally regularized reconstruction scheme in Section 2.3.

2.2.1 Matrix recovery using nuclear norm minimization

The recovery of a low-rank matrix $\mathbf{\Gamma}$ from few of its linear measurements is currently a hot topic in signal processing. The recent theoretical results indicate that a matrix $\mathbf{\Gamma} \in \mathbb{R}^{m \times n}$ of rank $r; r \leq \min(m, n)$ can be perfectly recovered from its

measurements $\mathbf{b} = \mathcal{A}(\mathbf{\Gamma})$ by solving the constrained optimization problem [43, 54]:

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 \text{ such that } \text{rank}(\mathbf{\Gamma}) \leq r. \quad (2.4)$$

The rank constraint is an effective means of regularizing the inverse problem since it significantly reduces the number of degrees of freedom. Specifically, the number of degrees of freedom in representing $m \times n$ matrices of rank r is $r(m + n - r)$, which is much smaller than mn . Recht et. al. have shown that this approach perfectly recovers the matrix with a high probability, if the random measurement ensemble is used and the number of measurements exceeds a constant (two to four) times the number of degrees of freedom [54]. Reformulating the above constrained optimization problem using Lagrange's multipliers, we get

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \lambda \text{rank}(\mathbf{\Gamma}). \quad (2.5)$$

Since the rank penalty is non-convex, it is often replaced with the nuclear norm, which is the closest convex relaxation. The nuclear norm of an r -rank matrix $\mathbf{\Gamma} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, denoted by $\|\mathbf{\Gamma}\|_*$, is the sum of the singular values of $\mathbf{\Gamma}$ ($\|\mathbf{\Gamma}\|_* = \sum_i(\Sigma_{i,i})$). With this relaxation, the recovery of the matrix is simplified as

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}} \underbrace{\|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \lambda \|\mathbf{\Gamma}\|_*}_{c(\mathbf{\Gamma})} \quad (2.6)$$

2.3 k - t SLR: Formulation

We introduce the proposed algorithm in two steps to facilitate its easy understanding. We will first introduce the reconstruction of the spatio-temporal signal as a spectrally regularized matrix recovery problem in Section 2.3.1. This scheme is then further constrained using additional sparsity priors to improve the recovery rate in

Section 2.3.2.

2.3.1 Regularized matrix recovery using spectral priors

We recover the matrix $\mathbf{\Gamma}$ from the undersampled $k - t$ space data as a spectrally regularized optimization problem, similar to (2.6):

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \lambda \varphi(\mathbf{\Gamma}), \quad (2.7)$$

where $\varphi(\mathbf{\Gamma})$ is an appropriate spectral penalty¹. We use the general class of Schatten p-functionals, specified by

$$\varphi(\mathbf{\Gamma}) = (\|\mathbf{\Gamma}\|_p)^p = \sum_{i=1}^{\min\{m,n\}} \sigma_i^p. \quad (2.8)$$

Here, $\mathbf{\Gamma} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$ is the singular value decomposition of $\mathbf{\Gamma}$ and $\mathbf{\Sigma} = \text{diag}([\sigma_0, \sigma_1, \dots, \sigma_{r-1}])$. The above spectral penalty simplifies to the nuclear norm for $p = 1$. When $p \leq 1$, this penalty ceases to be a norm and is non-convex. The use of similar non-convex semi-norms are well-studied in the context of vector recovery; they are found to significantly improve the reconstruction of the signal from fewer measurements, in comparison to the standard ℓ_1 semi-norms [55–59]. Majumdar et. al [60] introduced the non-Convex Schatten p-norm for denoising and 2D MRI. However, the optimization algorithm in [60] is very different from our approach. In addition to providing rapid convergence, our algorithm is also capable of using sparsity penalties.

Note that the cost function, specified by (2.8), does not depend explicitly on the temporal basis functions or its spatial weights as in the case of current two-step KLT schemes. However, the optimization algorithm to minimize (2.8) iteratively updates the temporal basis functions and spatial weights, which are essentially the column vectors of \mathbf{V} and \mathbf{U} respectively. The optimization algorithm is discussed in detail in

¹Such cost functions are termed as spectral penalties since they are functions of the singular values of the matrix.

Section 2.4.

2.3.2 Regularized matrix recovery using spectral and sparsity priors

In dynamic imaging applications, the images in the time series may have sparse wavelet coefficients or sparse gradients. In addition, if the intensity profiles of the voxels are periodic (eg. cardiac cine), the columns of $\mathbf{\Gamma}$ may be sparse in the Fourier domain. We propose to additionally exploit the sparsity of the signal in specified basis sets along with the low-rank property to further improve the recovery rate. Specifically, we consider the simple example of recovering an r -rank matrix $\mathbf{\Gamma} \in \mathbb{R}^{m \times n}$ that has atmost N non-zero entries in a specified basis: $\|\mathbf{\Phi}^H \mathbf{\Gamma} \mathbf{\Psi}\|_{\ell_0} \leq N$. Here, $\mathbf{\Phi}$ and $\mathbf{\Psi}$ are transformations or operators that sparsify the row-space and column space of $\mathbf{\Gamma}$, respectively. For example, $\mathbf{\Phi}$ can be chosen to be the 2-D wavelet transform to sparsify each of the images in the time series, while $\mathbf{\Psi}$ can be a 1-D Fourier transform to exploit the pseudo-periodic nature of motion. The set of matrices that satisfy both the rank and the sparsity constraints are far smaller in dimension than the class of matrices that satisfy only one of the constraints. For example, consider an r -rank matrix $\mathbf{\Gamma}$ whose right and left singular vectors are k_1 and k_2 sparse. The number of degrees of freedom of such r -rank matrices is given by $r(k_1 + k_2 - r)$. If $k_1 \ll m$ and $k_2 \ll n$, the use of this prior knowledge, along with the low-rank constraint, can significantly reduce the number of measurements required to recover the matrix. To exploit the sparsity and low-rank properties of the matrix, we formulate the problem as

$$\begin{aligned} \mathbf{\Gamma}^* &= \arg \min_{\mathbf{\Gamma}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 \\ &\text{s.t. } \{\text{rank}(\mathbf{\Gamma}) \leq r, \|\mathbf{\Phi}^H \mathbf{\Gamma} \mathbf{\Psi}\|_{\ell_0} < K\} \end{aligned} \quad (2.9)$$

Rewriting the above constrained optimization problem using Lagrange's multipliers and relaxing the penalties, we obtain

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \lambda_1 \varphi(\mathbf{\Gamma}) + \lambda_2 \psi(\mathbf{\Gamma}), \quad (2.10)$$

where $\psi(\mathbf{\Gamma}) = \|\Phi^H \mathbf{\Gamma} \Psi\|_{\ell_1}$ is a surrogate for the ℓ_0 term and $\varphi(\mathbf{\Gamma}) = \|\mathbf{\Gamma}\|_p^p$. When $p \geq 1$, the cost function is convex and hence will have a unique minimum.

While it is straight-forward to use this scheme to exploit the sparsity in different transform domains, it cannot be used for non-separable total variation (TV) penalties. Exploiting the sparsity of the gradient has proven to be very powerful in various image recovery application and is shown to provide comparable or better performance than most other transform domain schemes [61]. To adapt this scheme for TV regularization, we consider a collection of transforms/operators on $\mathbf{\Gamma}$, indicated by $\Phi_i^H \mathbf{\Gamma} \Psi_i, i = 0, \dots, q - 1$, and specify the non-separable penalty as

$$\psi(\mathbf{\Gamma}) = \left\| \left\| \sum_{i=0}^{q-1} |\Phi_i^H \mathbf{\Gamma} \Psi_i|^2 \right\|_{\ell_1} \right\|. \quad (2.11)$$

The total variation norm of the entire volume can be obtained by setting $q = 3$, $\Phi_0 = \mathbf{D}_x$; $\Psi_0 = \mathbf{I}$, $\Phi_1 = \mathbf{D}_y$; $\Psi_1 = \mathbf{I}$, and $\Phi_2 = \mathbf{I}$; $\Psi_2 = \mathbf{D}_t$, where \mathbf{D}_x , \mathbf{D}_y and \mathbf{D}_t are the finite difference matrices along x , y , and t respectively. Note that the above expression simplifies to the standard ℓ_1 penalty, when the number of transforms/operators is $q = 1$.

The proposed scheme is well posed since the sparsifying transforms/operators $\Phi_i, \Psi_i; i = 0, \dots, q - 1$ are incoherent with the Fourier sampling operator. We do not need the additional assumption of the right and the left singular vectors of $\mathbf{\Gamma}$ to be in-coherent with the operator that picks the samples/matrix entries of $\mathbf{\Gamma}$ as in [41] to make the problem well posed.

2.4 Optimization Algorithm

It is not straightforward to extend the current nuclear norm minimization schemes [42,50,51] to solve (2.10), since it uses both sparsity and spectral penalties. We introduce a novel variable splitting algorithm for the efficient recovery of the matrix using (2.10). We pose the regularized matrix recovery scheme as a constrained minimization problem using variable splitting:

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \lambda_1 \varphi(\mathbf{R}) + \lambda_2 \left\| \sqrt{\sum_{i=0}^{q-1} \|\mathbf{S}_i\|^2} \right\|_{\ell_1} \quad \text{s.t. } \mathbf{\Gamma} = \mathbf{R}; \quad \mathbf{S}_i = \mathbf{\Phi}_i^H \mathbf{\Gamma} \mathbf{\Psi}_i; \quad i = 0, \dots, q$$

Here, \mathbf{R} and $\mathbf{S}_i; i = 0, \dots, q - 1$ are auxiliary variables, which are also determined during the optimization process. The rationale behind the above decomposition is that the constrained optimization problem is simpler to solve than its unconstrained version, specified by (2.10). We solve (2.12) using the penalty method, where we minimize

$$\begin{aligned} \mathcal{D}_{\beta_1, \beta_2}(\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i) &= \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \lambda_1 \varphi(\mathbf{R}) + \lambda_2 \left\| \sqrt{\sum_{i=0}^{q-1} \|\mathbf{S}_i\|^2} \right\|_{\ell_1} \\ &\quad + \frac{\beta_1}{2} \|\mathbf{\Gamma} - \mathbf{R}\|^2 + \frac{\beta_2}{2} \sum_{i=0}^{q-1} \|\mathbf{\Phi}_i^H \mathbf{\Gamma} \mathbf{\Psi}_i - \mathbf{S}_i\|^2 \end{aligned} \quad (2.13)$$

with respect to $\mathbf{\Gamma}$, \mathbf{R} and $\mathbf{S}_i; i = 0, \dots, q - 1$. The second row of (2.13) are the penalties introduced to enforce the constraints $\mathbf{\Gamma} = \mathbf{R}$ and $\mathbf{S}_i = \mathbf{\Phi}_i^H \mathbf{\Gamma} \mathbf{\Psi}_i; i = 0, \dots, q - 1$. The solution of the above problem tends to that of (2.12), when $\beta_1, \beta_2 \rightarrow \infty$. We solve (2.13) using a three-step alternating minimization scheme below (2.14,2.15,2.16), where we solve a variable of interest assuming the rest to be known:

$$\mathbf{\Gamma}_{n+1} = \arg \min_{\mathbf{\Gamma}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|^2 + \frac{\beta_1}{2} \|\mathbf{\Gamma} - \mathbf{R}_n\|^2 +$$

$$\frac{\beta_2}{2} \sum_{i=0}^{q-1} \|\Phi_i^H \Gamma \Psi_i - \mathbf{S}_{i,n}\|^2, \quad (2.14)$$

$$\mathbf{R}_{n+1} = \arg \min_{\mathbf{R}} \|\Gamma_{n+1} - \mathbf{R}\|^2 + 2\lambda_1/\beta_1 \varphi(\mathbf{R}), \quad (2.15)$$

$$\mathbf{S}_{i,n+1} = \arg \min_{\{\mathbf{S}_i\}} \sum_{i=0}^{q-1} \|\Phi_i^H \Gamma_{n+1} \Psi_i - \mathbf{S}_i\|^2 + 2\lambda_2/\beta_2 \left\| \sqrt{\sum_{i=0}^{q-1} \|\mathbf{S}_i\|^2} \right\|_{\ell_1}; i = 0, \dots, q-1 \quad (2.16)$$

Similar alternating directions methods are widely used in compressed sensing and TV minimization [53, 62]. The first sub-problem (2.14) is quadratic and hence can be solved analytically as

$$\Gamma_{n+1} = \left(\mathcal{A}^T \mathcal{A} + \frac{\beta_1}{2} \mathcal{I} + \frac{\beta_2}{2} \sum_{i=0}^{q-1} \mathcal{Q}_i^T \mathcal{Q}_i \right)^{-1} \left(\mathcal{A}^T \mathbf{b} + \frac{\beta_1}{2} \mathbf{R} + \frac{\beta_2}{2} \sum_{i=0}^{q-1} \mathbf{S}_i \right) = \mathcal{T}(\mathbf{R}, \mathbf{S}_i), \quad (2.17)$$

where the operator \mathcal{Q}_i is defined as $\mathcal{Q}_i(\Gamma) = \Phi_i^H \Gamma \Psi_i; i = 0, \dots, q-1$. This step can be efficiently evaluated in the Fourier domain, if the measurements are Fourier samples on a Cartesian grid [52, 53]. We instead rely on solving (2.14) using a few conjugate gradient steps, since we are dealing with non-Cartesian sampling problems.

The second sub problem is of the similar form of standard nuclear norm minimization problems. The iterative singular value thresholding (IST) scheme used in nuclear norm minimization can be generalized to the case that has non convex spectral penalties. The generalization in this regard, would lead to obtaining \mathbf{R}_{n+1} as a singular value thresholding of Γ_{n+1} , specified by $\mathcal{S}_{\lambda_1/\beta_1}$:

$$\mathbf{R}_{n+1} = (\mathcal{S}_{\lambda_1/\beta_1} \circ \mathcal{T})(\mathbf{R}_n, \mathbf{S}_{i,n}), \quad (2.18)$$

where the singular value shrinkage is specified by,

$$\mathcal{S}_{\lambda_1/\beta_1}(\Gamma_{n+1}) = \sum_{i=0}^{\min(m,n)} (\sigma_i - \lambda \sigma_i^{p-1}/\beta)_+ \mathbf{u}_i \mathbf{v}_i^*, \quad (2.19)$$

Here, \mathbf{u}_i , \mathbf{v}_i and σ_i are the singular vectors and values of Γ_{n+1} , respectively. The

thresholding function is defined as

$$(\sigma)_+ = \begin{cases} \sigma & \text{if } \sigma \geq 0 \\ 0 & \text{else} \end{cases} \quad (2.20)$$

Note that, when $p = 1$, the expression in (2.19) simplifies to the shrinkage scheme used for nuclear norm minimization problems.

The solution to the third sub problem (2.16) requires the joint processing of all the terms $\mathcal{Q}_i(\mathbf{\Gamma}_{n+1}); i = 0, \dots, q-1$, such that the magnitude, specified by $\sum_{i=0}^{q-1} \|\mathcal{Q}_i(\mathbf{\Gamma}_{n+1})\|^2$, is reduced:

$$\mathbf{S}_{i,n+1} = \frac{\mathcal{Q}_i(\mathbf{\Gamma}_{n+1})}{\sum_{i=0}^{q-1} \|\mathcal{Q}_i(\mathbf{\Gamma}_{n+1})\|^2} \cdot \left(\sum_{i=0}^{q-1} \|\mathcal{Q}_i(\mathbf{\Gamma}_{n+1})\|^2 - \frac{\lambda_2}{\beta_2} \right)_+ = \kappa_{\lambda_2/\beta_2}(\mathbf{\Gamma}_{n+1}), \quad (2.21)$$

This approach is termed as multidimensional shrinkage of $\{\mathcal{Q}_i(\mathbf{\Gamma}_{n+1}), i = 0, \dots, q-1\}$ [53, 62].

The convergence of the above three step alternating minimization scheme as the penalty parameters $\beta_1, \beta_2 \rightarrow \infty$ is well known [63]. The three-step optimization scheme involves update rules based on the operators $\mathcal{T}, \mathcal{S}_{\lambda_1/\beta_1}, \kappa_{\lambda_2/\beta_2}$. Clearly, we are interested in the convergence of this iterative scheme to its fixed point, specified by $\mathbf{R}^*, \mathbf{S}_i^*; i = 0, \dots, q-1$. Following the proofs in [52], it can be shown that the three-step scheme converges to the global minimum of $\mathcal{D}_{\beta_1, \beta_2}(\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i)$ for any fixed $\beta_1 > 0, \beta_2 > 0$. The argument proceeds by showing that the operator \mathcal{T} and the shrinkage operations $\mathcal{S}_{\lambda_1/\beta_1}, \kappa_{\lambda_2/\beta_2}$ are non-expansive; (i.e., $\|\mathcal{S}_{\lambda_1/\beta_1}(\mathbf{\Gamma}_1) - \mathcal{S}_{\lambda_1/\beta_1}(\mathbf{\Gamma}_2)\|^2 \leq \|\mathbf{\Gamma}_1 - \mathbf{\Gamma}_2\|^2$)² and ($\|\kappa_{\lambda_2/\beta_2}(\mathbf{\Gamma}_1) - \kappa_{\lambda_2/\beta_2}(\mathbf{\Gamma}_2)\|^2 \leq \|\mathbf{\Gamma}_1 - \mathbf{\Gamma}_2\|^2$). Since these operators are non-expansive, the above iterative algorithm to update the auxiliary variables \mathbf{R} and \mathbf{S}_i will decrease the distances $\|\mathbf{R}_n - \mathbf{R}^*\|_F, \|\mathbf{S}_{i,n} - \mathbf{S}_i^*\|_F$ respectively at each iteration; here $(\mathbf{R}^*, \mathbf{S}_i^*; i = 0, \dots, q-1)$ is the optimal solution. This implies that $\mathbf{R}_n \rightarrow \mathbf{R}^*$ and

²The shrinkage operation is non-expansive if the spectral norm is convex.

$\mathbf{S}_{i,n} \rightarrow \mathbf{S}_i^*$ as $n \rightarrow \infty$.

High values of β_1, β_2 are needed for the solution of $\mathcal{D}_{\beta_1, \beta_2}(\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i)$ to yield a good approximation for the original minimization scheme in (2.12), as discussed before. However, the quadratic problem specified by (2.14) will become ill-conditioned for high values of β_1, β_2 , resulting in poor convergence. We propose to use a continuation strategy to overcome the tradeoff between computational complexity and accuracy. Specifically, we will start with very small values of β_1, β_2 , when the algorithm converges very fast to $\{\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i\}_{(\beta_1, \beta_2)_n}; i = 0, \dots, q - 1$, which is the solution of $\mathcal{D}_{\beta_1, \beta_2}$. To improve the quality of the approximation, we will then increase $(\beta_1, \beta_2)_n$ to obtain $(\beta_1, \beta_2)_{n+1}$ and initialize the algorithm with $\{\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i\}_{(\beta_1, \beta_2)_{n+1}} = \{\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i\}_{(\beta_1, \beta_2)_n}; i = 0, \dots, q - 1$. We observe that the continuation strategy significantly improves the convergence of the algorithm. Similar continuation strategies are widely used in similar algorithms for total variation minimization and compressed sensing [53, 62].

To summarize, the regularized matrix recovery scheme as a constrained minimization problem using variable splitting framework involves the following three step algorithm with a continuation strategy:

Variable splitting with continuation: Set $p = 0; (\beta_1)_0, (\beta_2)_0 > 0; \mathbf{R} = \mathbf{0}; \mathbf{S}_i = \mathbf{0}; i = 0, \dots, q - 1; \mathbf{\Gamma} = \mathbf{0}$

Repeat

Repeat

Update $\mathbf{\Gamma}$ by solving (2.14) using the CG scheme;

Shrinkage: $\mathbf{R} = \mathcal{S}_{\lambda_1 / (\beta_1)_p}(\mathbf{\Gamma});$

Shrinkage: $\mathbf{S}_i = \kappa_{\lambda_2 / (\beta_2)_p}(\mathbf{\Gamma}); i = 0, \dots, q - 1;$

Until stopping criterion is satisfied.

$(\beta_1)_{p+1} = (\beta_1)_p * \text{INC_FACTOR}_1; (\beta_2)_{p+1} = (\beta_2)_p * \text{INC_FACTOR}_2;$

$p = p + 1$

Until $\mathbf{R} \approx \mathbf{\Gamma}$ and $\mathbf{S}_i = \mathbf{\Phi}_i^H \mathbf{\Gamma} \mathbf{\Psi}_i; i = 0, \dots, q - 1$

Note that the above algorithm involves two loops. The parameters β_1, β_2 are incremented in the outer loop, while the minimization of $\mathcal{D}_\beta(\mathbf{\Gamma}, \mathbf{R}, \mathbf{S}_i); i = 0, \dots, q - 1$ is performed in the inner loop. We terminate the inner iteration when the stopping criterion, specified by

$$d_n = \left| \frac{\mathcal{D}_{(\beta_1, \beta_2)_p}(\mathbf{\Gamma}_n, \mathbf{R}_n) - \mathcal{D}_{(\beta_1, \beta_2)_p}(\mathbf{\Gamma}_{n-1}, \mathbf{R}_{n-1})}{\mathcal{D}_{(\beta_1, \beta_2)_p}(\mathbf{\Gamma}_n, \mathbf{R}_n)} \right| < \text{TOLERANCE}. \quad (2.22)$$

is satisfied.

The above discussed theoretical guarantees on the convergence are not valid for the non-convex spectral penalties (i.e when $p < 1$). However, we did not experience issues with convergence in our practical experiments; we obtained monotonic reduction in the cost function and the algorithm converged to a good minimum, independent of the initialization. The main reason for the good convergence performance may be attributed to the continuation scheme.

2.4.1 Implementation

The computationally expensive component of the algorithm is the singular value decomposition required for (2.15). The usual dynamic MRI data sizes are 128x128x70, resulting in the matrix $\mathbf{\Gamma}$ of size 16384 x 70. To minimize the computational complexity, we first determine the right singular vectors and the singular values as the eigen decomposition of $\mathbf{\Gamma}^H \mathbf{\Gamma}$. The eigen decomposition of this 70x70 matrix takes less than 0.1 seconds in MATLAB. The left singular vectors are then obtained using a simple least squares scheme, using the known singular values and left singular vectors. We realized the entire algorithm, described by (2.14)-(2.16), in MATLAB using Jacket [64] on a Linux workstation with eight cores and a NVIDIA Tesla graphical processing unit. We observe that the execution time for the reconstruction of the largest data (128x128x70) is approximately eight to ten minutes. We focus on the




Figure 2.2: The PINCAT phantom used to validate the proposed scheme. Three distinct spatial frames at different instances of the contrast uptake are shown in (a)-(c). The $x - t$ cross section of the dataset corresponding to the arrow in (c) is shown in (d). The temporal location of the frames shown in (a) - (c) are marked by dotted lines in (d).

total variation sparsity prior as explained in Section 2.3.2. However, the proposed algorithm is general enough to exploit the sparsity in any transform/operator domain.

2.5 Materials and Methods

2.5.1 Datasets

We study the utility of the proposed k-t SLR scheme in accelerating cardiac perfusion MRI. To validate the method, we use (a) the physiologically improved non uniform cardiac torso (PINCAT) numerical phantom [65,66] and (b) in-vivo cardiac perfusion MRI data. We set the parameters of the PINCAT phantom to obtain realistic cardiac perfusion dynamics and contrast variations due to bolus passage, while accounting for respiration with variability in breathing motion. The contrast variations due to bolus passage are realistically modeled in regions of the right ventricle (RV), left ventricle (LV) and the left ventricle myocardium. To obtain a realistic model, we use the Biot-Savart's law to simulate the spatial distribution of the magnetic flux of the receiver coil [67]. We consider a single coil that is placed on the chest and has the maximum sensitivity to the FOV containing the heart. Here, we use a single slice and assume a temporal resolution of one heart-beat, acquired during the diastolic phase (where the cardiac motion is minimal). The time series data consists

of 70 time frames. We observe that the predominant motion (due to respiration) is in the superior-inferior direction with a low degree of through plane motion in the anterior-posterior direction. The spatial matrix size is 128×128 , which corresponds to a spatial resolution of $1.5 \times 1.5 \text{ mm}^2$. A few slices of this dataset are shown in Figure 2.2.

The in-vivo data was acquired on a 3T Siemens scanner with a saturation-recovery sequence (TR/TE=2.5/1 ms, saturation recovery time=100 ms) at the University of Utah. The study was approved by the institutional review board and written consent was obtained from the subject before the acquisition. The data from a single slice was acquired on a Cartesian grid with a k-space matrix of 90×190 (phase-encodes x frequency encodes) at a temporal resolution of one heart-beat. The subject was instructed to hold the breath for as long as possible. However, the data had significant motion as the subject was not capable of holding the breath for the entire imaging duration.

2.5.2 Comparisons against different methods

We compare the k-t SLR scheme against (a) two-step KLT schemes with different number of phase encodes in the training data (b) the k-t FOCUSS scheme, which relies on sparsity in the x-f space, and (c) variants of the k-t SLR scheme, which rely on only the TV penalty and the spectral penalty alone. Using these comparisons, we mainly seek to verify the following claims:

1. Posing the dynamic reconstruction problem as a spectrally regularized matrix recovery problem provides improved reconstructions over two-step KLT schemes. To verify this claim, we focus on the comparisons between the spectrally regularized matrix recovery scheme (only low rank prior; $p = 0.1$; $\lambda_2 = 0$) and the two step KLT method [36–39] with different training data settings at different accelerations.

2. The exploitation of the sparsity priors, along with the low rank structure, can improve the reconstructions. To verify this claim, we focus on the comparisons of the k-t SLR scheme against regularized schemes that rely only on the spectral ($\lambda_2 = 0$) or TV ($\lambda_1 = 0$) penalty.
3. The k-t SLR scheme can outperform regularized schemes that rely on the sparsity in x-f space. The k-t FOCUSS scheme is known to provide comparable or better performance over all dynamic imaging schemes that use the sparsity in x-f space. We hence compare the k-t SLR scheme against k-t FOCUSS.

The reconstructions are evaluated at a range of acceleration factors denoted by A , which is defined as

$$A = \frac{\# \text{ PEs in fully sampled dataset}}{\# \text{ acquired PEs}}; \quad (2.23)$$

it is the ratio of the number of acquired phase encodes in the fully sampled dataset to the number of phase encodes used to reconstruct the dataset. We quantify the performance of the algorithms using the signal to error ratio (SER) specified as

$$\text{SER} = -10 \log_{10} \frac{\|\mathbf{\Gamma}_{\text{rec}} - \mathbf{\Gamma}_{\text{orig}}\|_F^2}{\|\mathbf{\Gamma}_{\text{orig}}\|_F^2}, \quad (2.24)$$

where $\|\cdot\|_F$ is the Frobenius norm. While this measure provides a quantitative index of performance, it is notorious in being insensitive to artifacts and other distortions. Hence, we also show specific reconstructed frames and the time series data to enable visual comparisons.

The two-step KLT schemes assume a dual density Cartesian sampling pattern. Specifically, the central k-space samples are acquired at the Nyquist temporal sampling rate, while the outer k-space are sampled with a lower-density as shown in Figure 2.5. We consider the KLT scheme with different number of phase encodes

in the training data to analyze the performance dependence of the scheme on the number of samples in the training data. Here, we denote the size of the training data by N_t . The regularized reconstruction schemes such as k-t FOCUSS, k-t SLR, and its variants (spectral penalty alone, TV penalty alone) are capable of accounting for arbitrary non-Cartesian sampling patterns. For these schemes, we consider a radial trajectory with uniform angular spacing; the angular spacing between the spokes is chosen to obtain the specified acceleration factor. The trajectory is rotated by a small random angle in each temporal frame to make the measurements incoherent. By using the equi-angular spacing within each frame, we ensure that the entire k-space is covered uniformly. By considering a small random angle rotation, we not only maintain incoherency, which is required for k-t SLR, spectral penalty and k-t FOCUSS schemes; but also ensure that there are not any sudden jumps across the samples acquired over time, as these jumps could not be optimal for the reconstruction based on only the TV penalty.

We use the NUFFT approximation [68] to realize the \mathcal{A} operator. See Figure 2.5 for an illustration. We add zero mean Gaussian random noise to the measurements in the PINCAT comparisons such that the signal to noise ratio is 46 dB. In the in-vivo comparisons, we resample the uniformly sampled Cartesian data. Hence, we approximate the above radial trajectory with its closest Cartesian trajectory. Specifically, we approximate each k-space location with its nearest neighbor on the Cartesian grid. We chose to use the non convex ($p = 0.1$) spectral penalty to exploit the low rank structure because of its superior performance of suppressing singular values associated with artifacts as opposed to the ($p = 1$), nuclear norm penalty; (See figure (2.4), where the spectral penalty obtains a consistent increase in the SER over nuclear norm at a range of accelerations).

The regularization parameters of the penalized schemes (k-t FOCUSS, k-t SLR, low rank penalty alone, and TV penalty alone) have to be optimized to enable fair

comparisons between the different methods. We determine the optimal regularization parameters such that the SER of the reconstructions are maximized. Similarly, the model order (number of temporal basis functions) of the two-step KLT schemes are chosen such that the SER is maximized. We rely on the fully sampled dataset to compute the SER. Alternate risk functions, which closely approximate the signal to error ratio, have been introduced by [69] for cases when ground truth is not available. We plan to use such risk functions for the selection of the regularization parameters and model-order in the future. We initialize the regularized reconstruction schemes with the gridding solution and iterate the algorithms until convergence.

2.6 Results

We initially demonstrate the utility of the continuation scheme in accelerating the convergence. We then perform quantitative and qualitative comparisons of the proposed scheme with the different methods listed in section V-B on the PINCAT and invivo cardiac perfusion data sets to verify our claims.

2.6.1 Convergence of the algorithm

We first study the convergence of the optimization algorithm in the context of the PINCAT phantom, sampled with 20 k-space spokes/frame; $A = 6.4$ in Figure 2.3. Here, we plot the decrease in the cost function, specified by (2.10), and the improvement in SER of the reconstructed data as a function of the number of iterations. It is seen that for lower values of β_1 and β_2 , the algorithm converges quite fast to the solution of $\mathcal{D}_{\beta_1, \beta_2}$. However, this corresponds to a low SER since the constraints in (2.12) are not satisfied. Increasing the parameters β_1 and β_2 ensures that the constraints are satisfied, but results in slow convergence. We observe that the continuation scheme, where β_1 and β_2 are gradually increased starting from low values, provides a significantly improved convergence rate with good accuracy. In our experiments, we set the TOLERANCE= 10^{-6} to ensure good convergence.

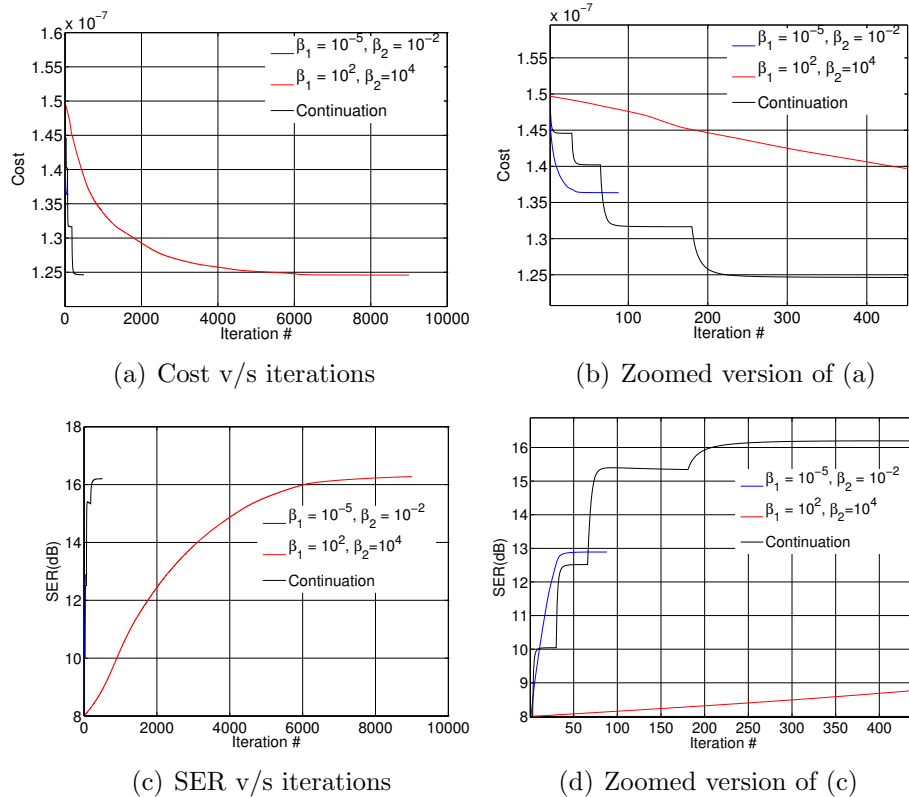


Figure 2.3: Convergence of the proposed continuation scheme as a function of β_1 and β_2 . (a) indicates the evolution of the cost function specified by (2.10), while (b) is zoomed version of (a). The change in signal to error ratio as a function of the iterations is shown in (c) with its zoomed version in (d). Note that the convergence of the algorithm is very slow if these parameters are chosen as high values, which is needed for the constraints in (2.12) to be satisfied. In contrast, the algorithm converges very fast, when these parameters are set to low values. However, the solution of $\mathcal{D}_{\beta_1, \beta_2}$ is a poor approximation for the solution of (2.10). We observe that by properly selecting a continuation scheme, it is possible to significantly improve the convergence rate, while maintaining the accuracy.

2.6.2 Comparisons on the PINCAT phantom

We plot the SER v/s acceleration (A) for the various reconstruction schemes in Figure 2.4. It is seen that k-t SLR method consistently out-performs k-t FOCUSS and the classical KLT-based algorithms by 2-4 dB at most accelerations. We observe that the TV regularization scheme provides comparable SER to k-t SLR at lower acceleration factors, but the performance of the TV algorithm degrades significantly as

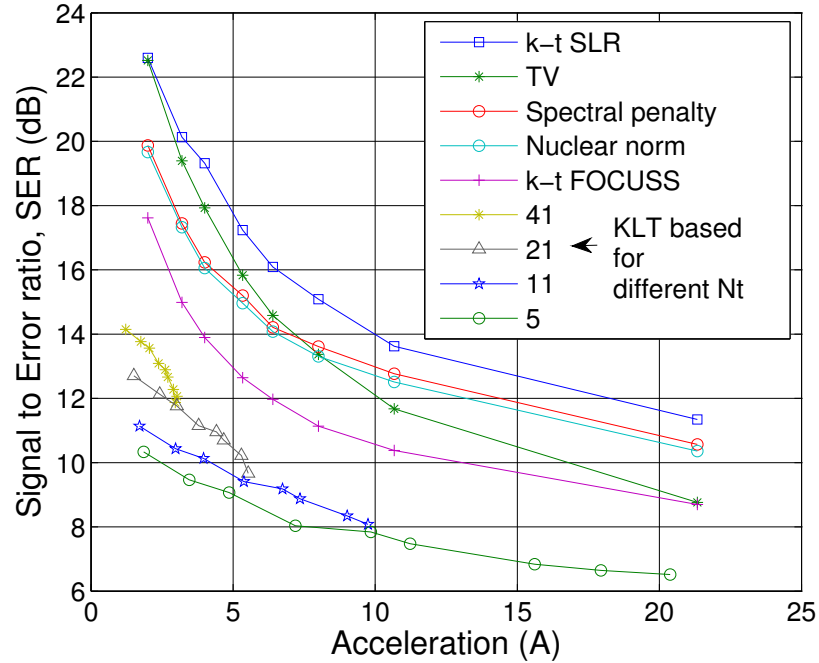


Figure 2.4: Decay of SER as a function of acceleration on the PINCAT data. Note that the k-t SLR scheme provides an improvement of around 2-4 dB over k-t FOCUSS and two-step KLT based schemes at most accelerations. It is seen that the TV scheme provides reconstructions that are similar in SER to the k-t SLR scheme at low accelerations ($A < 4$). However, at higher accelerations, the TV reconstructions exhibit significant over-smoothing and loss of spatial details as seen from Figure 2.6.

the acceleration increase. To enable visual comparisons, we show the reconstructions of the different approaches at $A \approx 5$ in Figure 2.6. The improved reconstructions offered by k-t SLR can be easily appreciated. We now specifically focus on verifying our claims.

2.6.2.1 Utility of the proposed spectrally regularized scheme over two-step KLT methods

We compare the reconstructions of the spectrally regularized algorithm ($p = 0.1; \lambda_2 = 0$) with the two-step KLT approach in Figure 2.5. We set $A \approx 3$ and consider two different choices of the training data. Note that the accuracy of the

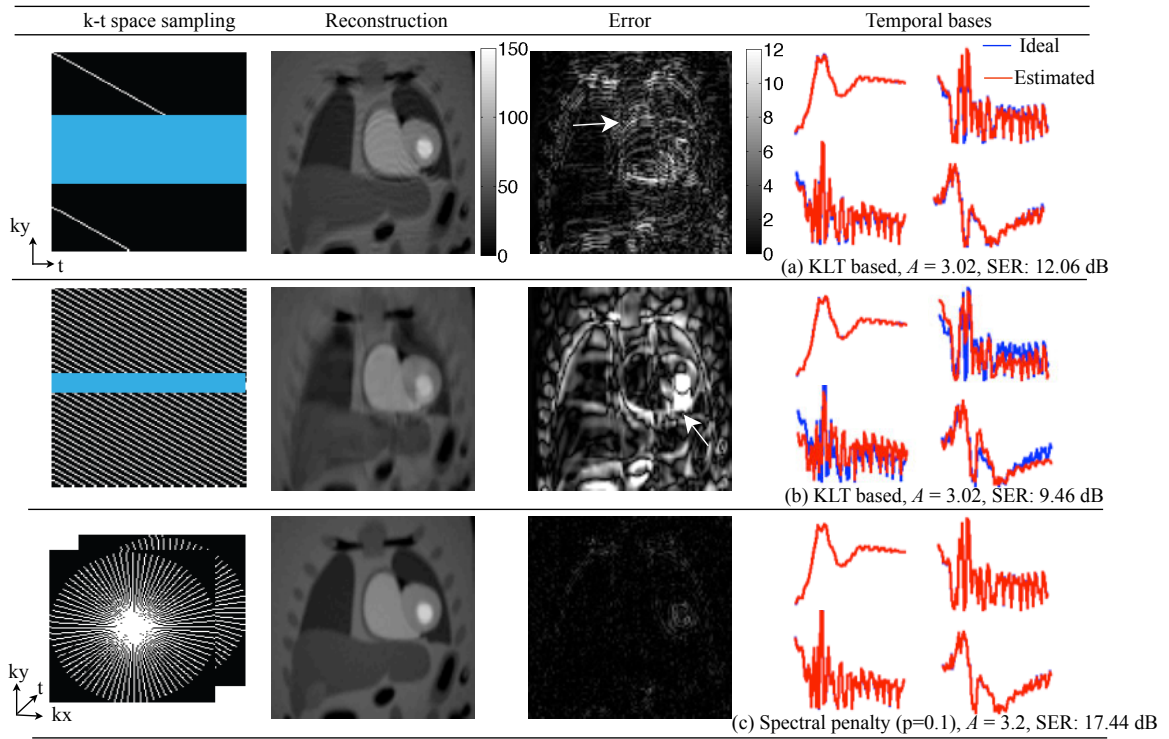


Figure 2.5: Comparison of the two-step KLT schemes (two top rows) with the spectrally regularized reconstruction scheme ($p = 0.1$, $\lambda_2 = 0$), shown in the bottom row. The sampling pattern, the peak LV frame of the reconstructed dataset, the corresponding error image (shown at the same scale in all the insets), and the estimated temporal basis functions ($v_i(t)$, $i = 0$ to 3) overlaid on the actual temporal basis functions are shown in each column. Note that the classical KLT based schemes experience a tradeoff between spatial aliasing and accuracy of temporal modeling. The first row correspond to $Nt = 41$, where the basis functions are estimated correctly. However, the sparse sampling of outer k-space regions results in spatial aliasing, indicated by the dotted arrow. When the number of phase encodes in training data is reduced to $Nt = 5$ in the second row, the temporal basis functions fail to capture the dynamics; this often results in inaccurate temporal modeling of the cardiac motion, especially in regions with significant respiratory motion (denoted by the solid arrow). The spectrally regularized reconstruction scheme, along with the radial sampling pattern, is capable of accurately estimating the temporal bases and spatial weights directly from the undersampled data. The significantly decreased errors with the spectral regularization scheme proves the utility in jointly estimating the temporal basis functions and its spatial weights.

temporal basis functions estimated with the two-step KLT schemes are dependent on the number of phase encodes in the training data. It is seen from the second row

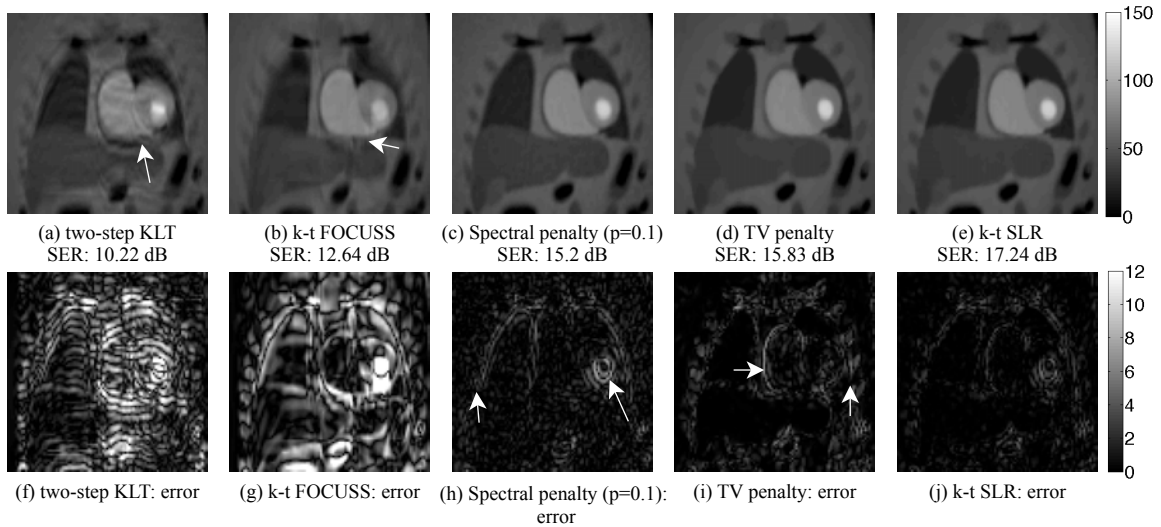


Figure 2.6: Performance evaluation of k-t SLR in comparison with different schemes on the PINCAT phantom: We compare the k-t SLR (fifth column) against the best two-step KLT scheme (first column), k-t FOCUSS (second column) and its own variants i.e, using the spectral penalty alone and the TV penalty alone (third and fourth columns respectively). The two-step KLT scheme assumes $A = 5.28$, while all the other methods provides an acceleration of $A = 5.33$. The reconstructed peak LV uptake frame and the corresponding error images shown in the top and bottom rows respectively. The two-step KLT scheme exhibits incorrect temporal modeling and spatial aliasing, indicated by the arrow in (a)). Since the sparsity in the $x - f$ space is disturbed in the presence of respiratory motion, the k-t FOCUSS reconstructions results in aliasing in regions with significant inter frame motion (see arrow in (b)). The use of the spectral penalty alone resulting in temporal smoothing (dotted arrow in (h)) and residual streaking artifacts due to aliasing (solid arrow in (h)). The use of the TV scheme alone suppresses the spatial aliasing artifacts, while it loses important spatial details due to over smoothing. For instance, the border between the myocardium and the blood pool are smeared, indicated by the dotted arrow in (i) and the details of the ribs are smeared (solid arrow in (i)). By combining the benefits of both low rank and TV schemes, the k-t SLR scheme provides more accurate reconstructions.

of Figure 2.5 that the estimate of the temporal basis functions are poor when the number of phase encodes in the training data is less, resulting in degradations in the temporal dynamics. While the accuracy of the temporal basis functions is improved when the number of phase-encodes in the training data are increased, it comes at the expense of lower density in outer k-space; the lower k-space density results in

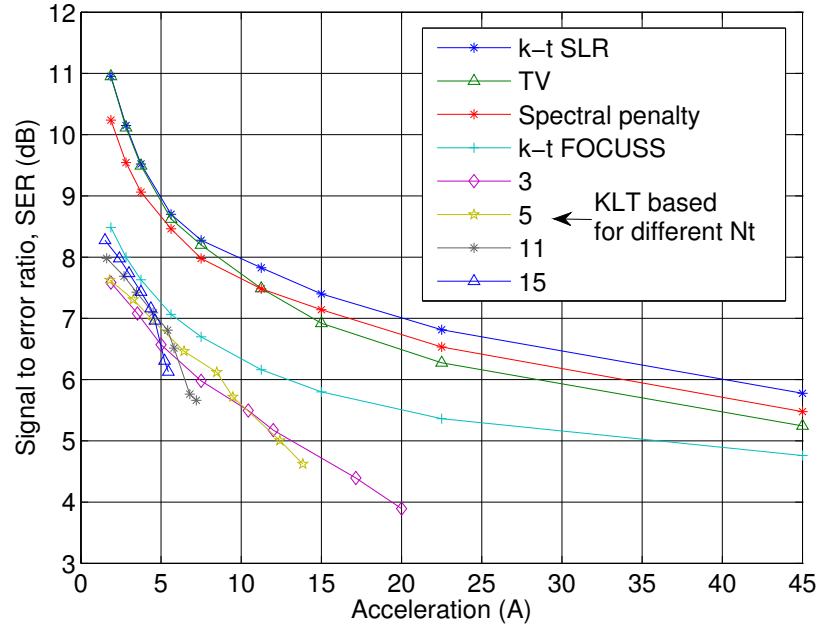


Figure 2.7: Quantitative comparison of the different schemes at a range of accelerations on in-vivo data. It is seen that the k-t SLR consistently provides significantly higher SER over classical KLT based schemes and k-t FOCUSS at all accelerations. The trend of all the methods are similar to the PINCAT comparisons. The use of TV penalty alone is seen to provide comparable SER to k-t SLR until $A \approx 7$. However, note that the SER of the TV scheme is observed to drop at higher accelerations due to over-smoothing, as seen from Figure 2.8.

significant spatial aliasing artifacts in the reconstructions, as seen from the first row of Figure 2.5. Since the spectrally regularized reconstruction algorithm estimates the temporal bases and the spatial weights directly from the undersampled data, the estimates are more accurate as seen from the last row of Figure 2.5.

2.6.2.2 Advantage of exploiting total variation prior, along with the spectral penalty

We direct the readers attention to the last three columns of Figure 2.6 where we study the regularized schemes with TV penalty only, spectral penalty only, and k-t SLR at $A = 5.33$. It is seen that the TV algorithm over smooths the edges of the myocardium in (d)). In contrast, the use of spectral penalty alone results

in reconstructions with unsuppressed spatial aliasing and temporal smoothing (see the residual streaking artifacts in (c) and errors due to the temporal smoothing in (h)). The k-t SLR method, which relies on both spectral and sparsity penalties, significantly reduces these artifacts. It provides a 2 dB improvement in SER over the methods that rely on only spectral or sparsity penalties.

2.6.2.3 Comparison of k-t SLR with k-t FOCUSS (model based $x - f$ scheme)

The second and fourth columns of Figure 2.6 shows the reconstructions of k-t FOCUSS and k-t SLR at $A = 5.33$, using the same k-space trajectory. Since k-t FOCUSS relies on the $x - f$ space sparsity that is degraded in the presence of breathing motion and contrast variations due to bolus passage, the reconstructions exhibits significant aliasing artifacts in regions with high interframe motion.

2.6.3 Comparisons on the in vivo data

We plot the SER of the in-vivo reconstructions as a function of the acceleration in Figure 2.7. The trend is consistent with the PINCAT comparisons. Specifically, the k-t SLR scheme provides a consistent 1-2 dB performance improvement over classical KLT based method and k-t FOCUSS at most acceleration factors. The visual comparisons of all the methods at $A \approx 11$ is shown in Figure 2.8. In figure 2.8, the time profiles of regions within the blood pool and the myocardium are routinely studied and form the basis for perfusion quantification. To evaluate the accuracy in determining these profiles, we also show the plots of the time series of specific regions within the blood pool and the myocardium; For consistently plotting these time series, we initially perform a registration step on the fully sampled data, such that the chosen regions are stationary across the time frames. The deformations from the registration step are then used to wrap the reconstructions. Next, we plot the average signal intensity of the regions at each time frame and obtain the plots in

figure 2.8(e),(f).

Similar to our findings with the PINCAT phantom, we find significant performance improvement of k-t SLR in comparison to the other methods in figure (2.8) that verify our claims in V.B. Specifically, the utility of k-t SLR in obtaining close to accurate time profiles in the blood pool and myocardial regions is of great clinical importance. Any inaccuracies here, could lead to false analysis in the subsequent quantification stages. For instance, the blood pool region time profiles are used to determine the arterial input function, which forms the key component of the model fitting stage in the perfusion quantification. The methods of KLT, k-t FOCUSS and spectral penalty provide inaccurate time profiles in this regard. While TV provides good blood pool curves, it loses its accuracy in determining the time profiles within the clinically relevant myocardial region due to over smoothing; In contrast, k-t SLR provides a close match of its time profiles with that of the fully sampled data.

2.7 Discussion

The quantitative comparisons of the different algorithms on numerical simulations and in-vivo perfusion MRI data clearly demonstrates the ability of the k-t SLR scheme in significantly accelerating cardiac perfusion MRI, while introducing few artifacts. Specifically, it provides consistently improved results over current state-of-the-art approaches such as two-step KLT algorithm and the k-t FOCUSS method, which relies on sparsity in x-f space. Since the proposed scheme learns the temporal basis functions from the data itself, and does not make any assumptions on the x-f space structure; it is capable of exploiting the correlations in the data, even when the dynamics are not periodic. This property makes the proposed framework applicable to arbitrary dynamic imaging problems.

We have cast the low rank property in a non convex form, while the TV sparsity in the convex form. The non convex spectral penalty in our experiments have shown to provide consistently better performance over the convex nuclear norm at a range

of accelerations. The main reason of using the convex TV norm is to be relatively robust to stair case like artifacts that are usually dominant, when one uses the non convex TV norm. The non convex TV sparsity penalty introduce these artifacts in locally smooth regions; and the voxel time series in our application of perfusion MRI are usually smooth along time.

From our results, it is observed that using spectral penalty alone results in significant aliasing artifacts, while the use of TV penalty alone results in considerable spatial smoothing. By constraining the reconstructions, the proposed scheme is capable of providing improved reconstructions. In our results, we observe that the performance of the TV scheme is comparable to that of k-t SLR at lower accelerations. However, the SER of the TV scheme drops significantly at higher accelerations due to excessive spatial smoothing. This behavior is reported by earlier myocardial perfusion MRI schemes that only rely on the TV penalty [70].

In this work, the trade-off between the problem fidelity and convergence rate in the variable splitting strategy is addressed by the use of a continuation scheme. The continuation scheme can have some numerical instabilities at high values of β_1, β_2 . To address this instability, we plan to investigate an augmented Lagrangian strategy as proposed in [71] in the future.

We considered the reconstruction of perfusion dynamics from a single slice of the heart in this chapter. We expect to obtain significantly improved results by jointly recovering multiple slices from 3-D k-space acquisitions. The main reason is the significant redundancy in the temporal profiles between slices, which the k-t SLR scheme is capable of exploiting. We also plan to address the recovery using multichannel data. Several authors have used the spatial diversity of the coil sensitivity profiles to accelerate cardiac MRI [72, 73]. We expect that these extensions will enable k-t SLR to provide robust high-resolution perfusion MRI data from 8-12 slices with a temporal resolution of upto 2-3 frames per second.

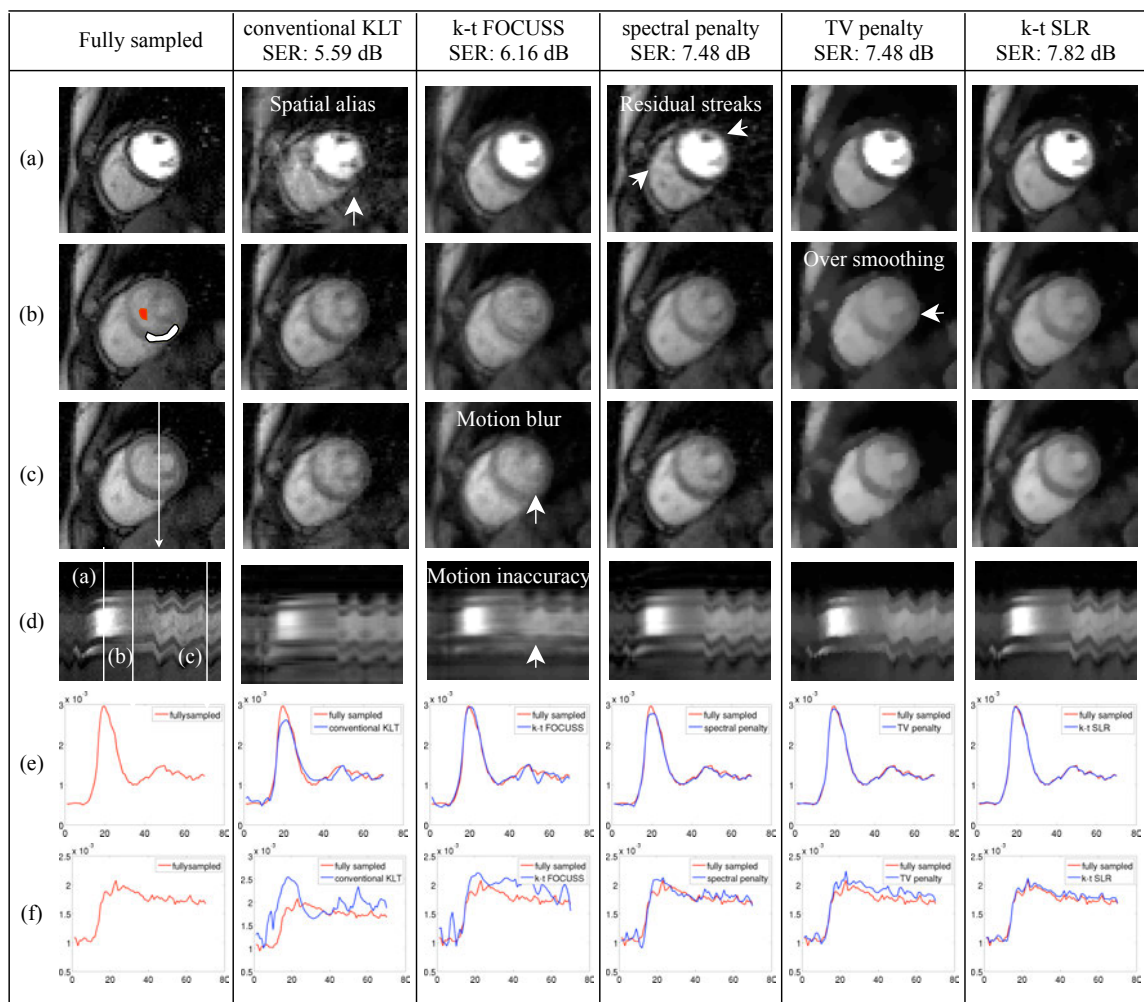


Figure 2.8: Comparisons on in-vivo data: The first column shows the reconstructions of the fully sampled data. Columns 2 to 6 show the reconstructions using the best two-step KLT scheme, k-t FOCUSS method, spectrally regularized, TV regularized, and the k-t SLR scheme, respectively. We choose $A = 11.2$ for all the methods except the two-step KLT, which is at an acceleration of $A = 10.2$. Rows (a) to (c) show different spatial frames. Row (d) shows the image time series plot corresponding to the arrow in (c); Rows (e) and (f) respectively show the averaged signal intensity of the blood pool and myo-cardial regions (denoted in (b)) for the registered reconstructions overlaid on the registered fully sampled data. We observe that the reconstructions with the two-step KLT scheme exhibit spatial aliasing (e and f). The k-t FOCUSS reconstructions exhibit shape distortions and motion inaccuracies. These artifacts can be seen from the time series plots in (d-f), shown in the third column. The spectrally regularized scheme with only the low rank constraint has residual aliasing artifacts as pointed by the arrows in (a), 4th column. This has smoothing along time as well, which can be seen from smoothening of the perfusion peaks in (e) and (f) in the fourth column. The TV penalty scheme has over spatial smoothing and blurring of important structures like the myocardium (see (b), fifth column), due to which the myocardial time series are inaccurate as seen in (f), fifth column. In contrast, k-t SLR in the last column provides efficient reconstructions, with good correlations of the blood and myocardial region time series with the fully sampled data.

CHAPTER 3

ACCELERATING FREE BREATHING MYOCARDIAL PERFUSION MRI USING MULTI COIL RADIAL K-T SLR

3.1 Introduction

Our studies in the previous chapter were based on retrospective resampling of a single coil Cartesian MPI dataset. In this chapter, we exploit the full power of the $k-t$ SLR algorithm by (a) extending it to account for multi-coil acquisitions and to handle different weights for TV in space and time as done in [70], (b) using radial k-space acquisitions, and (c) introducing a novel augmented Lagrangian optimization framework to significantly improve the convergence rate. To exploit the flexibility offered by radial sampling, we customize the sampling pattern to the proposed algorithm using $k-t$ radial sampling experiments on multi-coil data. We use the improved $k-t$ SLR algorithm to achieve quality reconstructions from undersampled radial free breathing MPI datasets. Such acceleration will enable the improvement of volume coverage. We design an experimental paradigm wherein accelerated reconstructions are performed with the multi-coil $k-t$ SLR algorithm by considering subsets of the acquired radial data. The reconstructions are tested by comparisons against the reference fully sampled datasets. We base our studies on rest and adenosine stress datasets acquired from two normal subjects and one patient with myocardial ischemia. We compare the $k-t$ SLR reconstructions with STCR (spatio-temporal constrained reconstruction) and $k-t$ SPARSE with SENSE.

3.2 Theory

3.3 Low rank model representation

In first pass myocardial perfusion imaging, the temporal profiles of pixels corresponding to specific anatomic regions (eg. myocardium, blood pool) are highly correlated. Hence, the temporal profiles of the pixels in the dynamic dataset $\gamma(\mathbf{x}, t)$ can be expressed as a weighted linear combination of few temporal basis functions

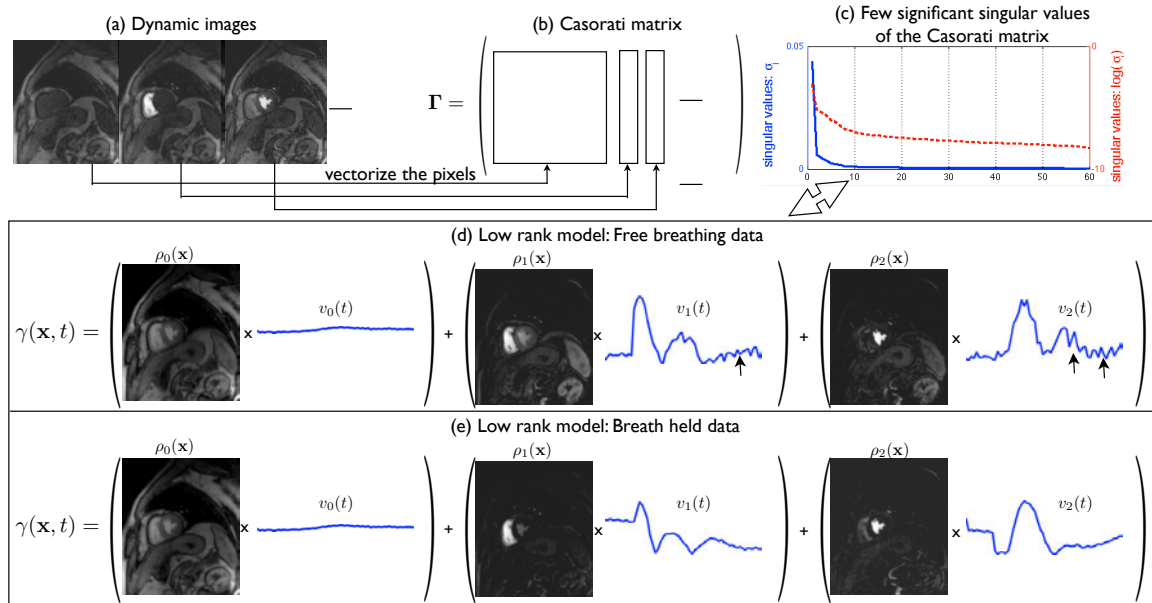


Figure 3.1: The low rank Casorati matrix representation of dynamic data: The pixels in each spatial frame of the dynamic data (a) are vectorized and represented column wise in the Casorati matrix (b). This matrix is low rank (c) which enables the decomposition of the dynamic data as a linear combination of few data derived orthogonal temporal bases (Eq. 3.1). Note from (d) and (e), how the bases adapt to model the respiratory motion in the free breathing data (see arrows in (d) that correspond to motion). Also note that the number of bases will increase with the motion induced temporal variations.

$v_i(t)$:

$$\gamma(\mathbf{x}, t) = \sum_{i=0}^{r-1} \rho_i(\mathbf{x})v_i(t); r \ll n, \quad (3.1)$$

where $\rho_i(\mathbf{x})$ represents the model coefficients and n the number of time frames. r is the number of temporal basis functions and $\mathbf{x} = (x, y)$ is the spatial location. This model thus accounts for the similarity between the time profiles of pixels in specific anatomic regions. The above model implies that the temporal profiles of the pixels lie in a low-dimensional space, which is equivalent to imposing a low-rank constraint on the Casorati matrix $\mathbf{\Gamma}$ [74] (see Fig. 1 and Eq. 3.2). The columns of $\mathbf{\Gamma}$ corresponds to the images at different time instants, while each row of $\mathbf{\Gamma}$ is the temporal profile of the corresponding pixel:

$$\mathbf{\Gamma}_{m \times n} = \begin{pmatrix} \gamma(\mathbf{x}_0, t_0) & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_0, t_{n-1}) \\ \gamma(\mathbf{x}_1, t_0) & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_1, t_{n-1}) \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \gamma(\mathbf{x}_{m-1}, t_0) & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_{m-1}, t_{n-1}) \end{pmatrix}; \quad (3.2)$$

where m denotes the number of voxels in each frame, and n is the total number of time frames. The correlations among the pixel time series result in linear dependencies between the rows of $\mathbf{\Gamma}$ thus resulting in the matrix having a low rank specified by r . Fig. 1 shows the low rank model representations of fully sampled free breathing and breath held MPI datasets. Note that the model is capable of adapting to motion induced intensity variations in the free breathing data. Furthermore, it is also important to note that the number of basis functions required to accurately represent the signal will increase with the motion-induced temporal variations. The proposed $k-t$ SLR algorithm jointly estimates the temporal basis functions $v_i(t)$ and the spatial weights ($\rho_i(\mathbf{x})$) in Eq. 3.1 from the undersampled $k-t$ measurements in an iterative framework by solving a spectrally regularized problem; this will be described in the

next section.

3.3.1 Radial $k-t$ SLR with parallel imaging

The undersampled radial acquisition of sensitivity weighted dynamic perfusion images can be modeled as: $\mathbf{b} = \mathcal{A}(\mathbf{\Gamma}) + \mathbf{n}$, where \mathbf{b} is a concatenated vector containing the measured non-Cartesian noisy $k-t$ measurements for each coil. \mathbf{n} is the additive white noise. $\mathbf{\Gamma}$ is the $m \times n$ Casorati matrix containing the dynamic data as defined in Eq. 3.2 (m is the total number of pixels in a frame and n is the number of time frames). The operator \mathcal{A} models the coil sensitivity encoding as well as Fourier encoding on the specified radial trajectory. We determine the radial sampling pattern that provides the best recovery with the $k-t$ SLR algorithm based on $k-t$ sampling experiments using multi-coil data (section 3.3); specifically, we employ a sampling scheme with golden ratio spacing between successive radial rays.

We formulate the recovery of $\mathbf{\Gamma}$ as a spectral and sparsity regularized optimization problem (Fig.2):

$$\mathbf{\Gamma}^* = \arg \min_{\mathbf{\Gamma}} \underbrace{\|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|_2^2}_{\text{data consistency}} + \underbrace{\lambda_1 \Phi(\mathbf{\Gamma})}_{\text{Schatten } p\text{-norm}} + \underbrace{\lambda_2 \Psi(\mathbf{\Gamma})}_{\text{spatiotemporal TV norm}}. \quad (3.3)$$

Here, the non-convex Schatten p -norm $\Phi(\mathbf{\Gamma})$ is the surrogate for the rank defined as:

$$\Phi(\mathbf{\Gamma}) = (\|\mathbf{\Gamma}\|_p)^p = \sum_{j=0}^{n-1} \sigma_j^p; p < 1; \quad (3.4)$$

where σ_j are the singular values of $\mathbf{\Gamma}$ (elements of the diagonal matrix $\mathbf{\Sigma}_{n \times n}$) in the singular value decomposition: $\mathbf{\Gamma} = \mathbf{U}_{m \times n} \mathbf{\Sigma}_{n \times n} \mathbf{V}_{n \times n}^H$. $\Psi(\mathbf{\Gamma})$ is the spatio-temporal total variation norm and is the surrogate for spatio-temporal smoothness of $\mathbf{\Gamma}$ defined as :

$$\Psi(\mathbf{\Gamma}) = \left\| \left\| \sqrt{|\nabla_x(\mathbf{\Gamma})|^2 + |\nabla_y(\mathbf{\Gamma})|^2 + \alpha |\nabla_t(\mathbf{\Gamma})|^2} \right\|_1 \right\|_1; \quad (3.5)$$

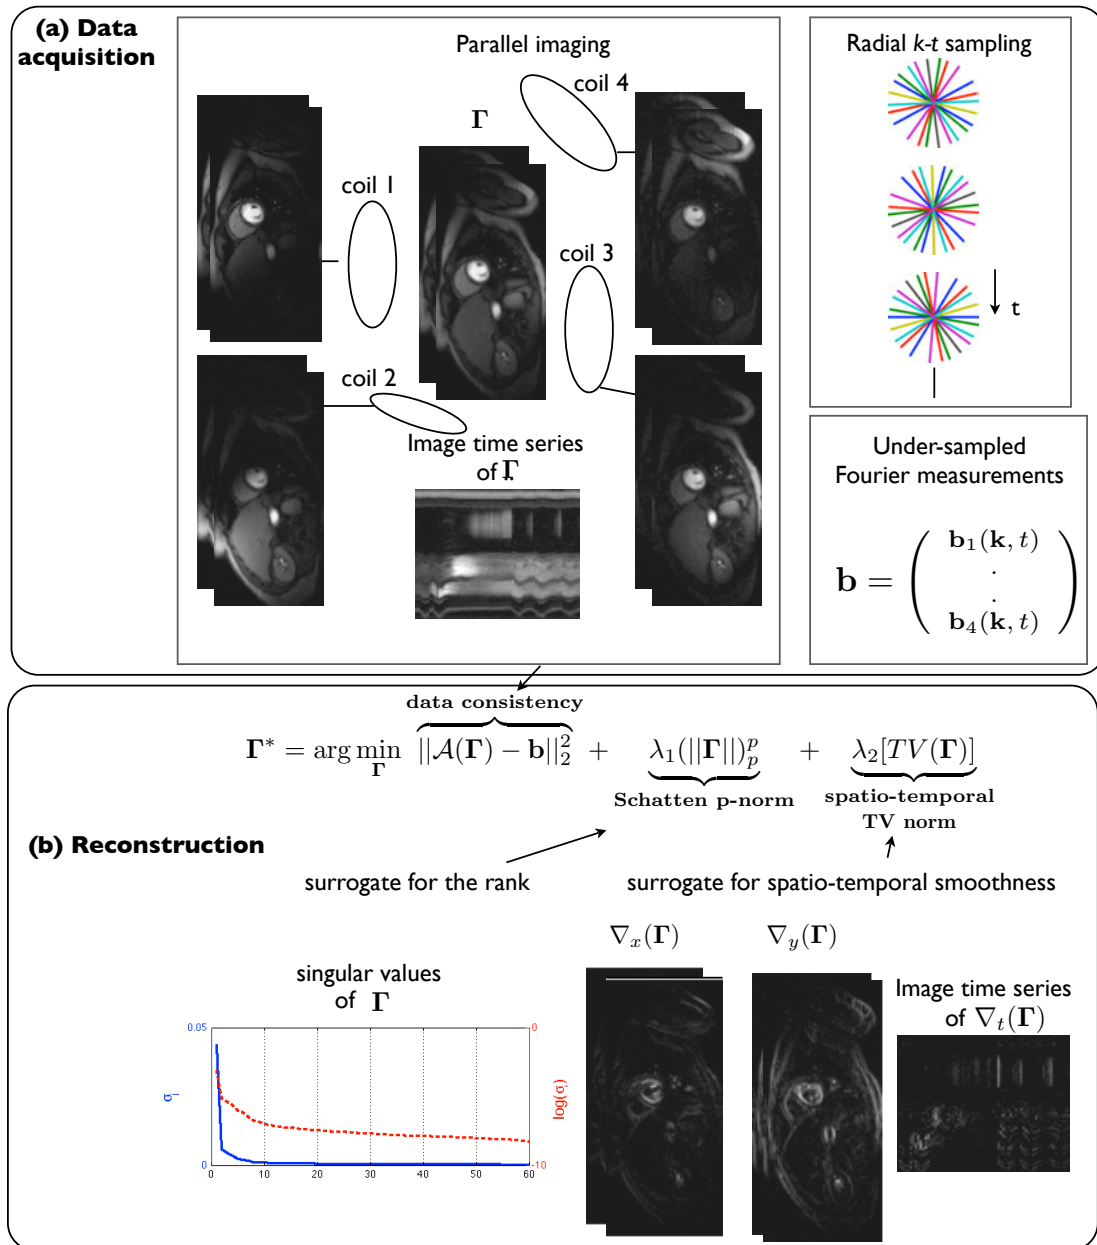


Figure 3.2: $k-t$ SLR with parallel MRI for accelerated imaging: The perfusion images are acquired using multiple coils and radial sampling with golden angle ray spacing (a). $k-t$ SLR exploits the low rank and smooth spatio-temporal properties of perfusion data by utilizing the non-convex Schatten p -norm ($p < 1$) and the spatio-temporal total variation norm (b). The reconstruction in (b) is formulated as a spectral and sparsity penalized optimization problem; the coil sensitivity encoding in combination with radial sampling improves data consistency in the formulation.

where $\nabla_x, \nabla_y, \nabla_t$ are the difference operators along the x, y and t dimensions. The factor $\alpha \geq 1$ controls the relative weight of the temporal and spatial gradients, and $\|\cdot\|_1$ denotes the ℓ_1 norm. λ_1 and λ_2 in Eq. 3.3 are the regularization parameters that control the balance between the two norms and the data fidelity.

3.3.2 Fast augmented Lagrangian (AL) algorithm

The penalty terms in Eq. 3.3 are non-differentiable. Hence, the use of gradient based schemes to solve Eq. 3.3 will result in prohibitively slow convergence. In addition, since a non-convex spectral penalty is used, this approach can result in the solution being trapped in the local minima of the criterion. To overcome these problems, we employ an augmented Lagrangian (AL) optimization algorithm with continuation [75].

A variable splitting technique is used to reformulate the unconstrained optimization problem in Eq. 3.3 to the constrained optimization problem in Eq. 3.6. This splitting enables us to decouple the non-quadratic penalties from the quadratic data term; the complex problem is decoupled into a sequence of simpler subproblems.

$$\arg \min_{\mathbf{\Gamma}, \mathbf{S}, \mathbf{T}} \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|_2^2 + \lambda_1 (\|\mathbf{S}\|_p)^p + \lambda_2 \left\| \sqrt{|\mathbf{T}_1|^2 + |\mathbf{T}_2|^2 + \alpha |\mathbf{T}_3|^2} \right\|_1; \quad (3.6)$$

$$\text{subject to, } \mathbf{\Gamma} = \mathbf{S}; \underbrace{\begin{bmatrix} \mathbf{T}_1 \\ \mathbf{T}_2 \\ \mathbf{T}_3 \end{bmatrix}}_{\mathbf{T}} = \underbrace{\begin{bmatrix} \nabla_x(\mathbf{\Gamma}) \\ \nabla_y(\mathbf{\Gamma}) \\ \nabla_t(\mathbf{\Gamma}) \end{bmatrix}}_{\nabla \cdot \mathbf{\Gamma}};$$

We now use the AL method to solve the above constrained optimization problem. Specifically, the constraints are enforced using quadratic penalty terms and Lagrange multiplier terms \mathbf{X} and \mathbf{Y} , as shown in the appendix. The strength of the quadratic penalty terms are specified by β_1 and β_2 , respectively. The main benefit in using the AL scheme is that β_1 and β_2 does not have to be taken to ∞ to enforce the constraints;

the algorithm converges slowly when β_1 and β_2 are high. We show in the appendix that the AL scheme simplifies to the algorithm shown below (also illustrated in Fig. 9). Note that the algorithm involves the alternation between simple steps, which are implemented efficiently. The pseudo code of the algorithm is given below:

Initialization: $\Gamma_0 = \mathcal{A}^T(\mathbf{b})$, $\beta_1 = \frac{1}{\max(\Sigma_0)}$, $\beta_2 = \frac{1}{|\max(\Gamma_0)|}$; Σ_0 is a matrix containing the singular values of Γ_0 ;

while ($|\text{cost}_n - \text{cost}_{n-1}|/|\text{cost}_n| < 10^{-6}$) ; stopping rule (cost as defined in Eq. (3.3);

$\Gamma_n \leftarrow$ (A.1); regularized SENSE problem solved by conjugate gradients;

$\mathbf{S}_n \leftarrow$ (A.2); singular value shrinkage;

$\mathbf{T}_n \leftarrow$ (A.3); total variation shrinkage;

$\mathbf{X}_n \leftarrow$ (A.4); linear update rule;

$\mathbf{Y}_n \leftarrow$ (A.5); linear update rule;

if ($|\text{cost}_n - \text{cost}_{n-1}|/|\text{cost}_n| < 10^{-1}$)

$\beta_1 = \beta_1 * 1.2, \beta_2 = \beta_2 * 1.2$; continuation

end

end

The algorithm employs a continuation strategy where the β_1, β_2 parameters are initialized to small values, and are gradually incremented when the cost in Eq. (3.3) stagnates to a threshold level of 10^{-1} . This strategy is similar to homotopic like continuation schemes employed to solve non convex problems [76]. We observed the continuation scheme to be a key aspect in avoiding convergence to local minima solutions. The proposed algorithm was implemented on a desktop system with 47 GB

RAM, 24 core Intel Xenon E5645 2.4 GHz processors, an NVIDIA Tesla C2075 (5 GB RAM) graphical processing unit, and Matlab R2012a 64 bit with AccelerEyes Jacket v2.2. Jacket is a library that enables computations on the GPU within Matlab using NVIDIA CUDA. Thanks to the fast convergence of the AL scheme and the fast GPU computations, the reconstructions of large multi-coil data sets (of sizes: $[N_x \times N_y \times N_t \times L : 256 \times 256 \times 60 \times 4]$) with the $k-t$ SLR algorithm takes about 1-2 minutes of run time. In this work we considered pre-interpolation of the radial data onto Cartesian grid points that were within 0.5 unit of a measured sample. This facilitated the use of Fast Fourier Transforms (FFTs) in the forward and backward models of the iterative algorithm. We did not observe any noticeable change in the quality of the reconstructions by using the preinterpolated data with FFTs when compared to using the nonuniform radial data with NUFFTs, INUFFTs in the iterative algorithm, as also noted in [70].

3.4 Materials and Methods

3.4.1 Multi-coil radial acquisition of free breathing myocardial perfusion data

Two normal subjects and one patient with cardiac disease were scanned at the University of Utah, in accordance with the institutional review board. Data was acquired with a perfusion radial FLASH saturation recovery sequence (TR/TE \approx 2.6/1.2 ms, 3 slices per beat, flip angle of 14 degrees, $2.3 \times 2.3 \times 8$ mm voxel size, FOV: 280 mm², Bandwidth 1002 Hz/pixel) on a Siemens 3T Trio scanner [77]. 72 radial rays equally spaced over π radians and with 256 samples per ray were acquired for a given time frame and a given slice. These rays were acquired in an interleaved manner in subsets of 6 rays each. The rays in successive frames were rotated by a uniform angle of $\pi/288$ radians, which corresponded to a period of 4 across time (see Fig. 3.4). Data was acquired with the Siemens cardiac coil array and combined into four coils. For coil sensitivity estimation, the complex valued k

space measurements from each coil were averaged along time. From the resultant time averaged image data, the complex valued sensitivity estimates were obtained by dividing each component coil image by the root of sum of absolute squared intensities from all the coils.

Rest data sets were acquired after a Gd bolus administration of 0.02 mmol/kg. Stress data sets were acquired with an adenosine infusion, where 0.03 mmol/kg of Gd contrast agent was injected after 3 minutes of infusion. A total of three rest and three stress data sets were used in this study. A SENSE based reconstruction with mild regularization based on spatio-temporal total variation (TV) constraints was used to resolve residual aliasing in the acquired data. The regularized SENSE reconstructions still contained background noise which was denoised by using a block matching 4-D (BM4D) denoising filter [78] - these denoised images formed the reference datasets.

3.4.2 Undersampled reconstruction with different algorithms

The acquisition described in the previous section had a compromise in the slice coverage (3 slices). In order to demonstrate that the slice coverage of such an acquisition could be further increased, retrospective accelerated experiments were performed. Undersampled reconstructions were performed by different reconstruction algorithms by considering subsets of the measured data. Specifically, the performance of the $k-t$ SLR, STCR [70], low rank, and $k-t$ SPARSE/SENSE [16] algorithms were compared. The comparisons were done at various acceleration levels by considering different numbers of subsets of the measured data that used 24 to 15 rays/frame. The quality of these reconstructions were evaluated against the above reference datasets. All the algorithms relied on the knowledge of coil sensitivities. STCR was implemented by considering $\lambda_1 = 0$ in $k-t$ SLR. The low rank penalty based reconstruction was implemented by considering $\lambda_2 = 0$ in $k-t$ SLR. $k-t$ SPARSE/SENSE was implemented by minimizing the l_1 norm of the signal in the spatial-spectral $(x-f)$ space.

All the algorithms were optimized for the regularization parameters that gave the maximum signal to error ratio (SER) between the reconstructions and the reference data:

$$\text{SER}_{\text{ROI}} = -10 \log_{10} \frac{\sum_{i=1}^n \left(\frac{\|\Gamma_{\text{recon},i} - \Gamma_{\text{ref},i}\|_F^2}{\|\Gamma_{\text{ref},i}\|_F^2} \right)}{n}; \quad (3.7)$$

where n is the number of time frames. During this optimization, the SER_{ROI} metric was evaluated only in a field of view that contained regions of the heart. This was motivated by recent findings in [79], and by our own experience in determining a quantitative metric that best describes the accuracy in reproducing the perfusion dynamics in different regions of the heart, and the visual quality in terms of preserving crispness of borders of heart, and minimizing visual artifacts due to reconstructions. The details of optimization of the regularization parameters in this work are described in the appendix section.

3.4.3 Simulations to determine an optimal radial sampling trajectory

The quality of the reconstructed data is dependent on the specific sampling strategy. With the objective of choosing a radial pattern that provides good recovery, the performance of different sampling trajectories were studied. As described in section 3.1, the reference 72 ray data sets were acquired using radial rays uniformly spaced within each frame, and uniform rotations across frames. To simulate undersampling, subsets of the acquired data were chosen based on the following three families of radial sampling trajectories (see Fig. 3.4):

1. uniform spacing of radial rays within a frame, with uniform rotations across frames.
2. completely random spacing of radial rays within a frame.
3. golden ratio spacing of $(\pi/1.618)$ between successive rays¹.

¹Since the rays from the 72 ray data were uniformly spaced, subsampling of rays was done such

The performance of the low rank, STCR, $k - t$ SLR reconstruction algorithms were compared with each of the above sampling scenarios. The performance was studied at different undersampling factors by considering 24, 21, 18, 15 rays/frame.

3.4.4 Metrics used for quantitative comparison

The reconstructions of the different algorithms were quantitatively compared based on the following three metrics (also see Fig. 7)

- Signal to Error ratio in a region of interest containing the heart (SER_{ROI}): As described in Eq.(5.9), this metric gives a measure of overall accuracy in reproducing the spatio-temporal dynamics in the regions of the heart.
- Normalized high frequency error metric (HFEN): The HFEN metric gives a measure of the quality of fine features, edges, and spatial blurring in the images. We adapt this metric from [80] which is defined as:

$$HFEN_{ROI} = -10 \log_{10} \frac{\sum_{i=1}^n \left(\frac{\|\text{LoG}(\mathbf{\Gamma}_{ref,i}) - \text{LoG}(\mathbf{\Gamma}_{recon,i})\|_2^2}{\|\text{LoG}(\mathbf{\Gamma}_{ref,i})\|_2^2} \right)}{n}, \quad (3.8)$$

where LoG is a Laplacian of Gaussian filter that capture edges. We use the same filter specifications as in [80]: kernel size of 15×15 pixels, with a standard deviation of 1.5 pixels. We evaluate this metric in a square region of interest containing the heart.

- Signal to error ratio of temporal curves in the left ventricle (SER_{TC}): This metric gives a measure of accuracy in reproducing the temporal dynamics in the left ventricular blood pool and myocardium. It also quantifies temporal blurring. To evaluate this metric consistently, the reference data sets were initially registered to estimate the deformation maps that correspond to breathing motion. For that they approximately follow the golden angle distribution.

registration, we employed the non-rigid demons registration algorithm [81] using the normalized cross correlation as the similarity metric. Starting from the second frame, the deformations were obtained by matching the n^{th} frame in the moving sequence to the $(n - 1)^{th}$ frame of the deformed sequence. The deformation maps were used to warp the reconstructions, after which the time intensity profiles in the region of interest of left ventricular blood pool and myocardium region of interests were evaluated. The SER_{TC} metric is evaluated as:

$$SER_{TC} = -10 \log_{10} \frac{\sum_{i=1}^k \left(\frac{\|TC(\mathcal{W} \cdot \mathbf{\Gamma}_{ref,i}) - TC(\mathcal{W} \cdot \mathbf{\Gamma}_{recon,i})\|_2^2}{\|TC(\mathcal{W} \cdot \mathbf{\Gamma}_{ref,i})\|_2^2} \right)}{k}, \quad (3.9)$$

where TC is an operator that extracts the time curves for a specified pixel in the left ventricle; k is the total number of pixels in the left ventricle. \mathcal{W} is an image warping operator that applies the deformation maps corresponding to breathing motion to the reconstructions.

3.4.5 Qualitative evaluation: clinical scoring

In addition to the quantitative validation, we also performed a qualitative analysis by obtaining clinical scores from a cardiologist with 15 years of cardiac MRI experience. Image quality and artifact assessment was performed on the reference images reconstructed from the 72 ray acquisition, and on the $k - t$ SLR images reconstructed from 24 ray subsampled data. The grading scale was (5-1, highest quality to lowest quality). All the images were presented as 4 image sets with each set containing time series of 3 slices for both stress and rest. Image sets from a patient with decreased perfusion and from two normal subjects were presented to the cardiologist in a blinded fashion in a random order.

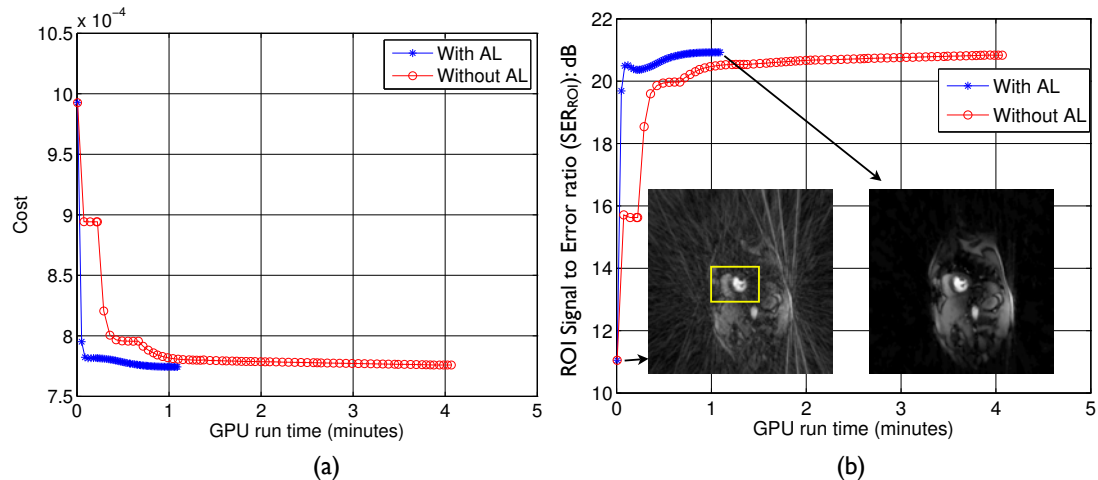


Figure 3.3: Convergence analysis: (a) Cost in Eq.3.3 v/s GPU run time, (b) Region of interest signal to Error ratio (SER_{ROI}) as defined in Eq.5.9 in (dB) v/s GPU run time. From (a), (b), it can be seen that $k-t$ SLR has a faster convergence (by a factor of 4) with the augmented Lagrangian (AL) algorithm in comparison to the previous implementation without the Lagrange multiplier updates. The converged reconstruction in (b) show that the radial streaking artifacts are fully resolved. The SER_{ROI} is evaluated in a square region of interest containing the heart as depicted in (b).

3.5 RESULTS

3.5.1 Convergence analysis

In Fig.3.3, the convergence behavior of the $k-t$ SLR algorithm is shown. Here, undersampled reconstructions were performed with golden ratio sampling using 24 rays/frame. As seen in Fig.3.3, our previous implementation of $k-t$ SLR [82] that relied solely on the increments of $(\beta_1, \beta_2 \text{ towards } \infty)$ had a slow convergence. In contrast, the proposed AL method had a faster convergence, and did not require high values of β_1, β_2 for convergence. The initial values of the continuation parameters were $\{\beta_1, \beta_2 = 10^3, 10^5\}$, while the final values (at convergence) with and without AL respectively were: $\{\beta_1, \beta_2 \approx 10^6, 10^8\}, \{\beta_1, \beta_2 \approx 10^9, 10^{11}\}$. From Fig.3.3, it can also be seen that the streaking artifacts were fully resolved with the proposed $k-t$ SLR

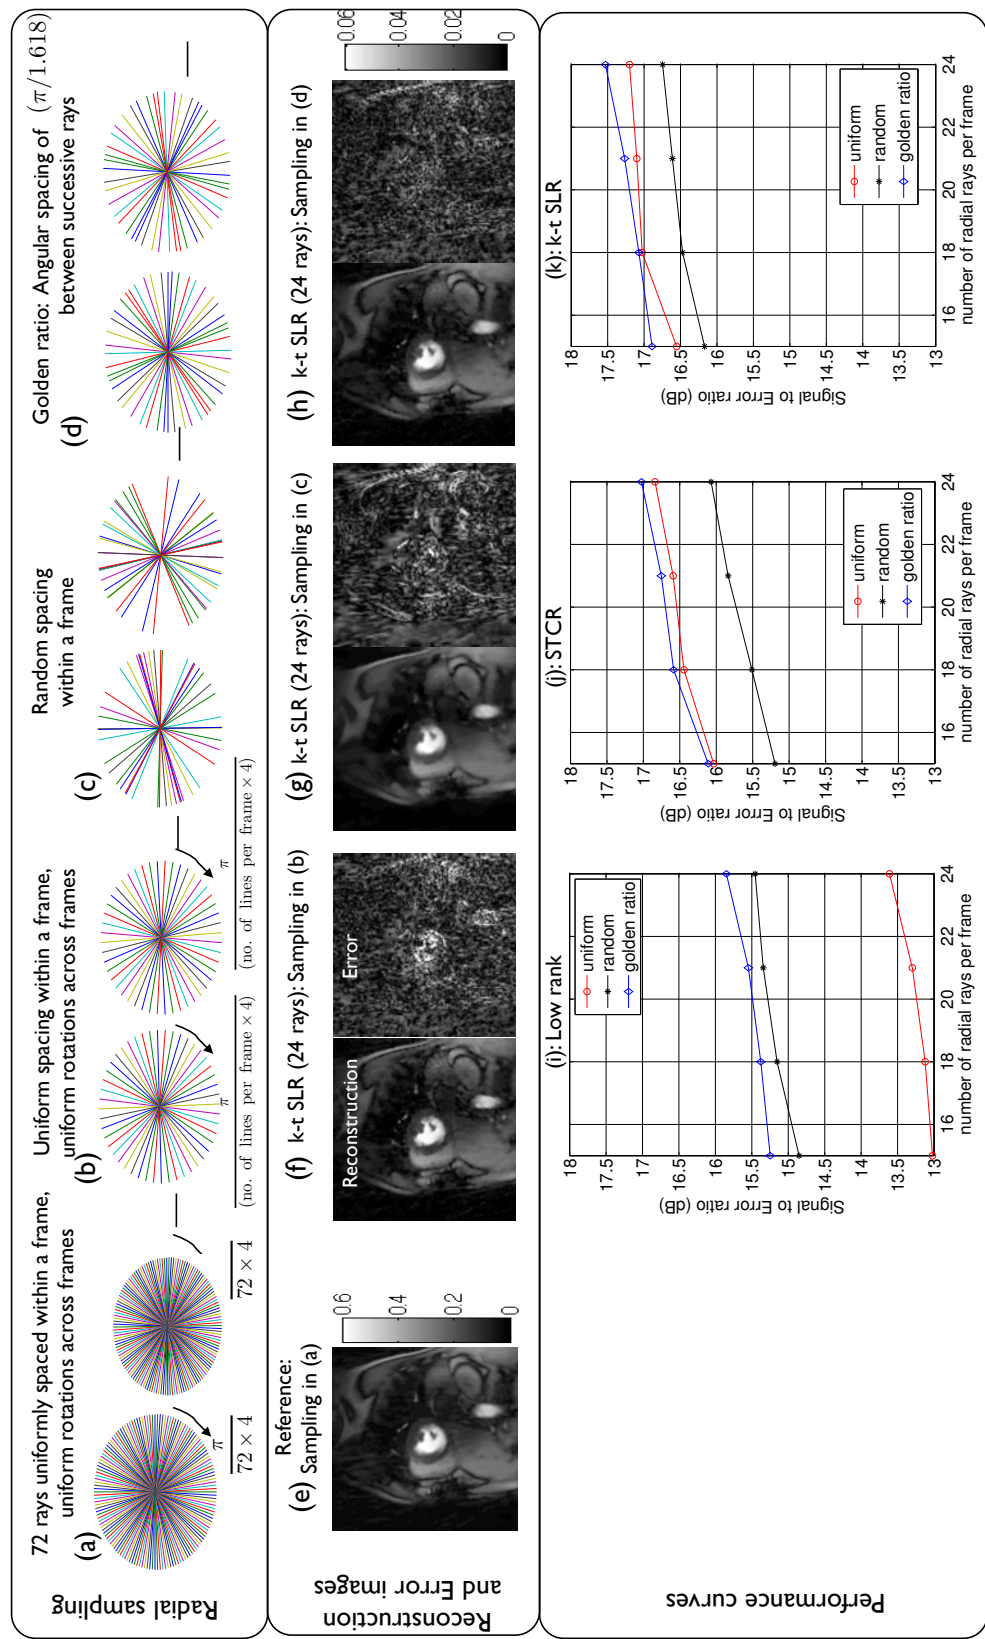


Figure 3.4: Performance of different radial sampling schemes: The different $k - t$ radial sampling schemes along with the corresponding reconstructions are respectively shown in the first and second rows. The third row shows the signal to error ratio plots for the reconstruction algorithms at different acceleration rates. As seen from the third row, the golden ratio sampling gave the best performance with all the algorithms. From the second row, with $k - t$ SLR, uniform sampling was suboptimal as the conditions of incoherent sampling were not met. Random sampling was suboptimal since some regions of k -space are under-sampled. The golden ratio sampling satisfied the requirements of incoherency and uniformity and was found to be optimal.

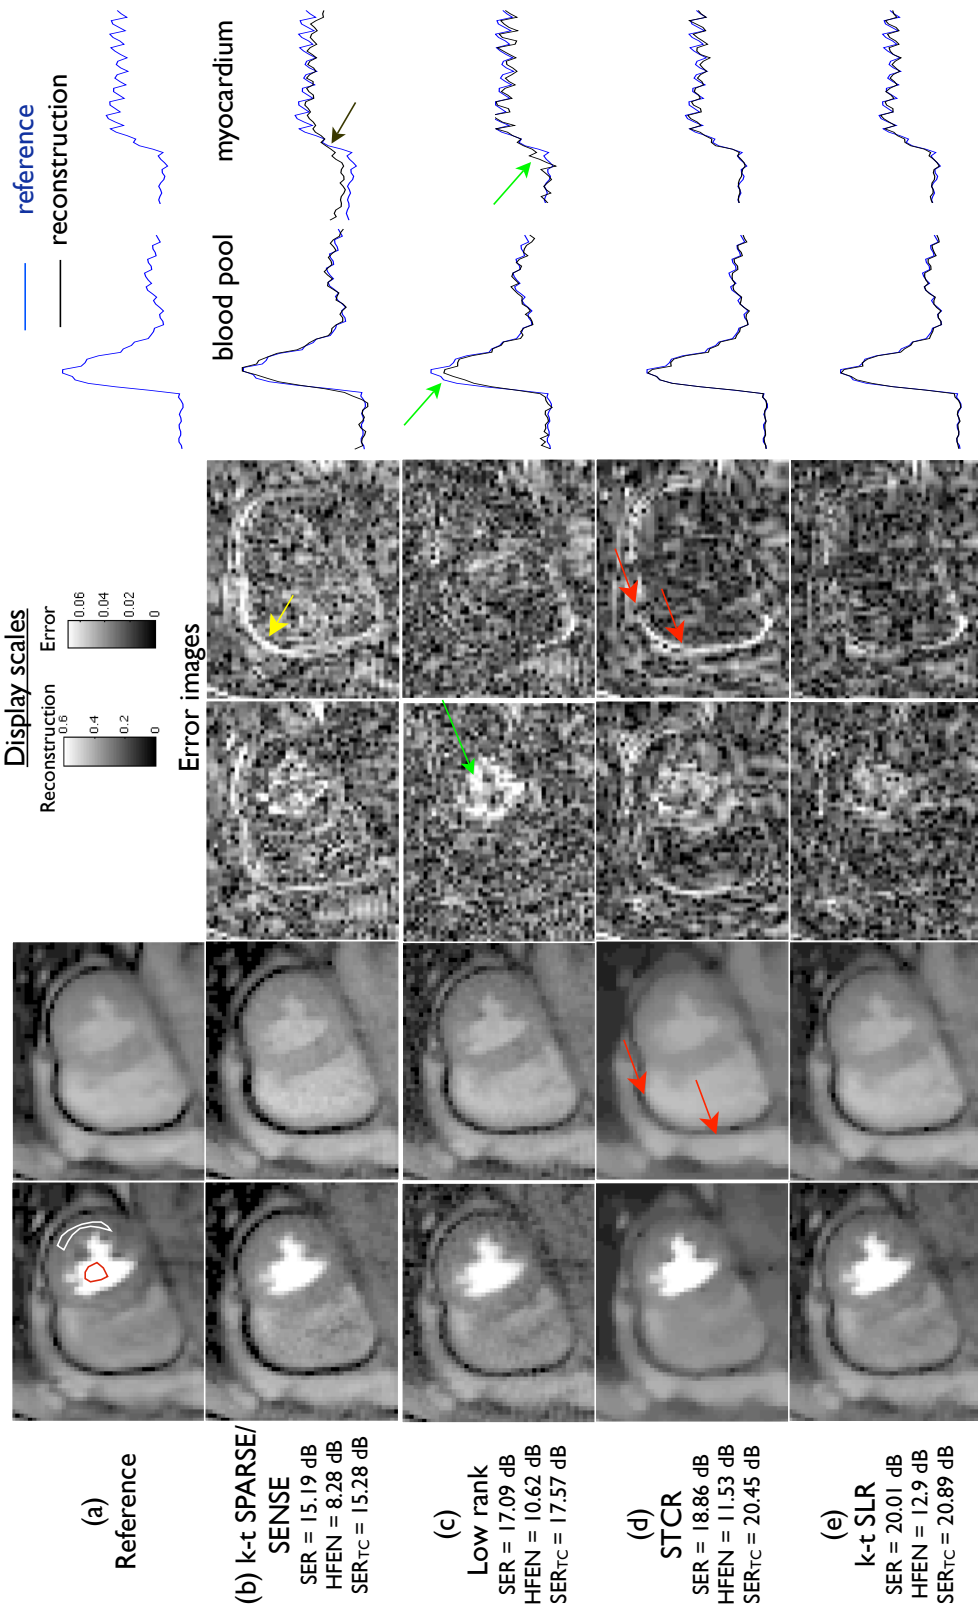


Figure 3.5: Comparisons of different MPI algorithms on a rest dataset with breathing motion. Each row shows the reconstructions, error images and time profile curves for the different algorithms. The image frames in the first and second columns respectively correspond to the peak left ventricular blood enhancement, and peak myocardial wall enhancement instants. The time curves are shown after averaging the signal intensity of the blood pool and myocardial regions (as denoted in (a)) for the registered reconstructions overlaid on the registered reference data. The $k-t$ SPARSE/SENSE method was sensitive to motion and resulted in temporal blurring (see arrows in (b)). The low rank model yielded noise enhancement and temporal blur due to ill-conditionness (arrows in (c)). STCR was relatively robust to motion, however suffered from loss of resolution especially near the edges (see arrows in (d)). $k-t$ SLR was found to maintain a good compromise between spatial and temporal blurs.

algorithm.

3.5.2 Simulations to determine an optimal radial trajectory

The comparisons of different radial sampling methods are shown in Fig. 4. From the SER plots, it is observed that the low rank method does not perform as well with uniform sampling. This is expected since uniform sampling results in more repeated k-space measurements at the same locations in k-space and has less incoherency. In contrast, the STCR method is insensitive to the pattern, as long as the completely random pattern is not used. When the low rank and total variation penalties are merged into $k - t$ SLR, the golden ratio patterns provide better results than the completely random sampling pattern or the uniform pattern. This observation is consistent with the findings reported in the context of standard compressed sensing [83], [84]. From these simulations, the golden ratio sampling pattern was found to be optimal, and therefore was used for undersampling in all the algorithms.

3.5.3 $k - t$ SLR compared to other MPI acceleration schemes

In Fig. 3.5, the comparisons of MPI reconstructions using different algorithms from 21 ray undersampled data are shown. These comparisons are from a rest acquisition on a normal subject with breathing motion. The $k - t$ SPARSE/SENSE method was observed to be sensitive to the breathing motion, and yielded motion related artifacts as depicted both in the temporal curves and the error images of Fig. 3.5 (b). The low rank reconstruction was more robust to breathing motion, when compared to $k - t$ SPARSE/SENSE. However, it had poor temporal fidelity especially during the contrast uptake frames. This is depicted from temporal curve blurring during contrast uptake, and also in the error images of Fig. 3.5(c). STCR had better temporal fidelity, thus preserving contrast dynamics and motion. However edge blurring and patchy artifacts were evident. The $k - t$ SLR algorithm preserved the temporal

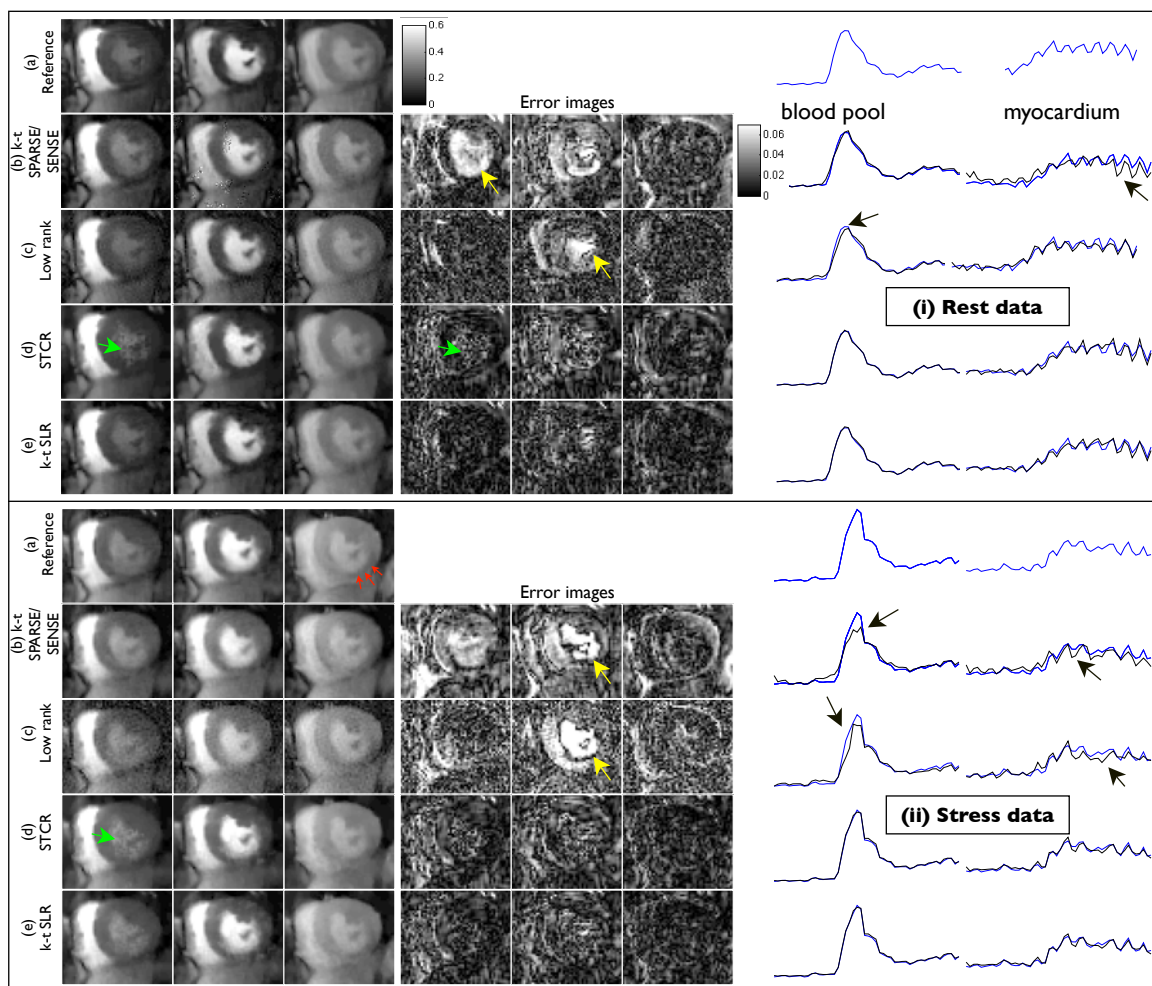
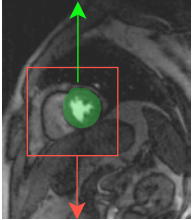


Figure 3.6: Example comparisons of different MPI acceleration algorithms using rest (i) and stress (ii) perfusion data from a patient with myocardial ischemia. The image frames in the first three columns respectively correspond to peak right ventricular blood enhancement, transition between right ventricle and left ventricle, and peak myocardial wall enhancement. During stress, the patient showed reduced contrast uptake in the inferior myocardium wall due to ischemia (see red arrows in (ii.a)). The time curves correspond to the regions depicted in i.a, ii.a). The $k-t$ SLR reconstructions were observed to be less sensitive to artifacts observed with other methods. Specifically $k-t$ SPARSE/SENSE yielded motion blur artifacts (see arrows in i.b, ii.b, and the time curves), low rank method had some temporal blur especially during peak contrast frames (see arrows in ii.c, ii.d, and the time curves). STCR exhibited patchy artifacts (see arrows in (i.d) and (ii.d)). $k-t$ SLR had better quality across all frames in comparison to the other methods. We also observe that the regions of low contrast uptake to be well preserved in the $k-t$ SLR reconstructions.

Subject id	Rest/Stress	Quantitative metric (dB)	k-t SPARSE/SENSE	Low rank	STCR	k-t SLR
Subject 1 (Patient with ischemia)	Rest	SER _{ROI}	15.85	18.43	20.19	21.96
		HFEN _{ROI}	7.55	11.82	12.19	13.00
		SER _{TC}	16.62	21.06	22.45	23.88
	Stress	SER _{ROI}	15.91	18.58	21.50	22.49
		HFEN _{ROI}	7.77	11.26	12.35	13.25
		SER _{TC}	14.62	16.58	22.73	23.37
Subject 2 (Normal)	Rest	SER _{ROI}	15.18	17.09	18.86	20.01
		HFEN _{ROI}	8.28	10.62	11.53	12.90
		SER _{TC}	15.28	17.57	20.45	20.89
	Stress	SER _{ROI}	15.86	17.88	19.1	19.56
		HFEN _{ROI}	8.09	10.72	11.21	11.68
		SER _{TC}	16.49	18.11	19.40	19.40
Subject 3 (Normal)	Rest	SER _{ROI}	16.09	17.08	19.28	19.98
		HFEN _{ROI}	7.46	9.22	11.56	12.08
		SER _{TC}	17.06	19.22	23.31	23.30
	Stress	SER _{ROI}	16.5	17.33	20.02	20.43
		HFEN _{ROI}	7.5	9.78	12.11	12.31
		SER _{TC}	17.43	20.52	23.35	23.31

Left ventricular region where SER_{TC} is evaluated



Region of interest (ROI) where SER_{ROI} and HFEN_{ROI} are evaluated

Figure 3.7: Quantitative comparison of reconstructions from undersampled radial data (21 rays/frame) using different algorithms.




Figure 3.8: Peak myocardial wall enhancement images from (a) reference (72 rays/frame), and (b) $k-t$ SLR (24 rays/frame) reconstructions. Subtle ischemic areas were identified by the cardiologist in the inferior myocardial wall on the second slice in both the reconstructions (see arrows in a,b). These regions were subsequently classified as infarcted regions from the gadolinium enhancement observed in the delayed enhancement image (c).

fidelity and had less smoothing of edges and less patchy artifacts as depicted in Fig. 3.5(e). From this figure, it is also observed that the performance of $k - t$ SLR is comparable to that of STCR in the frames corresponding to contrast uptake. This is expected since the presence of contrast makes the images more or less piece wise constant and the SNR high. However, $k - t$ SLR provides better suppression of artifacts in the pre and post contrast frames. Additionally, from Fig.3.5, it can be seen that the quantitative metrics correlate well with the visual comparisons. Specifically in comparison to $k - t$ SLR, the SER_{TC} metric were low with the $k - t$ SPARSE/SENSE and low rank methods due to temporal blurring, while the $HFEN_{ROI}$ metric in STCR was low due to spatial smoothing. The SER_{ROI} metric was higher with $k - t$ SLR than the other algorithms due to a better overall image quality.

In Fig. 3.9, the comparisons of the algorithms on rest and stress data sets from a patient with ischemia are shown. Under stress conditions, this patient exhibited a reduction in the uptake of the contrast dynamics in the inferior wall of the myocardium. The patient was able to breath normally during rest, however breathed heavily during stress. The reconstructions are shown using 21 radial rays/frame. We observed similar trade offs amongst the methods. Specifically, $k - t$ SPARSE was sensitive to breathing, and the low rank method yielded reduced temporal fidelity during contrast uptake. STCR showed patchy artifacts in some frames, but was robust to motion. $k - t$ SLR was robust to motion and had less patchy artifacts. Similar to Fig. 3.5, all the quantitative metrics correlate well with the visual observations.

Figure.3.7 summarizes the quantitative comparisons of all the algorithms across all the datasets using 21 rays/frame. From this figure, it is observed that the performance of $k - t$ SLR was consistently better than the other algorithms across all the datasets.

3.5.4 Qualitative evaluation by a cardiologist

The clinical scores are presented in table 1. With the patient, the clinician was able

to identify the ischemia in the inferior myocardium wall in both the reference and $k-t$ SLR reconstructions (see Fig. 3.8). More specifically, he confidently identified reduced blood flow in both the rest and stress scans of the reference data set. With the $k-t$ SLR images, he found the ischemic defect to be evident in the stress reconstruction, and border line positive in the rest reconstruction. After looking at the delayed enhancement images, the clinician classified these ischemic regions as infarction. With the normal subjects, the quality scores of the $k-t$ SLR reconstructions were in close agreement with the reference images. No dark rim artifacts were observed in all the images. Slight residual streaking artifacts were present with the $k-t$ SLR method in one of the normal subjects. However, these artifacts were outside the field of view of the heart, which the cardiologist was not very concerned about.

3.5.5 Prospectively undersampled radial data:

An experiment with multi slice acquisition using prospectively undersampled radial data was performed. Data was acquired with the same FLASH sequence as described above, although 5 slices were imaged after each saturation pulse [70]. Seven slices per beat, and 30 rays/frame with a golden angle spacing between successive rays were acquired using a 32 coil cardiac array. Here, golden ratio sampling was employed based on our $k-t$ sampling experiments. A healthy subject was scanned during rest conditions as described above. The subject was not able to maintain a steady breath hold during the entire scan, and the data contained motion. Prior to reconstruction, a PCA based coil compression strategy [85] was used to compress the 32 coil data set to four principal component data set. The PCA coil compression was used to facilitate reconstruction without incurring a high computation cost and memory requirements associated with the 32 coils. As reported in [86], and also from our observation, about 4-5 principal components were enough to capture most of the variance in the data, which motivated us to use 4 coils for processing this data.

Subject id	Reconstruction	Rest	Stress	Presence of disease
		Quality score	Quality score	
Subject 1	Reference	4	4	positive
Subject 1	k-t SLR	4	4	positive
Subject 2	Reference	4	4	negative
Subject 2	k-t SLR	3.75	3.75	negative
Subject 3	Reference	4	4	negative
Subject 3	k-t SLR	4	4	negative

Table 3.1: Quality scores from a cardiologist on three subjects

3.5.5.1 $k - t$ SLR reconstructed multi-slice images

In the below figure, we show the $k-t$ SLR reconstructed images from the above acquisition experiment with multiple (seven) slices coverage using 30 rays/frame. We observe that the $k-t$ SLR reconstructions provided improved heart coverage with good spatio-temporal resolution.

Acquiring multiple slices during a single saturation pulse has shown to be more efficient as it minimizes the time spent for saturation [9, 70]. However, this comes at an expense of different slices having different saturation recovery times, and hence different intensities as seen in the figure below. As recommended in [70], the slices could be analyzed separately and be scaled to have similar intensities.

3.6 Discussion

In this chapter, the feasibility of $k - t$ SLR in providing robust free breathing MPI reconstructions at high acceleration levels was evaluated. This study considered retrospectively accelerating free breathing MPI datasets that were acquired using 72 radial rays/frame at 2.3 mm^2 in plane resolution and 3 slices. The results of obtaining good fidelity $k - t$ SLR reconstructions from highly undersampled data suggest that $k - t$ SLR could be used to improve the slice coverage and the spatial resolution




Figure 3.9: Example k-t SLR reconstructed multi slice 2D first-pass myocardial perfusion images on a healthy subject. The numbers on the images indicate the time instants of the dynamic perfusion acquisition. The subject could not maintain a steady breathhold during the entire scan, due to which the data contained motion content. These reconstructions demonstrate good spatial and temporal resolution with extended volume coverage of the heart.

in a prospective acquisition. For a reliable quantification of perfusion parameters from free breathing MPI data, it should be noted that the breathing motion should be compensated. In this work, we used a basic non-rigid registration algorithm to correct for the breathing motion, and analyzed the temporal curves in the registered reconstructions. The results show that the temporal profiles from the undersampled $k - t$ SLR reconstructions match well with the temporal profiles from the reference reconstructions. This suggests that the good temporal fidelity of the $k - t$ SLR reconstructions may lead to reliable estimation of quantitative perfusion parameters.

All the reconstruction algorithms considered in this work were based on non-linear reconstruction strategies. Unlike the classical linear k-t reconstruction strategies, the performance of these algorithms cannot be characterized by a single point spread function [87]. To analyze the performance of the nonlinear algorithms, this study relied on quantitative metrics that gave a measure of overall spatio-temporal fidelity, image sharpness, and temporal accuracy. Our comparisons against MPI accelerated schemes show that the $k - t$ SLR scheme is capable of reducing motion blurring and edge smoothing artifacts. In general, the $k - t$ SLR algorithm benefited from total variation sparsity regularization in being robust to temporal and spatial smoothing. The STCR method performed well in regions where the signal was piece wise constant both spatially and temporally - or equivalently during peak contrast frames and datasets with less motion. However, it was observed during the pre and post contrast frames, STCR yielded patchy artifacts and edge blurring. In such scenarios, $k - t$ SLR provided more robust and natural textures and less edge blurring. During the peak contrast frames, the use of the low rank regularizer alone yielded temporal blur, which was minimized with $k - t$ SLR.

During this study, we observed that performance evaluation using the noisy SENSE based reconstructions as reference datasets were not conclusive since the noise in these reconstructions were higher than the subtle differences between the different recon-

struction methods. To address this, we denoised the SENSE based reconstructions using the block matching 4D (BM4D) denoising algorithm. The BM4D algorithm is reported to give state of the art denoising performance. The algorithm performs denoising by exploiting nonlocal similarities of spatio-temporal patches. Since the BM4D algorithm is very different from all the reconstruction algorithms considered in this work, the performance comparisons are free from any bias. In addition, the BM4D algorithm has an automatic selection of parameters based on the estimation of the noise level, which minimizes the risk of subjectivity.

During the review of the contents of this chapter submitted as a paper, it was suggested that breath held data sets from a second injection could be used as reference ground truth images. The main challenge however would be to perform a good registration between the undersampled free breathing reconstructions and the breath held datasets for a head to head comparison; this is difficult due to out of plane motion, especially when only a few slices are imaged. In addition, any residual contrast from the first bolus may bias the comparisons. In this context, we believe that the usage of the free breathing datasets as reference sets would better fit to the goals of the current work.

In this study, optimization of λ_1, λ_2 in $k-t$ SLR was performed with a fixed value of $p = 0.1$ and $\alpha = 4$. The choices of p and α were motivated by empirical observations and worked well in practice for free breathing MPI data. A thorough search in a 4-dimensional space of $p, \alpha, \lambda_1, \lambda_2$ could improve the $k-t$ SLR reconstructions. The automatic tuning of regularization parameters for iterative nonlinear reconstruction algorithms is an actively researched area. There exists some strategies such as cross validation [88], and Stein unbiased risk estimator (SURE) methods [89]. In the future, we plan to investigate the adaption of one of these methods to our setting.

The performance of all the nonlinear iterative reconstruction algorithms in this study were evaluated based on quantitative metrics that gave measures of image

sharpness, temporal blurring, and overall mean square error. Recently, the resolution of reconstructed images from nonlinear algorithms were characterized by determining the local point spread functions at every image pixel [90]. Such an analysis could be adapted to our setting to characterize the resolution of the images from the different algorithms.

The augmented Lagrangian optimization algorithm used in this study was found to provide fast convergence. Speed up factors of about four were observed when compared to the previous implementation of $k-t$ SLR. In this study, we used a simple sum of squares approach [91] to estimate the coil sensitivities from time averaged data. In the future, we plan to consider other extensions for better estimations such as moving window approaches for time varying sensitivities, and/or joint estimation of the sensitivities along with the reconstructions [92].

The current study has limitations in that data from only three subjects were used for analysis. To fully evaluate the clinical feasibility of $k-t$ SLR and draw statistical conclusions, a study with a cohort of patient datasets is needed with validation against gold standard coronary x-ray angiography.

3.7 Appendix

3.7.1 Augmented Lagrangian (AL) algorithm steps

In this section, the derivation of the AL subproblems is described. Referring to Eq. 3.6, the constraints are enforced using Lagrange multiplier terms and quadratic penalties. The resulting optimization objective (termed as the AL function) is specified by:

$$\begin{aligned} \mathcal{D}_{\beta_1, \beta_2}(\mathbf{\Gamma}, \mathbf{S}, \mathbf{T}; \mathbf{X}, \mathbf{Y}) = & \quad (3.10) \\ & \|\mathcal{A}(\mathbf{\Gamma}) - \mathbf{b}\|_2^2 + \lambda_1 (\|\mathbf{S}\|_p)^p + \lambda_2 \left\| \sqrt{|\mathbf{T}_1|^2 + |\mathbf{T}_2|^2 + \alpha |\mathbf{T}_3|^2} \right\|_1 \\ & + \frac{\beta_1}{2} \|\mathbf{\Gamma} - \mathbf{S}\|_2^2 + \frac{\beta_2}{2} \|\nabla \cdot \mathbf{\Gamma} - \mathbf{T}\|_2^2 \end{aligned}$$

$$+\beta_1 \langle \mathbf{X}, \mathbf{\Gamma} - \mathbf{S} \rangle + \beta_2 \langle \mathbf{Y}, \nabla \cdot \mathbf{\Gamma} - \mathbf{T} \rangle;$$

where \mathbf{X}, \mathbf{Y} are matrices of Lagrange multipliers. β_1 and β_2 are the penalty parameters that determine the equivalence of Eq. 3.10 to Eq. 3.6, and hence the original problem in Eq. 3.3. In our earlier implementation [82], we only relied on the quadratic penalty terms in Eq. 3.10 (second line of Eq. 3.10) to enforce the constraints in Eq. 3.6 due to which β_1 and β_2 were tended to ∞ and resulted in slow convergence. The main advantage of using the Lagrange terms (last line of Eq. 3.10) rather than enforcing the constraints using penalties alone is that the parameters β_1, β_2 need not tend to ∞ for the constraints in Eq. 3.6 to hold, which allows for a faster convergence.

All the five variables in Eq. 3.10 are estimated using an alternating minimization algorithm. Specifically, we minimize the AL objective function in Eq. 3.10 alternately with respect to one variable at a time, assuming the other to be fixed. This approach simplifies the original problem to a sequence of well understood sub-problems. These subproblems are shown in Fig.9. In essence, the algorithm cycles through: (a) regularized SENSE problem solved by conjugate gradient algorithm, (b) singular value shrinkage, (c) total variation shrinkage, and (d,e) linear update rules of the Lagrange multipliers. Additionally, a continuation strategy is employed where the parameters β_1 and β_2 are initialized with small values and are gradually incremented. This continuation strategy was observed to be a key aspect in avoiding convergence to local minima [19].

3.7.2 Choosing the regularization parameters

The $k - t$ SLR algorithm depends on four parameters: $\lambda_1, \lambda_2, \alpha, p$. Since it is impractical to tune for these parameters in a four dimensional space, we restrict ourselves to a simpler approach. The values of p and α were fixed to $p = 0.1$ and $\alpha = 4$ based on empirical observations on free breathing MPI data. With the fixed values of α, p , we tune for λ_1, λ_2 in a 2D space. We tuned for λ_1, λ_2 for the rest

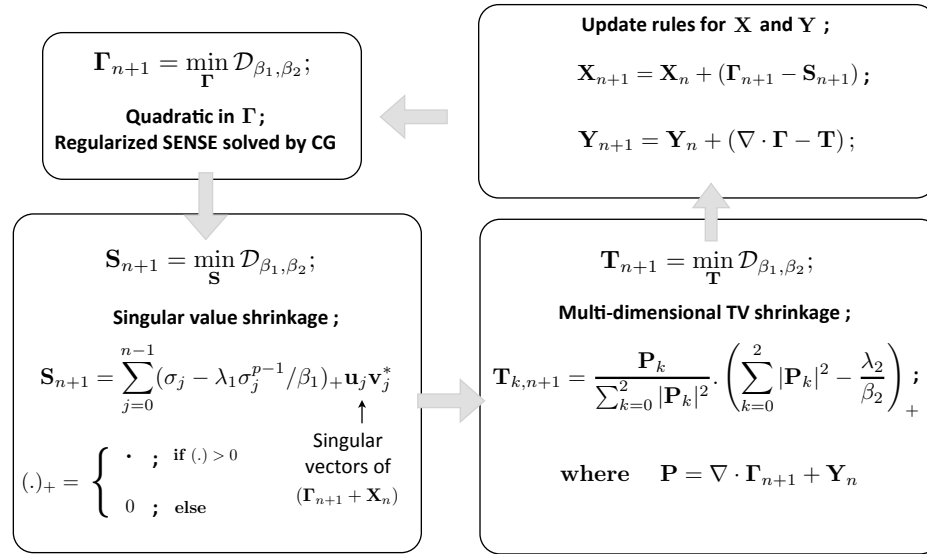


Figure 3.10: The augmented Lagrangian frame work with the different sub problems. The original problem in (3.3) is broken into a series of multiple simpler problems by using the augmented Lagrangian framework. Specifically, the algorithm iterates between the steps of regularized SENSE (that is solved by the method of conjugate gradient (CG)), singular value shrinkage, shrinkage and update rules for Lagrange multipliers. These steps are all solved by simple operations.

and stress datasets from a single subject. The parameter optimization is shown in Fig. 10, where a rest dataset is recovered using 21 radial rays/frame. The SER_{ROI} plot in Fig. 10 was evaluated for the values of λ_1, λ_2 in the window: $\underbrace{0.6554}_{256^2 \times 10^{-5}} \times [0, 6 \times 10^{-2}, 9 \times 10^{-2}, 3 \times 10^{-1}, 6 \times 10^{-1}, 9 \times 10^{-1}, 3, 6]$. The optimal values of λ_1, λ_2 were chosen such that the SER_{ROI} was the maximum. The values of λ_1, λ_2 varied slightly for the rest and stress datasets. Since all the subjects were scanned with the same protocol under shallow breathing, we used the same values of λ_1, λ_2 tuned for the first subject for reconstruction of all the other datasets. The total time spent for tuning the regularization parameters was approximately two hours.

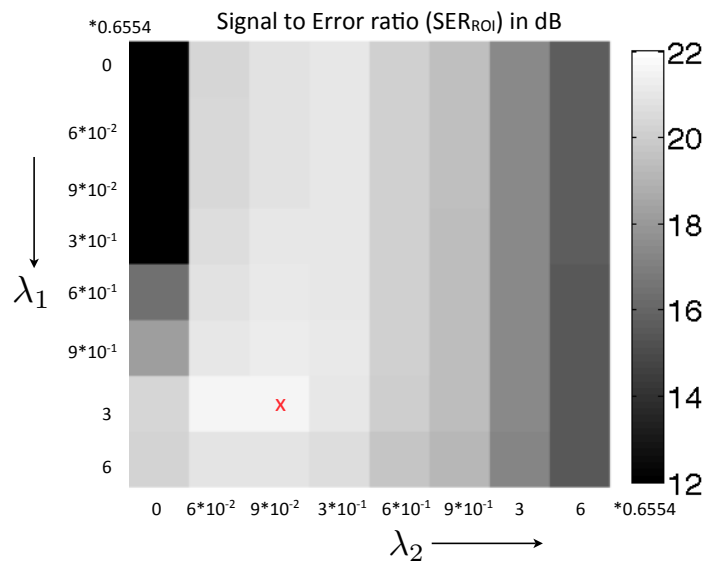


Figure 3.11: Tuning of the regularization parameters λ_1 and λ_2 . The SER was evaluated in a field of view containing the regions of the heart. The optimal parameters were chosen corresponding to the region where the SER between the reconstruction and the reference data set was maximum.

CHAPTER 4 BLIND COMPRESSIVE SENSING DYNAMIC MRI

4.1 Introduction

In the previous two chapters, we have developed improved low rank and sparsity based methods to address the limitations of modeling using predetermined bases such as Fourier bases (x-f structure/sparsity based methods). In the low rank models, since the basis functions are estimated from the data itself and no sparsity assumption is made on the coefficients, these schemes can be thought of blind linear models (BLM). These methods have been demonstrated to provide considerably improved results in perfusion [18, 93, 94] and other real time applications [95]. However, one challenge associated with this scheme is the degradation in performance in the presence of very large inter-frame motion. Specifically, large numbers of temporal basis functions are needed to accurately represent the temporal dynamics, thus restricting the possible acceleration. In such scenarios, these methods result in considerable spatio-temporal blurring at high accelerations [18, 96, 97]. The number of degrees of freedom in the low-rank representation is approximately¹ Mr , where M is the number of pixels and r is number of temporal basis functions or the rank. The dependence of the degrees of freedom on the number of temporal basis function is the main reason for the tradeoff between accuracy and achievable acceleration in applications with large motion.

In this chapter, we introduce a novel dynamic imaging scheme, termed as blind compressive sensing (BCS), to improve the recovery of dynamic imaging datasets with large inter-frame motion. Similar to classical CS schemes [15, 40, 98], the voxel intensity profiles are modeled as a sparse linear combination of basis functions in a dictionary. However, instead of assuming a fixed dictionary, the BCS scheme estimates the dictionary from the undersampled measurements itself. While this approach of estimating the coefficients and dictionary from the data is similar to BLM methods,

¹Assuming that the number of pixels is far greater than the number of frames, which is generally true in dynamic imaging applications.

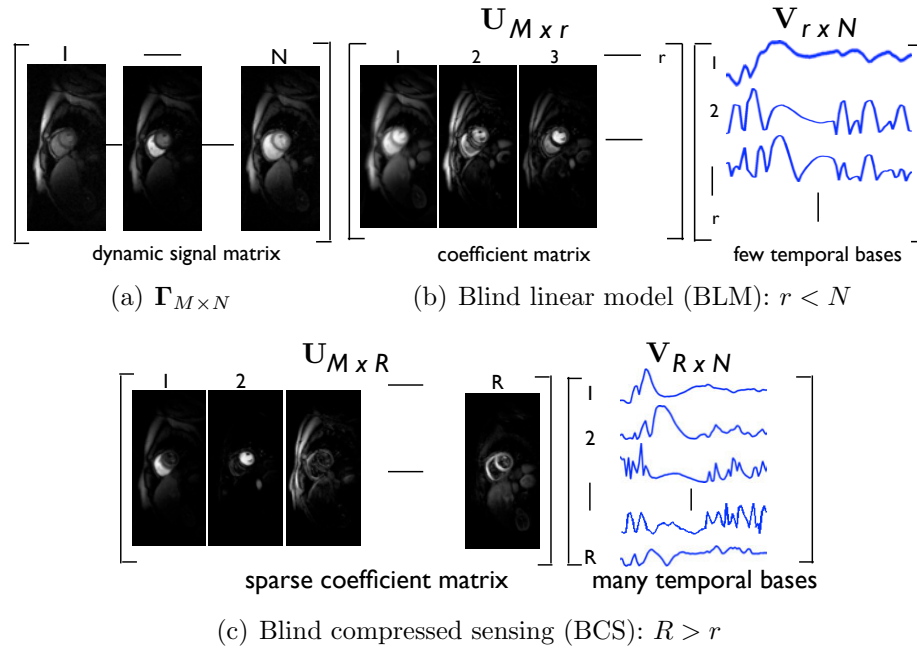


Figure 4.1: Comparison of blind compressed sensing (BCS) and blind linear model (BLM) representations of dynamic imaging data: The Casorati form of the dynamic signal Γ is shown in (a). The BLM and BCS decompositions of Γ are respectively shown in (b) and (c). BCS uses a large over-complete dictionary, unlike the orthogonal dictionary with few basis functions in BLM; ($R > r$). Note that the coefficients/ spatial weights in BCS are sparser than that of BLM. The temporal basis functions in the BCS dictionary are representative of specific regions, since they are not constrained to be orthogonal. For example, the 1st, 2nd columns of $\mathbf{U}_{M \times R}$ in BCS correspond respectively to the temporal dynamics of the right and left ventricles in this myocardial perfusion data with motion. We observe that only 4-5 coefficients per pixel are sufficient to represent the dataset.

the main difference is the sparsity assumption on the coefficients. In addition, the dictionary in BCS is much larger and the temporal basis functions are not constrained to be orthogonal (see figure 4.1). The significantly larger number of basis functions in the BCS dictionary considerably improves the approximation of the dynamic signal, especially for datasets with significant inter-frame motion. The number of degrees of freedom of the BCS scheme is $Mk + RN - 1$, where k is the average sparsity of the representation, R is the number of temporal basis functions in the dictionary, and N is the total number of time frames. However, in dynamic MRI, since $M \gg N$ the

degrees of freedom is dominated by the average sparsity k and not the dictionary size R , for reasonable dictionary sizes. In contrast to BLM, since the degrees of freedom in BCS is not heavily dependent on the number of basis functions, the representation is richer and hence provide an improved trade-off between accuracy and achievable acceleration.

An efficient computational algorithm to solve for the sparse coefficients and the dictionary is introduced in this chapter. In the BCS representation, the signal matrix $\mathbf{\Gamma}$ is modeled as the product $\mathbf{\Gamma} = \mathbf{UV}$, where \mathbf{U} is the sparse coefficient matrix \mathbf{V} is the temporal dictionary. The recovery is formulated as a constrained optimization problem, where the criterion is a linear combination of the data consistency term and a sparsity promoting ℓ_1 prior on \mathbf{U} , subject to a Frobenius norm (energy) constraint on \mathbf{V} . We solve for \mathbf{U} and \mathbf{V} using a majorize-minimize framework. Specifically, we decompose the original optimization problem into three simpler problems. An alternating minimization strategy is used, where we cycle through the minimization of three simpler problems. The comparison of the proposed algorithm with a scheme that alternates between sparse coding and dictionary estimation demonstrates the computational efficiency of the proposed framework; both methods converge to the same minimum, while the proposed scheme is approximately ten times faster. We also observe that the proposed scheme is less sensitive to initial guesses, compared to the extension of the K-SVD scheme [99] to under-sampled dynamic MRI setting. It is seen that the ℓ_1 sparsity norm and Frobenius norm dictionary constraint enables the attenuation of insignificant dictionary basis functions, compared with the ℓ_0 sparsity norm and column norm dictionary constraint used by most dictionary learning schemes. This implicit model order selection property is important in the under sampled setting since the number of basis functions that can be reliably estimated is dependent on the available data and the signal to noise ratio.

The proposed work has some similarities to [80], where a patch dictionary is

learned to exploit the correlations between image patches in a static image. The key difference is that the proposed scheme exploits the correlations between voxel time profiles in dynamic imaging rather than redundancies between image patches. The ℓ_0 norm sparsity constraints and unit column norm dictionary constraints are assumed in [80]. The adaptation of this formulation to our setting resulted in the learning of noisy basis functions at high acceleration factors. Similar to [100], the setting in [80] permits the reconstructed dataset to deviate from the sparse model. The denoising scheme is well-posed even in this relaxed setting since the authors assume overlapping patches; even if a patch does not have a sparse representation in the dictionary, the pixels in the patch are still constrained by the sparse representations of other patches containing them. Since there is no redundancy in our setting, the adaptation of the above scheme to our setting may also result in alias artifacts. Furthermore, the proposed numerical algorithm is very different from the optimization scheme in [80], where they alternate between a greedy K-SVD dictionary learning algorithm and a reconstruction update step admitting an efficient closed-form solution. We observe that the greedy approach is vulnerable to local minima in the dynamic imaging setting.

The proposed BCS setup has some key differences with the formulation in [101], where the recovery of several signals measured by the same sensing matrix is addressed; additional constraints on the dictionary were needed to ensure unique reconstruction in this setting. By contrast, we use different sensing matrices (sampling patterns) for different time frames, inspired by prior work in other dynamic MRI problems [16, 18, 40]. Our phase transition experiments show that we obtain good reconstructions without any additional constraints on the dictionary. Since the BCS scheme assumes that only very few basis functions are active at each voxel, this model can be thought of as a locally low-rank representation [96]. However, unlike [96], the BCS scheme does not estimate the basis functions for each neighborhood independently. Since it estimates \mathbf{V} from all voxels simultaneously, it is capable of exploiting

the correlations between voxels that are well separated in space (non-local correlations).

4.2 Dynamic MRI reconstruction using the BCS model

4.2.1 Dynamic image acquisition

The main goal of the chapter is to recover the dynamic dataset $\gamma(\mathbf{x}, t) : \mathbb{Z}^3 \rightarrow \mathbb{C}$ from its under-sampled Fourier measurements. We represent the dataset as the $M \times N$ Casorati matrix [74]:

$$\mathbf{\Gamma}_{M \times N} = \begin{pmatrix} \gamma(\mathbf{x}_1, t_1) & \cdot & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_1, t_N) \\ \gamma(\mathbf{x}_2, t_1) & \cdot & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_2, t_N) \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \gamma(\mathbf{x}_M, t_1) & \cdot & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_M, t_N) \end{pmatrix}. \quad (4.1)$$

Here, M is the number of voxels in the image and N is the number of image frames in the dataset. The columns of $\mathbf{\Gamma}$ correspond to the voxels of each time frame. We model the measurement process as

$$\mathbf{b}_i = \underbrace{(\mathcal{S}_i \circ \mathcal{F} \circ \mathcal{T}_i)}_{\mathcal{A}_i}(\mathbf{\Gamma}) + \mathbf{n}_i; \quad i = 1, \dots, N; \quad (4.2)$$

where, \mathbf{b}_i and \mathbf{n}_i are respectively the measurement and noise vectors at the i^{th} time instants. \mathcal{T}_i is an operator that extracts the i^{th} column of $\mathbf{\Gamma}$, which corresponds to the image at t_i . \mathcal{F} is the 2 dimensional Fourier transform and \mathcal{S}_i is the sampling operator that extracts the Fourier samples on the k-space trajectory corresponding to the i^{th} time frame. We consider different sampling trajectories for different time frames to improve the diversity.

4.2.2 The BCS representation

We model $\mathbf{\Gamma}$ as the product of a sparse coefficient matrix $\mathbf{U}_{M \times R}$ and a matrix $\mathbf{V}_{R \times N}$, which is a dictionary of temporal basis functions:

$$\underbrace{\begin{pmatrix} \gamma(\mathbf{x}_1, t_1) & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_1, t_N) \\ \gamma(\mathbf{x}_2, t_1) & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_2, t_N) \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \gamma(\mathbf{x}_M, t_1) & \cdot & \cdot & \cdot & \gamma(\mathbf{x}_M, t_N) \end{pmatrix}}_{\mathbf{\Gamma}_{M \times N}} = \underbrace{\begin{pmatrix} u_1(\mathbf{x}_1) & \cdot & \cdot & u_R(\mathbf{x}_1) \\ u_1(\mathbf{x}_2) & \cdot & \cdot & u_R(\mathbf{x}_2) \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ u_1(\mathbf{x}_M) & \cdot & \cdot & u_R(\mathbf{x}_M) \end{pmatrix}}_{\mathbf{U}_{M \times R}} \underbrace{\begin{pmatrix} v_1(t_1) & \cdot & \cdot & v_1(t_N) \\ v_2(t_1) & \cdot & \cdot & v_2(t_N) \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ v_R(t_1) & \cdot & \cdot & v_R(t_N) \end{pmatrix}}_{\mathbf{V}_{R \times N}}. \quad (4.3)$$

Here, R is the total number of basis functions in the dictionary. The model in (4.3) can also be expressed as the partially separable function (PSF) model [74, 94]:

$$\gamma(\mathbf{x}, t) = \sum_{i=1}^R u_i(\mathbf{x}) \mathbf{v}_i(t), \quad (4.4)$$

Here, $u_i(\mathbf{x})$ corresponds to the i^{th} column of \mathbf{U} and is termed as the i^{th} spatial weight. Similarly, $\mathbf{v}_i(t)$ corresponds to the i^{th} row of \mathbf{V} and is the i^{th} temporal basis function. The main difference with the traditional PSF setting is that the rows of \mathbf{U} are constrained to be sparse, which imply that there are very few non-zero entries; this also suggests that few of the temporal basis functions are sufficient to model the temporal profile at any specified voxel. The over-complete dictionary of basis functions

are estimated from the data itself and are not necessarily orthogonal. In figure 4.1, we demonstrate the differences between BLM (low-rank) and BCS representations of a cardiac perfusion MRI data set with motion. Note that the sparsity constraint encourages the formation of voxel groups that share similar temporal profiles. Since many more basis functions are present in the dictionary, the representation is richer than the BLM model. The sparsity assumption ensures that the richness of the model is not translated to increased degrees of freedom. The sparsity assumption also enables the suppression of noise and blurring artifacts, thus resulting in sharper reconstructions.

The degrees of freedom associated with BCS is approximately $Mk + RN - 1$, where k is the average sparsity of the coefficients and R is the number of basis functions in the dictionary. Since $M \gg N$, the degrees of freedom in the BCS representation is dominated by the average sparsity (k) and not the size of the dictionary (R), for realistic dictionary sizes. Since the overhead in learning the dictionary is low, it is much better to learn the dictionary from the under-sampled data rather than using a sub-optimal dictionary. Hence, we expect this scheme to provide superior results than classical compressive sensing schemes that use fixed dictionaries.

4.2.3 The objective function

We now address the recovery of the signal matrix $\mathbf{\Gamma}$, assuming the BCS model specified by (4.3). Similar to classical compressive sensing schemes, we replace the sparsity constraint by an ℓ_1 penalty. We pose the simultaneous estimation of \mathbf{U} and \mathbf{V} from the measurements as the constrained optimization problem:

$$\{\hat{\mathbf{U}}, \hat{\mathbf{V}}\} = \arg \min_{\mathbf{U}, \mathbf{V}} \left[\sum_{i=1}^N \|\mathcal{A}_i(\mathbf{UV}) - \mathbf{b}_i\|_2^2 \right] + \lambda \|\mathbf{U}\|_{\ell_1}; \quad (4.5)$$

such that $\|\mathbf{V}\|_F^2 \leq c$.

The first term in the objective function (4.5) ensures data consistency. The second

term is the sparsity promoting ℓ_1 norm on the entries of \mathbf{U} defined as the absolute sum of its matrix entries: \mathbf{U} : $\|\mathbf{U}\|_{\ell_1} = \sum_{i=1}^M \sum_{j=1}^R |u(i, j)|$. λ is the regularization parameter, and c is a constant that is specified apriori. The Frobenius norm constraint on \mathbf{V} is imposed to make the problem well posed; if this constraint is not used, the optimization scheme can end up with coefficients \mathbf{U} that are arbitrarily small in magnitude. While other constraints (e.g. unit norm constraints on rows) can also be used to make the problem well-posed, the Frobenius norm constraint along with the ℓ_1 sparsity penalty encourages a ranking of temporal basis functions. Specifically, important basis functions are assigned larger amplitudes, while un-important basis functions are allowed to decay to small amplitudes. We observe that the specific choice of c is not very important; if c is changed, the regularization parameter λ also has to be changed to yield similar results.

4.2.4 The optimization algorithm

The Lagrangian of the constrained optimization problem in (4.5) is specified by:

$$\begin{aligned} \mathcal{L}(\mathbf{U}, \mathbf{V}, \eta) = & \left[\sum_{i=1}^N \|\mathcal{A}_i(\mathbf{U}\mathbf{V}) - \mathbf{b}_i\|_2^2 \right] + \lambda \|\mathbf{U}\|_{\ell_1} + \\ & + \eta (\|\mathbf{V}\|_F^2 - c); \quad \eta \geq 0 \end{aligned} \quad (4.6)$$

where η is the Lagrange multiplier.

Since the ℓ_1 penalty on the coefficient matrix is a non differentiable function, we approximate it by the differentiable Huber induced penalty which smooths the ℓ_1 penalty.

$$\varphi_\beta(\mathbf{U}) = \sum_{i=1}^M \sum_{j=1}^R \psi_\beta(u_{i,j}), \quad (4.7)$$

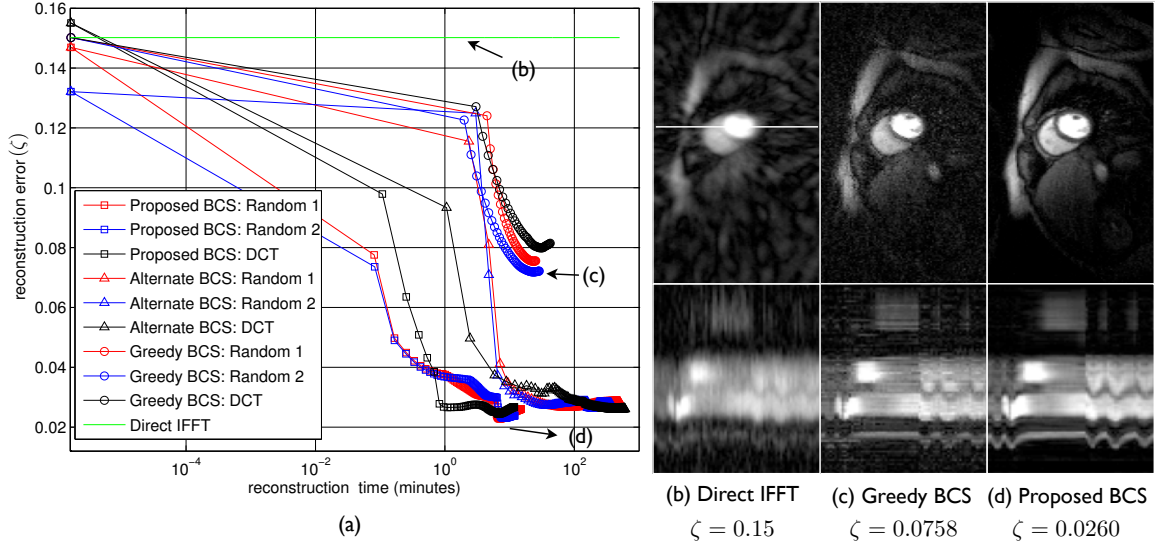


Figure 4.2: Comparison of different BCS schemes: In (a), we show the reconstruction error vs reconstruction time for the proposed BCS, alternate BCS, and the greedy BCS schemes. The free parameters of all the schemes were optimized to yield the lowest possible errors, while the dictionary sizes of all methods were fixed to 45 atoms. We plot the reconstruction error as a function of the CPU run time for the different schemes with different dictionary initializations. The proposed BCS and alternating BCS scheme converged to the same solution irrespective of the initialization. However, the proposed scheme is observed to be considerably faster; note that the alternating scheme takes around ten times more time to converge. It is also seen that the greedy BCS scheme converged to different solutions with different initializations, indicating the dependence of these schemes on local minima.

where, $u_{i,j}$ are the entries of \mathbf{U} and $\psi_\beta(u)$ is defined as:

$$\psi_\beta(x) = \begin{cases} |x| - 1/2\beta & \text{if } |x| \geq \frac{1}{\beta} \\ \beta |x|^2 / 2 & \text{else .} \end{cases} \quad (4.8)$$

The Lagrangian function obtained by replacing the ℓ_1 penalty in (4.6) by φ_β is:

$$\mathcal{D}_\beta(\mathbf{U}, \mathbf{V}, \eta) = \left[\sum_{i=1}^N \|\mathcal{A}_i(\mathbf{UV}) - \mathbf{b}_i\|_2^2 \right] + \lambda \varphi_\beta(\mathbf{U}) + \eta (\|\mathbf{V}\|_F^2 - c); \quad \eta \geq 0 \quad (4.9)$$

Note that $\varphi_\beta(\mathbf{U})$ is parametrized by the single parameter β . When $\beta \rightarrow \infty$, the Huber induced norm is equivalent to the original ℓ_1 penalty. Other ways to smooth the ℓ_1 norm have been proposed (eg: [102]).

We observe that the algorithm has slow convergence if we solve for (4.5) with $\beta \rightarrow \infty$. We use a continuation strategy to improve the convergence speed. Note that the Huber norm simplifies to the Frobenius norm when $\beta = 0$. This formulation (ignoring the constant c and optimization with respect to η) is similar to the one considered in [54]; according to the [54, Lemma 5], the solution of (4.9) is equivalent to the minimum nuclear norm solution. Here, we assume that the size of the dictionary R is greater than the rank of $\mathbf{\Gamma}$, which holds in most cases of practical interest. Thus, the problem converges to the well-defined nuclear norm solution when $\beta = 0$. Our earlier experiments show that the minimum nuclear norm solution already provides reasonable estimates with reduced aliasing [18]. Thus, the cost function is less vulnerable to local minimum when β is small. Hence, we propose a continuation strategy, where β is initialized to zero and is gradually increased to a large value. By slowly increasing β from zero, we expect to gradually truncate the small coefficients of \mathbf{U} , while re-learning the dictionary. Our experiments show that this approach considerably improves the convergence rate and avoids local minima issues.

We rely on the majorize-minimize framework to realize a fast algorithm. We start by majorizing the Huber norm in (4.9) as ² [19]:

$$\varphi_\beta(\mathbf{U}) = \min_{\mathbf{L}} \frac{\beta}{2} \|\mathbf{U} - \mathbf{L}\|_F^2 + \|\mathbf{L}\|_{\ell_1}, \quad (4.10)$$

where \mathbf{L} is an auxiliary variable. Substituting (4.10) in (4.9), we obtain the following

² Note that the right hand side of (4.10) is only guaranteed to majorize the Huber penalty $\varphi_\beta(\mathbf{U})$; it does not majorize the ℓ_1 norm of \mathbf{U} ; (as from (4.8) that $\psi_\beta(x)$ is lower than the ℓ_1 penalty by $1/2\beta$ when $|x| > 1/\beta$. Similarly $\psi_\beta < |x|/2$ when $|x| < 1/\beta$). This majorization in (4.10) later enables us to exploit simple shrinkage strategies that exist for the ℓ_1 norm; if the ℓ_1 penalty were used instead of the Huber penalty, it would have resulted in more complex expressions than in (4.16). For additional details, we refer the interested reader to [19].

modified Lagrange function, which is a function of four variables \mathbf{U} , \mathbf{V} , \mathbf{L} and η :

$$\begin{aligned} \mathcal{D}(\mathbf{U}, \mathbf{V}, \mathbf{L}, \eta) = & \sum_{i=1}^N [\|\mathcal{A}_i(\mathbf{U}\mathbf{V}) - \mathbf{b}_i\|_2^2] + \eta (\|\mathbf{V}\|_F^2 - c) + \\ & + \lambda \left[\|\mathbf{L}\|_{\ell_1} + \frac{\beta}{2} \|\mathbf{U} - \mathbf{L}\|_F^2 \right]; \end{aligned} \quad (4.11)$$

The above criterion is dependent on \mathbf{U} , \mathbf{V} , \mathbf{L} , and η and hence have to be solved for all of these variables. While this formulation may appear more complex than the original BCS scheme (4.5), this results in a simple algorithm. Specifically, we use an alternating minimization scheme to solve (4.11). At each step, we solve for a specific variable, assuming the other variables to be fixed; we systematically cycle through these subproblems until convergence. The subproblems are specified below.

$$\mathbf{L}_{n+1} = \arg \min_{\mathbf{L}} \|\mathbf{U}_n - \mathbf{L}\|_2^2 + \frac{2}{\beta} \|\mathbf{L}\|_{\ell_1}; \quad (4.12)$$

$$\begin{aligned} \mathbf{U}_{n+1} = \arg \min_{\mathbf{U}} \sum_{i=1}^N [\|\mathcal{A}_i(\mathbf{U}\mathbf{V}_n) - \mathbf{b}_i\|_2^2] + \\ + \frac{\lambda\beta}{2} \|\mathbf{U} - \mathbf{L}_{n+1}\|_2^2; \end{aligned} \quad (4.13)$$

$$\begin{aligned} \mathbf{V}_{n+1} = \arg \min_{\mathbf{V}} \sum_{i=1}^N [\|\mathcal{A}_i(\mathbf{U}_{n+1}\mathbf{V}) - \mathbf{b}_i\|_2^2] + \\ + \eta_n (\|\mathbf{V}\|_F^2 - c); \end{aligned} \quad (4.14)$$

We use a steepest ascent rule to update the Lagrange multiplier at each iteration.

$$\eta_{n+1} = (\eta_n + \|\mathbf{V}_{n+1}\|_F^2 - c)_+, \quad (4.15)$$

where ‘+’ represents the operator defined as $(\tau)_+ = \max\{0, \tau\}$, which is used to ensure the positivity constraint on η (see (4.9)).

Each of the sub-problems are relatively simple and can be solved efficiently, either using analytical schemes or simple optimization strategies. Specifically, (4.12) can be solved analytically as:

$$\mathbf{L}_{n+1} = \frac{\mathbf{U}_n}{|\mathbf{U}_n|} \left(|\mathbf{U}_n| - \frac{1}{\beta} \right)_+ ; \quad (4.16)$$

Since the problems in (4.13) and (4.14) are quadratic, we solve it using conjugate gradient (CG) algorithms.

Once $\|\mathbf{V}\|_F^2 \approx c$, we see that η stabilizes. Hence, we expect (4.14) to converge quickly. In contrast, the condition number of the \mathbf{U} sub-problem is dependent on β . Hence, the convergence of the algorithm will be slow at high values of β . In addition, the algorithm may converge to a local minimum if it is initialized directly with a large value of β . We use the above mentioned continuation approach to solve for simpler problems initially and progressively increase the complexity. Specifically, starting with an initialization of \mathbf{V} , the algorithm iterates between (4.12) and (4.15) in an inner loop, while progressively updating β starting with a small value in an outer loop. The inner loop is terminated when the cost in (4.6) stagnates. The outer loop is terminated when a large enough β is achieved. We define convergence as when the cost in (4.6) in the outer loop stagnates to a threshold of 10^{-5} . In general, with our experiments on dynamic MRI data, we observed convergence when the final value of β is approximately 10^{13} to 10^{15} times larger than the initial value of β .

4.3 Experimental Evaluation

We describe in sections (III. A-B) the algorithmic considerations of the proposed blind CS framework. We then perform phase transition experiments using numerical phantoms to empirically demonstrate the uniqueness of the blind CS framework (section III.C). We finally compare the reconstructions of blind CS against existing low

rank and compressed sensing schemes using invivo Cartesian and radial free breathing myocardial perfusion MRI datasets (section III.D).

4.3.1 Comparison of different BCS schemes

In this section, we compare the performance of the proposed scheme with two other potential BCS implementations. Specifically, we focus on the rate of convergence and the sensitivity to initial guesses of the following schemes:

- Proposed BCS: The proposed BCS formulation specified by (4.5) solved by optimizing \mathbf{U} and \mathbf{V} using the proposed majorize-minimize algorithm; the algorithm cycles through steps specified by (4.12)-(4.15).
- Alternating BCS: The proposed BCS formulation specified by (4.5) solved by alternatively optimizing for the sparse coefficients \mathbf{U} and the dictionary \mathbf{V} . Specifically, the sparse coding step (solving for \mathbf{U} , assuming a fixed \mathbf{V}) is performed using the state of the art augmented Lagrangian optimization algorithm [103]. The dictionary learning sub-problem solves for \mathbf{V} , assuming \mathbf{U} to be fixed. This is solved by iterating between a quadratic subproblem in \mathbf{V} (solved by a conjugate gradient algorithm), and a steepest ascent update rule for η (similar to (4.15)). The update of η ensures the Frobenius norm constraint on \mathbf{V} is satisfied at the end of the \mathbf{V} sub-problem. Both of the sparse coding and dictionary learning steps are iterated until convergence.
- Greedy BCS: We adapt the extension of the K-SVD scheme that was used for patch based 2D image recovery [80] to our setting of dynamic imaging. This scheme models the rows of $\mathbf{\Gamma}$ in the synthesis dictionary with temporal basis functions (as in (4.4)). Specifically, it solves the following optimization problem:

$$\{\hat{\mathbf{\Gamma}}, \hat{\mathbf{U}}, \hat{\mathbf{V}}\} = \arg \min_{\mathbf{\Gamma}, \mathbf{U}, \mathbf{V}} \|\mathbf{\Gamma} - \mathbf{UV}\|_2^2; \text{ such that}$$

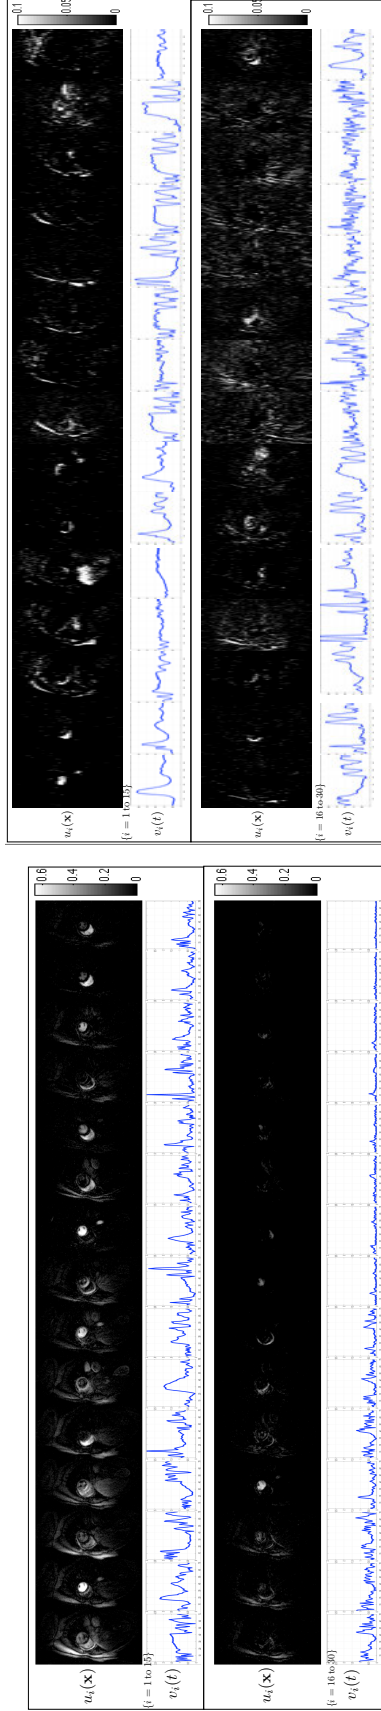
$$\begin{aligned}
\sum_{i=1}^N \|\mathcal{A}_i(\mathbf{\Gamma}) - \mathbf{b}_i\|_2^2 &< \sigma_n, \\
\|\mathbf{u}_k\|_0 &\leq j; \quad k = 1, \dots, M, \\
\|\mathbf{v}_q\|_2^2 &= 1; \quad q = 1, \dots, R;
\end{aligned} \tag{4.17}$$

where σ_n is the standard deviation of the measurement noise. Here the ℓ_0 norm is used to impose the sparsity constraints on the rows (indexed by k) of \mathbf{U} . The number of nonzero coefficients (or the sparsity level) of each row of \mathbf{U} is given by j . The unit column norm constraints are used on the elements of the dictionary to ensure well posedness (avoid scaling ambiguity). Starting with an initial estimate of the image data given by the zero filled inverse Fourier reconstruction $\mathbf{\Gamma}_{\text{init}}$, the BCS scheme in this setting iterates between a denoising/dealiasing step to update \mathbf{U} , \mathbf{V} , and an image reconstruction step to update $\mathbf{\Gamma}$. The denoising step involves dictionary learning and sparse coding with ℓ_0 minimization. It utilizes the K-SVD algorithm [99] which takes a greedy approach to update \mathbf{U} and \mathbf{V} . We implemented the K-SVD algorithm based on the codes available at the authors webpage [99]. The K-SVD implementation available online was modified to produce complex dictionaries. For sparse coding, we used the orthogonal matching pursuit algorithm (OMP). We used the approximation error threshold along with the sparsity threshold (upper bound on j) in OMP. The approximation error threshold was set to 10^{-6} . Our implementation also considered the pruning step described in [99, 100] to minimize local minima effects. Specifically, if similar basis functions were learnt, one of them was replaced with the voxel time profile that was least represented. In addition, if a basis function was not being used enough, it was replaced with the voxel time profile that was least represented. Other empirical heuristics such as varying the approximation error threshold in the OMP algorithm during the different iteration (alteration) steps may also be considered in the greedy BCS

scheme. In this work, we restrict ourselves to a fixed error threshold of 10^{-6} due to the difficulty of tuning for an optimal set of different error threshold values for different alteration steps.

In Fig. 4.2, we aim to recover a myocardial perfusion MRI dataset with considerable interframe motion ($N_x \times N_y \times N_t = 190 \times 90 \times 70$) from its undersampled $k - t$ measurements using the above three BCS schemes. We considered a noiseless simulation in this experiment for all the three BCS schemes. While resampling, we used a radial trajectory with 12 uniformly spaced rays within a frame with subsequent random rotations across frames to achieve incoherency. This corresponded to an acceleration of 7.5 fold. We used 45 basis functions in the dictionary. We compare the performance of the different BCS algorithms with different initializations of the dictionary \mathbf{V} . Specifically, we used dictionaries with random entries, and a dictionary with the discrete cosine transform (DCT) bases. To ensure fair comparisons, we optimized the parameters of all the three schemes: (i.e, regularization parameter λ in the proposed and alternating BCS schemes, as well as the sparsity level j in the greedy BCS scheme). These were chosen such that the normalized error between the reconstruction and the fully sampled data was minimal. A sparsity level of $j = 3$ was found to be optimal for the greedy BCS scheme. Further, in the greedy BCS scheme, after the first iteration, we initialized the K-SVD algorithm with the dictionary obtained from the previous iteration. We used the same stopping criterion in both the proposed and alternate BCS schemes: the iterations were terminated when the cost in (4.6) stagnated to a threshold of 10^{-5} . All the algorithms were run on a linux workstation with a 4 core Intel Xeon processor and 24 GB RAM.

From Fig. 4.2, we observe both the proposed and alternate BCS schemes to be robust to the choice of initial guess of the dictionary. They converged to almost the same solution with different initial guesses. However, the proposed BCS scheme converged to the solution significantly faster (atleast by a factor of 10 fold) compared



(a) Proposed BCS

(b) Greedy BCS

Figure 4.3: Model coefficients and dictionary bases. We show few of the estimated spatial coefficients $u_i(\mathbf{x})$ and its corresponding temporal bases $v_i(t)$ from 7.5 fold undersampled myocardial perfusion MRI data (data in Fig. 2). (a) corresponds to the estimates using the proposed BCS scheme, while (b) (in the next page) is estimated using the greedy BCS scheme. For consistent visualization, we sort the product entries $u_i(\mathbf{x})v_i(t)$ according to their ℓ_2 norm, and show the first 30 sorted terms. Note that the BCS basis functions are drastically different from exponential basis functions in the Fourier dictionary; they represent temporal characteristics specific to the dataset. It can also be seen that the energy of the basis functions in (a) varies considerably, depending on their relative importance. Since we rely on the ℓ_1 sparsity norm and Frobenius norm dictionary constraint, the representation will adjust the scaling of the dictionary basis functions $v_i(t)$ such that the $\|\mathbf{U}\|_{\ell_1}$ is minimized. Specifically, the ℓ_1 minimization optimization will ensure that basis functions used more frequently are assigned higher energies, while the less significant basis functions are assigned lower energy (see $v_{25}(t)$ to $v_{30}(t)$), hence providing an implicit model order selection. By contrast, the formulation of the greedy BCS scheme involves the setting of ℓ_0 sparsity norm and column norm dictionary constraint; the penalty is only dependent on the sparsity of \mathbf{U} . Unlike the proposed scheme, this does not provide an implicit model order selection, resulting in the preservation of noisy basis functions, whose coefficients capture the alias artifacts in the data. This explains the higher errors in the greedy BCS reconstructions in Fig. 2.

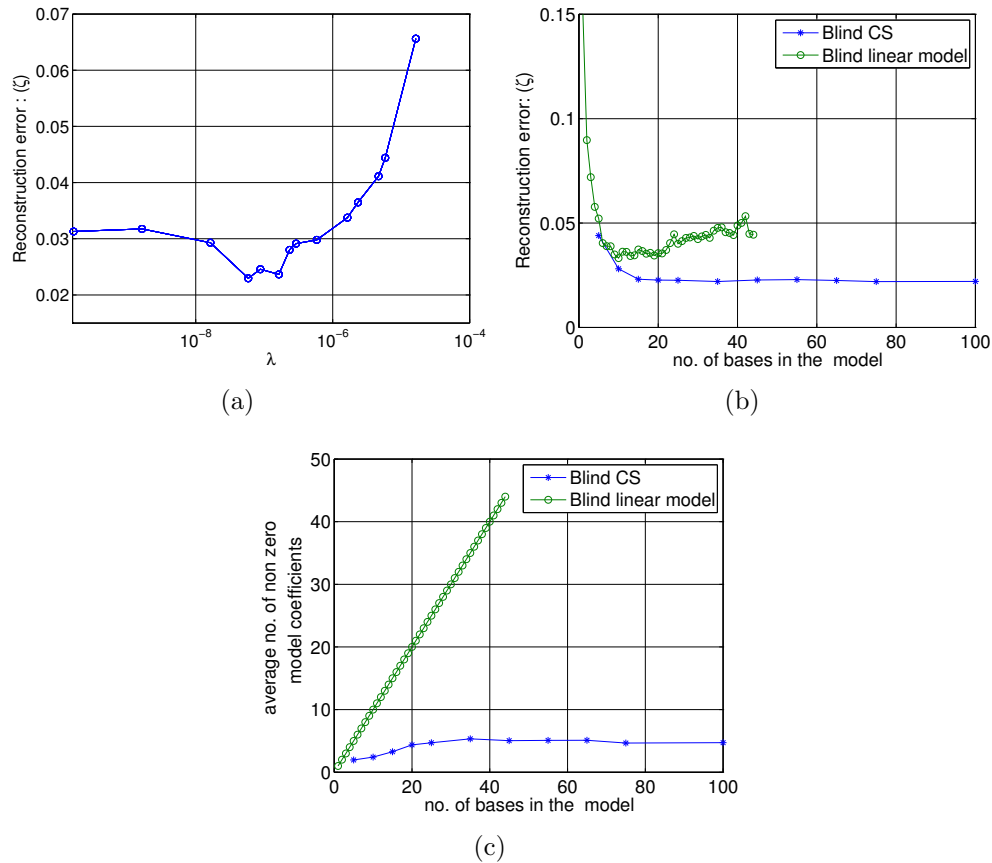


Figure 4.4: Blind CS model dependence on the regularization parameter and the dictionary size: (a) shows the reconstruction error (ζ) as a function of different λ in the BCS model. (b) and (c) respectively show the reconstruction error (ζ) and the average number of non zero model coefficients of the BCS and the BLM schemes as a function of the number of bases in the respective models. As depicted in (a), we optimize our choice of λ such that the error between the fully sampled data and the reconstruction is minimal. From (b), we observe that the BCS reconstruction error reduces with the dictionary size and hits a plateau after a size of 20 basis functions. This is in sharp contrast with the BLM scheme where the reconstructions errors increase when the basis functions are increased. The average number of BCS model coefficients unlike the BLM has a non-linear relation with the dictionary size reaching saturation to a number of 4-4.5. The plots in (b) and (c) depict that the BCS scheme is insensitive to dictionary size as long as a reasonable size (atleast 20 in this case) is chosen. We chose a dictionary size of 45 bases in the experiments considered in this chapter.

to the alternate BCS scheme. From Fig. 4.2, we observe the number of iterations for both the proposed and the alternate BCS schemes to be similar. However, since the alternate BCS scheme solves the sparse ℓ_1 minimization problem fully




Figure 4.5: The numerical phantoms Γ^j , which are used in the simulation study in figure 4.6. Here j is the number of non zero coefficients (sparsity levels) at each pixel. The top and bottom rows respectively show one spatial frame and the image time profile through the dotted white line. Note that the sparse decomposition provides considerable temporal detail even for a sparsity of one. This is possible since different temporal basis functions are active at each pixel.

during each iteration, it is more expensive than the proposed BCS scheme. On an average, an iteration of the alternate BCS scheme was ≈ 10 slower than an iteration of the proposed BCS scheme. From Fig. 4.2, we note the greedy BCS scheme to converge to different solutions for different initial guesses. Additionally, as noted in Fig. 4.2 c.d, the reconstructions with the proposed BCS scheme were better than the reconstructions with the greedy BCS scheme. Although the temporal dynamics were faithfully captured in the greedy BCS reconstructions, it suffered from noisy artifacts. This was due to modeling with noisy basis functions, which were learned by the algorithm from under sampled data (see Fig. 3). Note that this scheme uses the unit column norm constraints which has all the basis functions are ranked equally. In contrast, since the proposed scheme uses the ℓ_1 sparsity penalty and the Frobenius norm dictionary constraint, the energy of the learned bases functions varied considerably (see Fig. 4.3). With the proposed scheme, the ℓ_1 minimization optimization

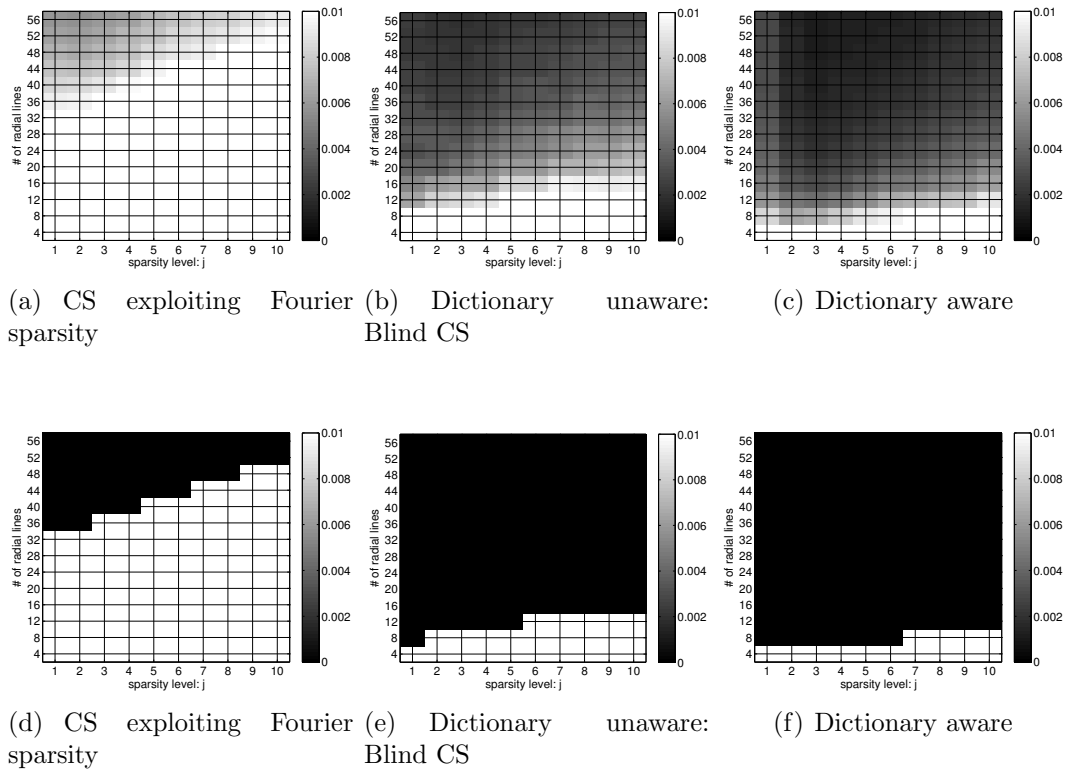


Figure 4.6: Phase transition behavior of various reconstruction schemes: Top row: Normalized reconstruction error ζ is shown at different acceleration factors (or equivalently different number of radial rays in each frame) for different values of j . Bottom row: ζ thresholded at 1 percent error; black represents 100 percent recovery. We study the ability of the algorithms to reliably recover each of the data sets Γ^j from different number of radial samples in kspace. The Γ^j , shown in Fig. 4.5 are the j sparse approximations of a myocardial perfusion MRI dataset with motion. As expected, the number of lines required to recover the dataset increases with the sparsity. The blind CS scheme outperformed the compressed sensing scheme considerably. The learned dictionary aware scheme yielded the best recovery rates. However due to a small overhead in estimating the dictionary, the dictionary unaware (blind CS) scheme was only marginally worse than the dictionary aware scheme.

ensures that the important basis functions (basis functions that are shared by several voxels) will have a higher energy. Similarly, the un-important noise-like basis functions that play active roles in fewer voxels will be attenuated, since the corresponding increase in $\|\mathbf{U}\|_{\ell_1}$ is small. Thus, the ℓ_1 penalty-Frobenius norm combination results

in a model order selection, which is more desirable than the ℓ_0 penalty-column norm combination. This choice is especially beneficial in the undersampled case since the number of basis functions that can be reliably recovered is dependent on the number of measurements and the signal to noise ratio.

4.3.2 Choice of parameters

The performance of the blind CS scheme depends on the choice of two parameters: regularization parameter λ and the number of bases in the dictionary R . Eventhough the criterion in (4.5) depends on c , varying it results in a renormalization of the dictionary elements and hence changing the value of λ . We set the value of c as 800 for both the numerical and invivo experiments. We now discuss the behavior of the blind CS model with respect to changes in λ and R .

4.3.2.1 Dependence on λ

We observe that if a low λ is used, the model coefficient matrix \mathbf{U} is less sparse. This results in representing each voxel profile using many temporal basis functions. Since the number of degrees of freedom on the scheme depends on the number of sparse coefficients, this approach often results in residual aliasing in datasets with large motion. In contrast, heavy regularization results in modeling the entire dynamic variations in the dataset using very few temporal basis functions; this often results in temporal blurring and loss of temporal detail. In the experiments in this chapter, we have access to the fully sampled ground truth data. As depicted in figure 4.4 (a), we choose the optimal λ such that the error between the reconstructions and the fully sampled ground truth data, specified by

$$\zeta = \left(\frac{\|\mathbf{\Gamma}_{\text{recon}} - \mathbf{\Gamma}_{\text{orig}}\|_F^2}{\|\mathbf{\Gamma}_{\text{orig}}\|_F^2} \right). \quad (4.18)$$

is minimized. Furthermore, in *invivo* experiments with myocardial perfusion MRI datasets, we optimize λ by evaluating the reconstruction error only in a field of view that contained regions of the heart (ζ_{ROI} , ROI: region of interest), specified by

$$\zeta_{\text{ROI}} = \left(\frac{\|\mathbf{\Gamma}_{\text{recon,ROI}} - \mathbf{\Gamma}_{\text{orig,ROI}}\|_F^2}{\|\mathbf{\Gamma}_{\text{orig,ROI}}\|_F^2} \right). \quad (4.19)$$

This metric is motivated by recent findings in [79], and by our own experience in determining a quantitative metric that best describes the accuracy in reproducing the perfusion dynamics in different regions of the heart, and the visual quality in terms of minimizing visual artifacts, and preserving crispness of borders of heart.

We realize that the above approach of choosing the regularization parameter is not feasible in practical applications, where the fully sampled reference data is not available. In these cases, one can rely on simple heuristics such as the L-curve strategy [104], or more sophisticated approaches for choosing the regularization parameters [105].

4.3.2.2 Dependence on the dictionary size

In figure 4.4.b & 4.4.c, we study the behavior of the BCS model as the number of basis functions in the model increase. We perform BCS reconstructions using dictionary sizes ranging from 5 to 100 temporal bases. The plot the reconstruction errors and the average number of non-zero model coefficients ³ as a function of the number of basis functions are shown in figures 4.4.b & 4.4.c, respectively. We observe that the BCS reconstructions are insensitive to the dictionary size beyond 20-25 basis functions. We attribute the insensitivity to number of basis functions to the combination of the ℓ_1 sparsity norm and the Frobenius norm constraint on the dictionary (see Fig. 4.3). Note that the number of basis functions that can be reliably estimated from under sampled data is limited by the number of measurements and the signal to noise

³Evaluated by performing the average of the number of non-zero coefficients in the rows of the matrix $\mathbf{U}_{M \times R}$ that was thresholded at 1 percent of the maximum value of \mathbf{U} .

ratio, unlike the classical dictionary learning setting where extensive training data is available. As discussed earlier (section III.A), the ℓ_1 sparsity norm and the Frobenius norm dictionary constraint allows the energy of the basis functions to be considerably different. Hence, the optimization scheme ranks the basis functions in terms of their energy, allowing the insignificant basis functions (which models the alias artifacts and noise) to decay to very small amplitudes. Based on these above observations, we fix the BCS dictionary size to 45 basis functions in the rest of the chapter. Note that since $45 < 70 =$ the number of time frames of the data, this is an undercomplete representation. From figure 4.4 (c), we observe that the average number of non zero model coefficients to be approximately constant ($\approx 4-4.5$) for dictionary sizes greater than 20 bases. The BCS model is also compared to the blind linear model (low-rank representation) in figures 4.4 (b & c) . The number of non zero model coefficients in the blind linear model grows linearly with the number of bases. This implies that the temporal bases modeling error artifacts and noise are also learned as the number of basis functions increase. This explains the higher reconstruction errors observed with the blind linear models as the number of basis functions increase beyond a limit.

4.3.3 Numerical simulations

To study the uniqueness of the proposed BCS formulation in (4.5), we evaluate the phase transition behavior of the algorithm on numerical phantoms. We generate dynamic phantoms with varying sparsity levels by performing dictionary learning on a fully sampled myocardial perfusion MRI dataset with motion ($N_x \times N_y \times N_t = 190 \times 90 \times 70$); i.e., $M = 17100$; $N = 70$. We use the K-SVD algorithm [99] to approximate the fully sampled Casorati matrix $\mathbf{\Gamma}_{M \times N}$ as a product of a sparse coefficient matrix $\mathbf{U}_{M \times R}^j$, and a learned dictionary $\mathbf{V}_{R \times N}^j$ by solving

$$\{\hat{\mathbf{U}}^j, \hat{\mathbf{V}}^j\} = \arg \min_{\mathbf{U}^j, \mathbf{V}^j} \|\mathbf{\Gamma} - \mathbf{U}^j \mathbf{V}^j\|_F^2$$

$$s.t. \quad \|\mathbf{u}_i\|_0 \leq j; \quad i = 1, 2, \dots, M,$$

$$\|\mathbf{v}_q\|_2 = 1; \quad q = 1, \dots, R; \quad (4.20)$$

Here, j denotes the number of non zero coefficients in each row of \mathbf{U}^j . We set the size of the dictionary as $R = 45$. We construct different dynamic phantoms corresponding to different values of j ranging from ($j = 1, 2, \dots, 10$) as $\mathbf{\Gamma}^j = \mathbf{U}^j \mathbf{V}^j$. Few of these phantoms are shown in figure 4.5. Note that the K-SVD model is somewhat inconsistent with our formulation since it relies on ℓ_0 penalty and uses the unit column norm constraint, compared to the ℓ_1 penalty and Frobenius norm constraint on the dictionary in our setting.

We perform experiments to reconstruct the spatio-temporal datasets $\mathbf{\Gamma}^j$ from $k - t$ measurements that are undersampled at different acceleration factors. Specifically, we employ a radial sampling trajectory with ' l ' number of uniformly spaced rays within a frame with subsequent random rotations across time frames; the random rotations ensure incoherent sampling. We consider different number of radial rays ranging from $l = 4, 8, 12, \dots, 56$ to simulate undersampling at different acceleration rates. The reconstructions were performed with three different schemes:

1. classical compressed sensing method, where the signal is assumed to be sparse in the temporal Fourier domain (CS) [40].
2. the proposed blind CS method, where the sparse coefficients and the dictionary are estimated from the measurements.
3. dictionary aware CS: this approach is similar to 1, except that the dictionary \mathbf{V}^j is assumed to be known. This case is included as an upper-limit for achievable acceleration.

The performance of the above schemes were compared by evaluating the normalized

reconstruction error metric ζ (4.18). All the above reconstruction schemes were optimized for their best performance by tuning the regularization parameters such that ζ was minimal.

The phase transition plots of the reconstruction schemes are shown in figure 4.6. We observe that the CS scheme using Fourier dictionary result in poor recovery rates in comparison to the other schemes. This is expected since the myocardial perfusion data is not sparse in the Fourier basis. As expected, the dictionary aware case (the exact dictionary in which the signal is sparse is pre-specified) provides the best results. However, we observe that the performance of the BCS scheme is only marginally worse than the dictionary aware scheme. As explained before, most of the degrees of freedom in the BCS representation is associated with the sparse coefficients. By contrast, the number of free parameters associated with the dictionary is comparatively far smaller since the number of voxels is far greater than the number of time frames. This clearly shows that the overhead in additionally estimating the dictionary is minimal in the dynamic imaging scenario. This property makes the proposed scheme readily applicable and very useful in dynamic imaging applications (e.g. myocardial perfusion, free breathing cine), where the signal is not sparse in pre-specified dictionaries.

4.3.4 Experiments on invivo datasets

4.3.4.1 Data acquisition and undersampling

We evaluate the performance of the BCS scheme by performing retrospective undersampling experiments on contrast enhanced dynamic MRI data. We consider one brain perfusion MRI dataset acquired using Cartesian sampling, and two free breathing myocardial perfusion MRI datasets that were acquired using Cartesian sampling, and radial sampling respectively.

The myocardial perfusion MRI datasets were obtained from subjects scanned on a Siemens 3T MRI at the University of Utah in accordance to the institute's review

board. The Cartesian dataset was acquired under rest conditions after a Gd bolus of 0.02 mmol/kg. The radial dataset was acquired under stress conditions where 0.03 mmol/kg of Gd contrast agent was injected after 3 minutes of adenosine infusion. The Cartesian dataset (phase \times frequency encodes \times time = $90 \times 190 \times 70$) was acquired using a saturation recovery FLASH sequence (3 slices, TR/TE = 2.5/1.5 ms, sat. recovery time = 100 ms). The motion in the data was due to improper gating and/or breathing; (see the ripples in the time profile in figure 4.7(c)). The radial data was acquired with a perfusion radial FLASH saturation recovery sequence (TR/TE 2.5/1.3 ms). 72 radial rays equally spaced over π radians and with 256 samples per ray were acquired for a given time frame. The rays in successive frames were rotated by a uniform angle of $\pi/288$ radians, which corresponds to a period of 4 across time. The acquired radial data corresponds to an acceleration factor of ≈ 3 compared to Nyquist. Since this dataset is slightly under sampled, we use a spatio-temporal total variation (TV) constrained reconstruction algorithm to generate the reference data in this case. We observe that this approach is capable of resolving the slight residual aliasing in the acquired data.

The single slice brain perfusion MRI dataset was obtained from a multi slice 2D dynamic contrast enhanced (DCE) patient scan at the University of Rochester. The patient had regions of tumor identified in the DCE study. The data corresponded to 60 time frames separated by TR=2sec; the matrix size was $128 \times 128 \times 60$.

Retrospective downsampling experiments were done using two different sampling schemes respectively for the Cartesian and radial acquisitions. Specifically, the Cartesian datasets were resampled using a radial trajectory with 12 uniformly spaced rays within a frame with subsequent random rotations across frames to achieve incoherency. This corresponds to a net acceleration level of 7.5 in the cardiac data, and 10.66 in the brain data. Retrospective undersampling of the cardiac radial data was done by considering 24 rays from the acquired 72 ray dataset. These rays were




Figure 4.7: Comparison of the proposed scheme with different methods on a retrospectively downsampled Cartesian myocardial perfusion data set with motion at 7.5 fold acceleration: A radial trajectory is used for downsampling. The trajectory for one frame is shown in (i). The trajectory is rotated by random shifts in each time frame. Reconstructions using different algorithms, along with the fully sampled data are shown in (i) to (v). (a-b), (c), (d-e), (f) respectively show few spatial frames, image time profile, corresponding error images, error in image time profile. The image time profile in (c) is through the dotted line in (i.b). The ripples in (i.c) correspond to the motion due to inconsistent gating and/or breathing. The location of the spatial frames along time is marked by the dotted lines in (i.c). We observe the BCS scheme to be robust to spatio-temporal blurring, compared to the low rank model; eg: see the white arrows, where the details of the papillary muscles are blurred in the Schatten p-norm reconstruction while maintained well with BCS. This is depicted in the error images as well, where BCS has diffused errors, while the low rank scheme (iii) have structured errors corresponding to the anatomy of the heart. The BCS scheme was also robust to the compromises observed with the CS scheme ; the latter was sensitive to breathing motion as depicted by the arrows in iv.




Figure 4.8: Comparisons of the different reconstructions schemes on a brain perfusion MRI dataset. The fully sampled data in (a) is retrospectively undersampled at a high acceleration of 10.66. The radial sampling mask for one frame is shown in (a), subsequent frames had the mask rotated by random angles. We show a spatial frame, the image time series, and the corresponding error images for all the reconstruction schemes. Note from (b,c), the low rank and CS schemes have artifacts in the form of spatiotemporal blur; the various fine features are blurred (see arrows). In contrast, the BCS scheme had crisper features, and superior spatiotemporal fidelity. The reconstruction error and the HFEN error numbers were also considerably less with the BCS scheme.

chosen such that they were approximately separated by the golden angle distance ($\pi/1.818$). The golden angle distribution ensured incoherent k-t sampling. The acquisition using 24 rays corresponds to an acceleration of ≈ 10.6 fold when compared to Nyquist. This acceleration can be capitalized to improve many factors in the scan (eg: increase the number of slices, improve the spatial resolution, improve quality in short duration scans such as systolic or ungated imaging).

4.3.4.2 Evaluation of blind CS against other reconstruction schemes

We compare the BCS algorithm against the following schemes:

- low rank promoting reconstruction using Schatten p -norm (Sp-N) ($p = 0.1$) minimization [18].
- compressed sensing (CS) exploiting temporal Fourier sparsity [40]

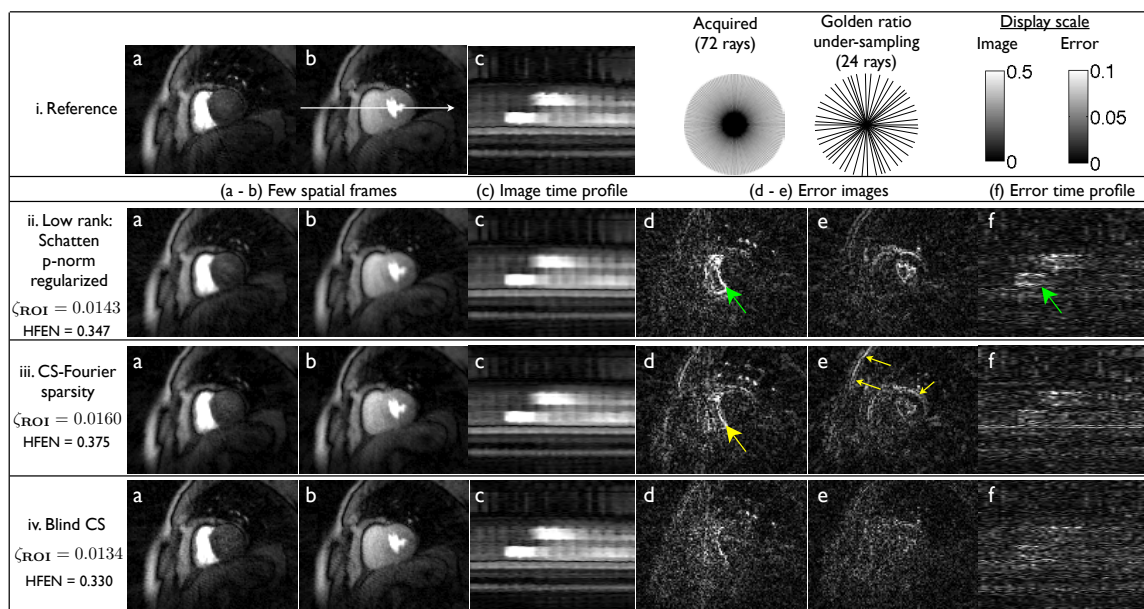


Figure 4.9: Comparisons of different reconstruction schemes on a stress myocardial perfusion MRI dataset with breathing motion: Retrospective sampling was considered by picking 24 radial rays/frame from the acquired 72 ray data; the rays closest to the golden ratio pattern was chosen. Few spatial frames, the corresponding image time profile, error frames, and error in image time profile are shown for all the schemes. We specifically observe loss of important borders and temporal blur with the low rank and CS schemes while the blind CS reconstructions have crisper borders and better temporal fidelity. Also note from the columns d,e,f that the errors in the BCS scheme are less concentrated at the edges, compared to the other methods. This indicates that the edge details and temporal dynamics are better preserved in the BCS reconstructions.

We compared different low-rank methods including two step low rank reconstruction [74], nuclear norm minimization [18], incremented rank power factorization (IRPF) [106], and observed that the Schatten p-norm minimization scheme provides comparable, or even better, results in most cases that we considered [29]. Hence we chose the Schatten p-norm reconstruction scheme in our comparisons. For a quantitative comparison amongst all the methods, we use the normalized reconstruction error metrics defined in (4.18, 4.19) and the high frequency normalized error norm metric (HFEN). The HFEN metric was used in [80] to quantify the quality of fine features, and the edges in the images, and is defined as:

$$HFEN = \frac{1}{N} \sum_{i=1}^N \frac{\|\text{LoG}(\mathbf{\Gamma}_{\text{recon},i}) - \text{LoG}(\mathbf{\Gamma}_{\text{ideal},i})\|_{\text{F}}^2}{\|\text{LoG}(\mathbf{\Gamma}_{\text{ideal},i})\|_{\text{F}}^2} \quad (4.21)$$

where LoG is a Laplacian of Gaussian filter that capture edges. We use the same filter specifications as in [80]: kernel size of 15×15 pixels, with a standard deviation of 1.5 pixels.

The comparisons on the Cartesian rest myocardial perfusion MRI dataset are shown in figure 4.7. We observe that the frames with significant motion content and contrast variations are considerably blurred with the low rank method. By contrast, the BCS scheme robustly recovers these regions with minimum spatio-temporal blur. The BCS scheme is more robust than the CS scheme. Specifically the former is robust to breathing motion, while the CS scheme results in motion blur (see arrows in figure 4.7 iv.e and iv.f).

Figure 4.8 shows the comparisons on the brain perfusion MRI dataset. We observe BCS to retain the subtle details and edges of the various structures in the brain. It shows superior spatio-temporal fidelity. In contrast, the CS and low rank schemes suffer from spatiotemporal blurring artifacts as depicted in Fig. 4.8.

In figure 4.9, we compare the various reconstruction schemes on the radial data acquired during stress conditions. We observe performance similar to figure 4.7. The low rank reconstructions exhibit reduced temporal fidelity. The reduced fidelity can result in inaccurate characterization of the contrast dynamics uptake. The CS reconstructions have considerable spatio-temporal blur. In particular, the borders of the heart and the papillary muscles are blurred with the CS scheme. By contrast, the blind CS scheme provides crisper images and are robust to spatio-temporal blur.

4.4 Towards high spatio-temporal resolution 3D DMRI of lung using Blind CS

In this section, we take a slight detour from MR-MPI, and evaluate the feasibilities of the novel blind CS scheme proposed in this chapter to enable improved dynamic MRI of the lung.

4.4.1 Preliminaries on 3D dynamic lung MRI

3-D dynamic MRI of the lung is a promising tool to assess lung function and mechanics. Compared to multi-slice 2D-DMRI, 3-D acquisitions enables the accurate estimation of lung volumes and its variations. It has recently been shown that the vital capacities estimated from 3D acquisitions are more correlated with spirometric measurements compared to 2D-DMRI [107]. However, its full potential is not clinically realized due to restricted spatio-temporal resolutions and volume coverage. To obtain whole lung coverage, a 3D FLASH scheme with Cartesian undersampling, view-sharing and parallel imaging was realized in [107] to achieve an isotropic resolution of 3.75 mm³ with 1 sec time resolution. However the reconstruction scheme was designed to image the dynamics of the lung during very slow, and controlled breathing conditions. In this section, we evaluate the feasibility of BCS in improving the tradeoffs in 3D dynamic MRI of lung. Since the bases functions in BCS are learnt from the data at hand, they are more representative of the temporal variations within

the data, and are expected to provide sparser representations than compressed sensing (CS) schemes that utilize predetermined bases. In addition, it does not require any assumptions on the breathing conditions. Additionally, we propose to combine BCS with parallel imaging and golden angle (GA) radial sampling; the combination offers superior incoherence properties

4.4.2 Retrospective undersampling of a 2D acquisition

: To test the feasibility of the BCS scheme in accelerating typical DMRI lung data, we performed retrospective undersampling experiments on a 2D free breathing dataset. Data from one coronal slice was acquired on an anesthetized swine using a TrueFISP sequence ($TR/TE = 138.62 / TE = 1.06$ msec, phase encodes: 128, Image matrix size after interpolation: 256×256 , FOV 320 mm², GRAPPA factor: 3) on a 3 T Siemens Trio with the body matrix coil array. The 28 second acquisition resulted in 200 2D images with a temporal footprint of 7images/sec. The reconstructed images were retrospectively undersampled using a golden angle radial k-t sampling pattern. Subsampling was performed by considering 40 to 10 spokes/frame. Image reconstructions with BCS, CS using Fourier bases, low rank (nuclear norm regularized), and view sharing were performed and compared.

4.4.3 Prospective 3D undersampling with stack of spokes GA radial acquisition

: A radial FLASH sequence with a 3D stack of spokes trajectory was used to image a healthy volunteer on a Siemens 3T Trio scanner with the body coil and spine coil arrays enabled. The sequence performs a radial readout in the read/phase (kx-ky) plane combined with a conventional 3D encoding step along the kz partition. The spacing between the spokes in each kx-ky plane is determined by golden angle of 111.25 degrees. 1000 radial spokes in each partition were continuously acquired for a total of 45 secs during normal breathing with (16 coronal partitions, $TR/TE = 2.84\text{ms}/1.24\text{ms}$,

FOV: 380mmx380mm, Base matrix size: 160x160, 2.37x2.37x4mm³). The data was binned at a time resolution of 0.72 seconds by considering 16 spokes/frame resulting in total of 62 frames. This corresponded to an acceleration of $R = 10$ fold, defined as ($R = \text{Base matrix size.}/\text{spokes}$). The BCS reconstruction was performed slice by slice after performing an inverse Fourier Transform along the kz dimension.

4.4.4 Results and discussion on dynamic lung MRI

From figure 4.10, we observe that the BCS scheme provided superior reconstructions compared to all the other methods. The view sharing method resulted in severe temporal blurring and artifacts due to normal breathing conditions. The CS scheme was found to be sensitive to motion artifacts, while the low rank reconstructions suffered from spatio-temporal blurring. The superior performance of BCS is attributed to the learned basis functions which capture underlying temporal dynamics (see example bases in fig.4.10). Feasible acceleration levels of upto 10 fold, (25 spokes/frame) where the mean square error was within 0.1 percent was achieved with the BCS scheme. In figure 4.11, we show every fourth slice from the 16 slice reconstructed dataset. The BCS reconstructions depicted adequate image quality in all the dynamic frames with minimal artifacts. The intensity variations of a pixel on the diaphragm depict the dynamics during breathing.

In this section, we showed preliminary feasibilities that BCS scheme utilizing learned dictionaries could significantly accelerate 3D dynamic lung MRI. During normal breathing with a 10 fold acceleration factor, spatial resolutions and time resolutions of upto (2.37mm², 0.72 sec) were achieved with a reasonable slice coverage (16 slices, 4 mm thickness). Further acceleration factors could be achieved by exploiting redundancies along the kz dimension, and undersampling the kz dimension. Future work includes systematic analysis under different breathing conditions (normal/deep/shallow) with comparisons against spirometry. Further analysis on multiple patient datasets with image quality evaluation using clinical scoring are required to

fully evaluate the reconstructions.

4.5 Discussion on Blind compressed sensing

In this chapter, we proposed a novel blind compressive sensing framework for accelerating dynamic MRI. Since the dictionary is learned from the measurements, we observe superior reconstructions compared to compressive sensing schemes that assume fixed dictionaries. Our numerical simulations and phase transition plots show that the overhead in additionally estimating the dictionary is only marginally higher than the case with known dictionary. This observation is valid in the dynamic imaging context since the number of non-zero coefficients (dependent on the number of pixels) is much higher than the size of the dictionary (dependent on the number of time frames).

We have also drawn similarities and important distinctions between the BCS scheme and blind linear models or low-rank methods. Our experiments show superior performance of the BCS scheme in comparison to the blind linear model. Specifically, better temporal fidelity, reduced spatial artifacts, sharper spatial features were distinctly observed with BCS when compared to blind linear model. These improvements can be attributed to the richness of the model in having an overcomplete set of learned temporal bases.

The proposed setting is fundamentally different from approaches that use dictionaries learnt from exemplar data and use them to recover similar images. The proposed setting learns the dictionaries jointly with the reconstruction directly from undersampled data. We observe the learnt temporal basis functions to be heavily dependent on respiration patterns, cardiac rate, timing of the bolus, gadolinium dosage, adenosine dosage, and the arterial input function (see from Fig. 3). Since these patterns would vary from subject to subject, the dictionaries learnt from the data at hand would be more beneficial in capturing subject specific patterns than dictionaries learnt from a data base of images.

The comparison of the proposed algorithm against an alternating scheme to minimize the same cost function, where the state of the art sparse coding scheme is alternated with dictionary learning, demonstrates the computational efficiency of the proposed optimization strategy. In addition, the proposed scheme is also seen to be fast and more robust to local minima than the extensions to the greedy K-SVD dictionary learning scheme. More importantly, the ability of the proposed scheme to accommodate Frobenius norm priors is seen to be advantageous in the context of dictionary learning from under sampled data; the number of basis functions that can be reliably learned is limited by the available measurements and signal to noise ratio in this setting. Specifically, this Frobenius norm constraint along with the ℓ_1 sparsity norm results in an implicit model order selection, where the insignificant basis functions are attenuated. We observe that the continuation approach in the majorize-minimize algorithm to be crucial in providing fast convergence. We plan to investigate solving the BCS problem with the augmented Lagrangian approach as proposed in [75, 108] to further improve the algorithm.

The quality of the BCS reconstructions depends on the sparsity regularization parameter λ . In general, in our experiments, the optimal value of λ (based on the metrics in (4.18) and (4.19)) did not vary much across datasets acquired with the same protocol (eg: rest cardiac perfusion MRI, shallow breathing datasets). So, in a practical setting, one could use the same λ tuned for one dataset (based on ground truth data) to recover other datasets from undersampled data that are acquired with the same protocol.

The proposed scheme can be extended in several possible directions. For example, the BCS signal representation can be further constrained by imposing the sparsity of \mathbf{U} in a fixed transform domain (e.g. wavelet, total variation domain) to further reduce the degrees of freedom. Since such priors are complementary to the redundancy between the intensity profiles of the voxels exploited by BCS, their use can

provide additional gains. This approach is similar in philosophy to [18], where we demonstrated the utility in combining low-rank models with smoothness priors. The adaptation of [109], where the authors used dictionaries with three-dimensional atoms, may be better than the 1-D dictionaries used in this work. Similarly, the use of motion compensation within the reconstruction scheme as in [32, 110] can also improve the results. The algorithm was observed to provide good performance with radial sampling trajectories. However, more work is required to evaluate the performance of the algorithm with different sampling trajectories.




Figure 4.10: Comparisons on 2D data with retrospective undersampling: BCS depicted superior quality in terms of minimal temporal and motion blurring. This is attributed to the usage of learned bases that represent the temporal dynamics (see example BCS bases in the last row that depict patterns from breathing and cardiac pulsations).

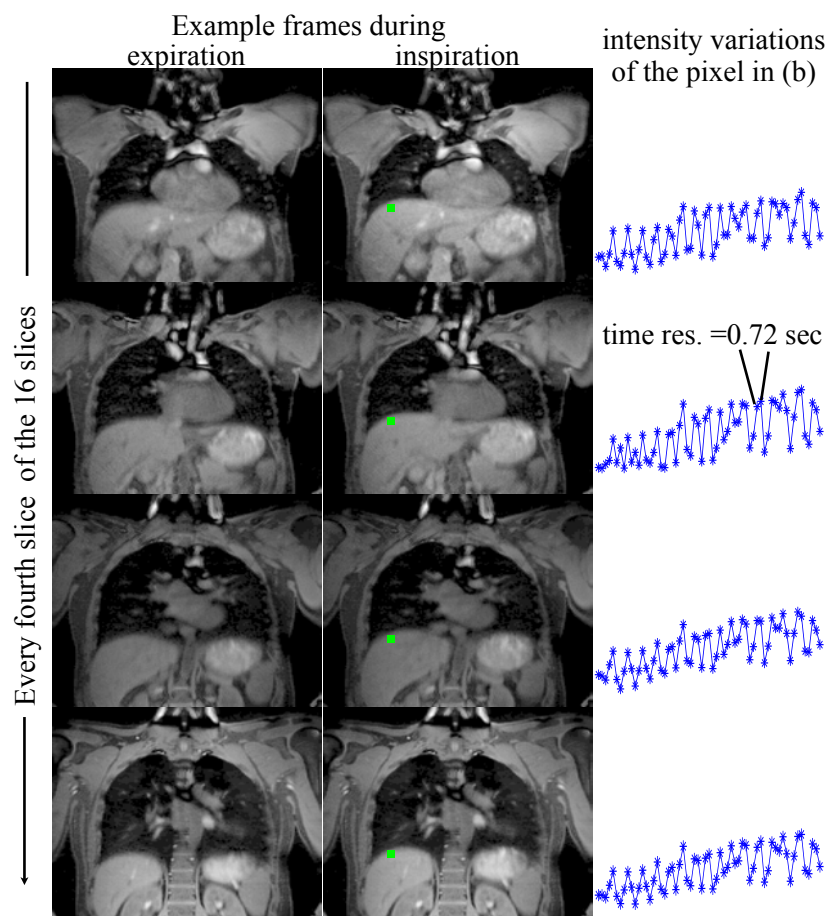


Figure 4.11: Example 3D dynamic images obtained with BCS using 16 spokes/frame with 16 slices, at spatial resolutions of $2.37 \times 2.37 \times 4 \text{ mm}^3$, and time resolution of 0.72sec.

CHAPTER 5

A NOVEL FRAMEWORK FOR MOTION COMPENSATED COMPRESSED SENSING MYOCARDIAL PERFUSION MRI

5.1 Introduction

In this chapter, we develop a novel method that improves the performance of existing compressed sensing (CS) schemes in the presence of motion. In the context of myocardial perfusion MRI, CS recovery based on exploiting sparse representations in transform domains such as temporal Fourier domain [16], temporal total variation [17], and temporal PCA domains [39] have been proposed. These methods demonstrated successful recovery when the inter-frame motion is minimal. A challenge with these methods is the sensitivity to large inter-frame motion. The motion decreases the sparsity of the representation as a result of which the reconstructions suffer from severe temporal blurring and motion related artifacts at high acceleration factors.

To overcome the challenges associated with motion artifacts, recently researchers have proposed schemes to estimate the motion and compensate for it during the reconstruction. Pederson et. al proposed to unify the reconstruction of the images and the motion compensation into a single algorithm [111]. They represented the contrast variations using a parametric perfusion model, while motion was modeled as a modulation of a 2-D displacement field, which is estimated from two images acquired at end inspiration and end expiration. The fewer degrees of freedom in this model may be restrictive in practical perfusion imaging applications. Jung et. al, have extended their $k-t$ FOCUSS scheme with motion estimation and compensation for cardiac cine MRI [110]. This scheme approximate the dynamic images as the deformation of fully sampled reference frames, collected before and after the dynamic acquisition. The residuals are then reconstructed from under-sampled k -space data using $k-t$ FOCUSS. Unlike cine MRI, the contrast of the dynamic images are significantly different from the reference images. Hence, the subtraction of the deformed

reference image may not generate sparse residuals. Moreover, more complex mutual information similarity measures may be needed for the registration. Similar to [110], Usman et.al in [112] developed a motion estimation scheme customized for cardiac cine imaging. Otazo et al. in [113] partially corrected for the motion in myocardial perfusion MRI using a rigid deformation model, where all the frames from a preliminary CS reconstruction were mapped to a single fully sampled reference image to estimate the motion. However, registering image frames to a single reference image may be suboptimal as image contrast varies significantly across time- frames [114]. Fessler recently introduced an elegant energy minimization framework to reconstruct a static image of a moving organ from its measurements [115]. The formulation of the problem as a unified energy minimization scheme enables the appreciation of the tradeoffs in the modeling. However, this scheme is not designed to recover image time series with dynamic contrast variations.

In this chapter, we propose a novel framework to jointly estimate motion and dynamic images from undersampled data. We model the motion as an elastic deformation, whose parameters are also estimated from the data. We assume the myocardial perfusion dynamics to be sparse in transform domains, once the motion is removed. This model is considerably less constrained than the parametric scheme used in [111]. We introduce an efficient variable splitting framework with continuation to decouple the problem into simpler sub-problems. The novelties enabled by this optimization are (a) a generalized formulation capable of handling any temporal sparsifying transform (such as temporal Fourier, temporal gradient, temporal PCA), (b) derivation of a reference dataset that is free of motion from the measurements themselves (c) efficient decoupling of the motion estimation problem from the reconstruction problem. Unlike existing MC-CS schemes, the proposed scheme does not require fully sampled prescans or navigators for motion estimation. Since we do not model the dynamic frames as deformations of pre-contrast reference images, our approach is robust to

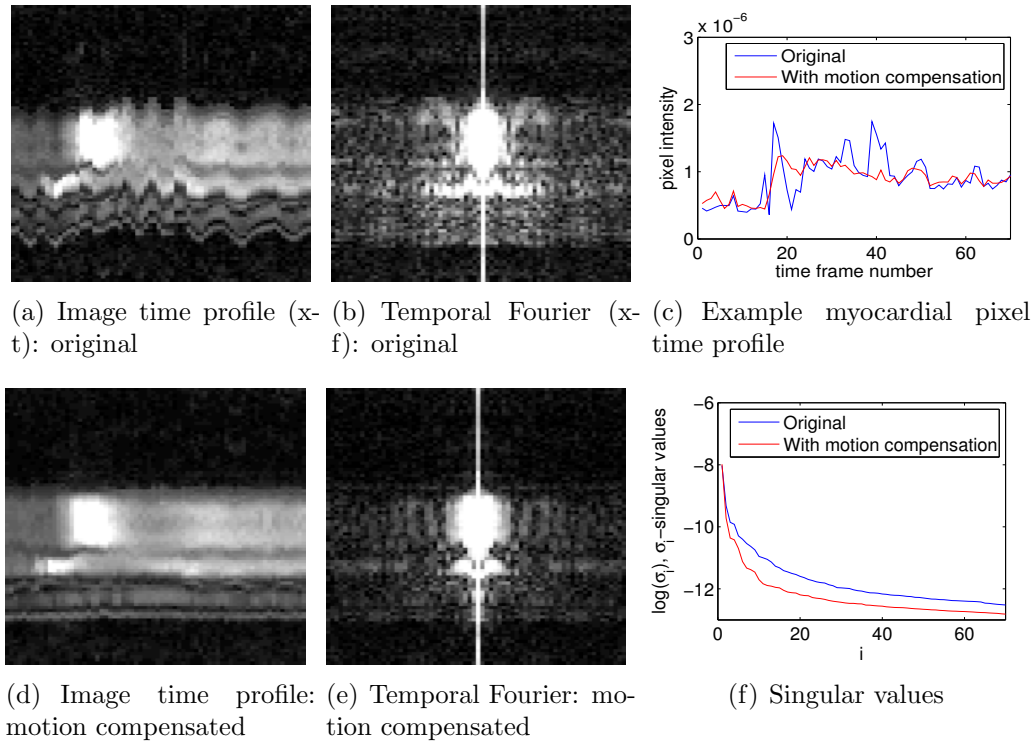


Figure 5.1: Free breathing myocardial perfusion MRI data representation in transform domains with and without motion compensation: We show the image time profile of a free breathing MPI dataset in (a); the ripples here correspond to interframe motion largely due to breathing. The corresponding profile of the motion compensated (MC) dataset in (d) show that profile is largely free of the ripples. From (b) and (e), it can be seen that the motion compensated dataset takes a sparser representation compared to the original data as the temporal harmonics corresponding to the motion are compensated in (e). From (c), it can be seen that the pixel time profiles are more piece-wise smooth compared with the MC data compared to the original, hence the former has more sparse temporal gradients. In (f), it can be seen that the number of significant singular values are reduced in the MC dataset compared to the original.

contrast variations due to bolus passage, in comparison to [110]. Our motion estimation scheme estimates the deformation by registering the dynamic data to a reference dataset that is free of respiratory motion, which is derived from the measurements themselves in the variable splitting optimization.

In the next two sections, we describe the formulation of the MC-CS problem, the proposed variable splitting framework and the resulting optimization algorithm. In

the last two sections of this chapter, we present results on the feasibility of the MC-CS algorithm to improve reconstruction quality of free breathing myocardial perfusion MRI data based on retrospective resampling experiments on fully sampled numerical phantom and in-vivo datasets. We also demonstrate the feasibility in improving free breathing reconstructions of prospectively accelerated radial data from subjects imaged during adenosine stress perfusion.

5.2 Motion compensated compressed sensing (MC-CS)

5.2.1 Dynamic image acquisition

The main objective in this chapter is to recover the dynamic dataset $f(\mathbf{x}, t); \mathbf{x} \in (x, y)$ from its undersampled Fourier noisy measurements $b(\mathbf{k}_i, t_i)$. The measurement process in dynamic MRI can be modeled as:

$$\mathbf{b}_i = \int_{\mathbf{x}} f(\mathbf{x}, t_i) \exp(-j\mathbf{k}_i^T \mathbf{x}) d\mathbf{x} + \mathbf{n}_i; \quad i = 0, \dots, s-1. \quad (5.1)$$

Here, (\mathbf{k}_i, t_i) indicates the i^{th} sampling location. We denote the set of sampling locations as $\Xi = \{(\mathbf{k}_i, t_i), i = 0, \dots, s-1\}$. The above expression can be rewritten in the vector form as

$$\mathbf{b} = \mathcal{A}(f) + \mathbf{n}, \quad (5.2)$$

where \mathcal{A} is an operator that evaluates the Fourier Transform on the sampling locations specified in Ξ .

5.2.2 MC-CS model

In a setting of the subject maintaining a perfect breath-hold, the voxel time profiles in myocardial perfusion MRI follow certain characteristics that make it amenable

to compressed sensing. As demonstrated in figure 5.1, the data takes a sparse representation in transform domains such as temporal Fourier transform (due to low temporal bandwidth), temporal total variation (due to reasonably smooth voxel time profiles), temporal principal component analysis (PCA) (as many voxels share similar time profiles). However, the sparse representations in these transform domains are significantly disturbed in the presence of significant inter-frame motion, that could be caused by respiratory motion or inconsistent gating (see fig. 5.1). In this context, we propose to simultaneously recover the motion characterized by $\theta(\mathbf{x}, t)$ and the dynamic images $f(\mathbf{x}, t)$ from under-sampled data $b(\mathbf{k}, t)$ using the following minimization scheme:

$$\{f^*, \theta^*\} = \min_{f, \theta} \|\mathcal{A}(f) - \mathbf{b}\|_2^2 + \lambda \|\Phi(\mathcal{T}_\theta \cdot f)\|_{\ell_1}; \quad (5.3)$$

Here \mathcal{A} is the Fourier sampling operator as defined in (5.2). $\theta(\mathbf{x}, t)$ are the motion parameters that describe pixel wise displacements, and \mathcal{T} is an image warping operator based on bilinear interpolation. $(\mathcal{T}_{\theta(\mathbf{x}, t)} \cdot f)$ is the motion compensated version of f . In this work, we consider a non-rigid deformation model to model the motion. Φ can be any temporal sparsifying transform such as the temporal Fourier transform, temporal finite difference transform, temporal PCA transform. Note that the regularization term promotes sparsity of the motion compensated dataset $(\mathcal{T}_\theta \cdot f)$ rather than f . The rapid variations in f , induced by inter-frame motion will be captured by $\theta(\mathbf{x}, t)$, which is also estimated during the joint estimation scheme from the undersampled measurements.

5.3 MC-CS: Optimization algorithm

We use the approach of variable splitting to decouple the original problem in (5.3) to simpler sub-problems. Specifically, we split the deformation from the ℓ_1 norm by introducing an auxiliary variable g . This enables us to reformulate the unconstrained

problem in (5.3) to a constrained one below (5.4):

$$\begin{aligned} \min_{f, \theta, g} \|\mathcal{A}(f) - \mathbf{b}\|_2^2 + \lambda \|\Phi(g)\|_{\ell_1}; \\ \text{s.t.}, \mathcal{T}_\theta \cdot f = g; \end{aligned} \quad (5.4)$$

We solve (5.4) using the penalty based method, where the constraint is relaxed and the corresponding quadratic violation is penalized as:

$$\min_{f, \theta, g} \|\mathcal{A}(f) - \mathbf{b}\|_2^2 + \lambda \left(\|\Phi(g)\|_{\ell_1} + \frac{\beta}{2} \|\mathcal{T}_\theta \cdot f - g\|_2^2 \right); \quad (5.5)$$

where β is the penalty parameter that determines the equivalence of (5.5) to (5.4). When β approaches ∞ , the solution of (5.5) tends to that of (5.4) and hence (5.3). The cost in (5.4) has to be now minimized with respect to three variables f, θ, g ; we solve it by using an alternating minimization scheme, where each variable is solved by assuming the rest to be known. This results in the following sub problems:

5.3.1 g sub-problem (Temporal denoising/dealiasing)

With f and θ fixed, the minimization of (5.5) with respect to g is a denoising/dealiasing problem; Specifically, it involves denoising/dealiasing of the motion compensated dataset $\underbrace{\mathcal{T}_\theta \cdot f}_q$ by promoting ℓ_1 sparsity in the transform domain Φ :

$$\min_g \frac{2}{\beta} (\|\Phi(g)\|_{\ell_1}) + \|\underbrace{\mathcal{T}_\theta \cdot f}_q - g\|_2^2; \quad (5.6)$$

Note that, for a low value of β , the solution in g will be biased towards the sparsity based regularization term and vice-versa. To solve (5.6), we employ a fast penalty based optimization algorithm with continuation. This algorithm involves alteration of a shrinkage step and an analytical update of g . The shrinkage steps and the analytical

updates can be derived accordingly to a specific chosen transform Φ (such as temporal Fourier, temporal total variation, or temporal PCA).

5.3.2 θ sub-problem (Motion estimation)

Assuming the variables f and g in (5.5) to be known, we solve for the motion parameters as:

$$\min_{\theta} \|\mathcal{T}_{\theta} \cdot f - g\|_2^2; \quad (5.7)$$

This is a registration problem with a least squares similarity metric, where the dynamic scene $f(\mathbf{x}, t)$ is registered frame by frame with a reference scene $g(\mathbf{x}, t)$. Note that the reference series is derived from the measurements itself (obtained from (5.6)); we do not require additional high resolution reference frames. The temporal profiles of the reference dataset g is significantly more smooth compared to f . This approach enables us to decouple the effects of smooth perfusion induced contrast changes and the more rapid changes resulting from respiratory motion.

To solve the above registration subproblem, one could employ any standard image registration algorithm. In this work, we used the optical flow based demons non-rigid registration algorithm [116, 117], which is available as an open source package [118]. In the demons algorithm, a displacement vector is assigned to each pixel, and the deformation is estimated such that the displacement field that gives for each pixel on the reference image its corresponding location on the target image is smooth. It uses a gaussian based regularizer to penalize irregular deformations, and has information from the gradient of the target and reference images as its driving force. To summarize, the demons algorithm alternate between the following steps in an iterative mode:

Summary of the demons registration algorithm

Input: $f(\mathbf{x}, t_i)$ is the target image, and $g(\mathbf{x}, t_i)$ is the reference image for a specified

time instant t_i ; $\mathbf{x} \in (x, y)$

Initialization: $\theta(\mathbf{x}, t_i) = \theta_{\text{init}}(\mathbf{x}, t_i)$;

while (iterations < Maximum iteration number) ;

$m(\mathbf{x}, t_i) = \mathcal{T}_{\theta(\mathbf{x}, t_i)} \cdot f(\mathbf{x}, t_i)$; (a) Warp the target image using the linear interpolation warping operator \mathcal{T} .

$u(\mathbf{x}, t_i) = \frac{(\mathcal{T}_{\theta} f - g) \nabla(g)}{|\nabla g|^2 + \alpha^2 (\mathcal{T}_{\theta} f - g)^2} + \frac{(\mathcal{T}_{\theta} f - g) \nabla(\mathcal{T}_{\theta} f)}{|\nabla(\mathcal{T}_{\theta} f)|^2 + \alpha^2 (\mathcal{T}_{\theta} f - g)^2}$; (b) Update the displacement based on demon forces from the gradients of the reference and target images.

$u = u * \mathcal{G}_{\sigma}(\mathbf{x})$; (c) Smooth the displacement field by an isotropic Gaussian filter (\mathcal{G}) having a zero mean, and a standard deviation of σ pixels.

$\theta = \theta + u$; (d) Update the transformation field

end

The displacement field in (a) has been adapted from [117], where the demon force comprise of (i) internal edge based forces from the static and reference images (i.e., the image gradients on the reference, target images that gives the relationship between the neighboring points), and the (ii) external force governed by the difference in the pixel intensities of the reference and target images. The algorithm has two free parameters that affect the performance of registration: (a) the value of σ that controls the width of the Gaussian regularizer which in turn controls the regularization of the deformation field, and (b) α which governs the force strength. Small values of α would correct for large deformations, while larger values of α corrects for finer deformations [117].

5.3.3 f sub-problem (reconstruction update)

Assuming θ , and g fixed in (5.5), the minimization with respect to f reduces to:

$$\min_f \|\mathcal{A}(f) - \mathbf{b}\|_2^2 + \frac{\lambda\beta}{2} \|\mathcal{T}_\theta \cdot f - g\|_2^2; \quad (5.8)$$

This is a quadratic problem which we solve by using the method of conjugate gradients (CG).

5.3.4 Continuation strategy to alternate between the sub-problems

The simultaneous estimation of f, θ in (5.3) makes the problem non-convex. The usage of the variable splitting strategy described above enabled the decomposition of the original problem into simpler subproblems in (5.6), (5.7), (5.8). To ensure efficient cycling between these sub problems, we utilize the following continuation strategies:

5.3.4.1 Continuation over the penalty parameter (β)

The cost in (5.5) should be solved for a very large value of β to ensure that the constraint in (5.4) is satisfied. At a high value of β , the problem gets highly ill-conditioned, and the resulting algorithm will have poor convergence properties. On the other hand a low value of β will ensure fast convergence but with the accuracy compromised. We incorporate a continuation strategy over β that has been successfully used in related works in the context of compressed sensing [53], [62]. Specifically, we start with a low values of β to solve the modified cost in (5.5) in an inner loop. Once a stopping criterion is met in the inner loop, we increment the value of β in an outer loop and again solve (5.5). This is repeated until the constraint in (5.4) is satisfied. In addition to improving the convergence of the algorithm, the continuation over β enables the derivation of the reference scene g , which is free of motion, and is subsequently used as the reference scene in the image registration subproblem. Specifically, during the initial iterations, with a low value β , the solution to the g

subproblem (temporal denoising/dealiasing problem in (5.6)) would be heavily regularized, which results in smoothing of the rapidly varying inter-frame oscillations while preserving the slowly varying dynamics due to the contrast changes in f . As β is increased, the solution in g subproblem gets closer to the data-term in (5.6) i.e, g tends towards $\mathcal{T}_\theta \cdot f$, and as described above, the algorithm is terminated when this constraint is satisfied.

5.3.4.2 Continuation over the deformation force strength parameter (α)

The motion parameter estimation problem in (5.7) is itself a non-convex problem. Additionally when iterated along with the reconstruction (5.8) and denoising (5.6) problems, there is a possibility that the estimated motion parameters could get stuck in bad local minima resulting in the algorithm to diverge. To avoid such scenarios, we propose to use a strategy of a coarse to fine motion correction strategy by adaptively tuning the demons registration algorithm parameters during the iterations. To this end, we adapt the recommendation prescribed in [117] by correcting for bulk motion using a small value of the force strength parameter α , and as the iterations proceed, we gradually increase α , and correct for finer motion changes. Specifically, with an initial guess of $\theta = 0$, the sub-problem in (5.7) is solved in an outer loop starting with a small value of α , and the solution in θ is refined by gradually incrementing α towards high values; After the first outer loop, the sub-problem in (5.7) is initialized with the motion estimates obtained from the previous iteration. The following pseudo code summarizes the continuation strategies that we adapt to solve the cost in (5.5):

Initialization: $f = f_{init}, \theta = 0, \alpha > 0, \beta > 0;$

for out = 1 to $\text{Max}_{\text{outer}}$ iterations

for in = 1 to $\text{Max}_{\text{inner}}$ iterations

```

 $f_n \leftarrow (5.8)$ ; CG reconstruction update with motion compensation;
 $g_n \leftarrow (5.6)$ ; temporal denoising/de-aliasing;
    if ( $|\text{cost}_n - \text{cost}_{n-1}|/|\text{cost}_n| < 10^{-3}$ ); cost as defined in (5.3);
        break; break the inner loop;
    end
end
while ( $\|\theta_n - \theta_{n-1}\|_2^2 / \|\theta_n\|_2^2 > 10^{-2}$ ); Check if the motion parameter update
has converged;
     $\theta_n \leftarrow (5.7)$ ; Motion estimation (Non-rigid image registration using the
demons algorithm);
    end
 $\alpha = \alpha \times 2$ ; continuation over the deformation force strength parameter;
 $\beta = \beta \times 10$ ; continuation over the penalty parameter;
end

```

5.3.5 Convergence analysis

In this section, we study the convergence behavior of the proposed MC-CS algorithm. We consider a myocardial perfusion MRI dataset with significant motion content, and aim to recover it from its under-sampled Fourier measurements by using the MC-CS algorithm (see section IV.A for the data specifications). For under-sampling, we consider resampling of the dynamic k-t data using a golden-angle radial k-t trajectory where the successive radial rays were spaced by the golden angle (111.25 degrees); here, we considered 20 radial rays that corresponded to an acceleration factor of 4.5. In this section, we employ Φ to be the temporal gradient (x-TV) operator. We now demonstrate the role of the continuation, and discuss the algorithm's dependence on

the initial guess of the reconstruction.

5.3.5.1 Continuation over β

In figure 5.2, we demonstrate the role of incrementing the parameter β gradually. From the figure, it can be seen that the gradual increment ensures the derivation of the reference dataset free of motion, which is used as the reference dataset for the image registration sub-problem. At convergence it can be seen that the constraint in eq. 5.4 is met.

5.3.5.2 Continuation over α

In figure 5.3 , we show the cost in (5.3) for the cases of a fixed value of $\alpha = 4$, and a continuation over α . In the latter, α is initialized to 4, and is incremented by a factor of 4 as described in the pseudo code above. From the figure, it can be seen that the algorithm converge to almost the same solution in both these cases. However, the method with continuation showed a faster convergence. Specifically the one with continuation had a reconstruction time of 30 minutes while the regime with fixed α took about 53 minutes to reconstruct.

5.3.5.3 Sensitivity to initial guess

In figure 5.4, we study the behavior of the algorithm to different initial guesses of f . Specifically, we study the evolution of the SER with the algorithm using initial guesses obtained from the fully-sampled ground truth data, zero filled direct IFFT reconstructed data, and a spatially regularized total variation reconstructed data. We found that the algorithm was more or less robust to the choice of the initialization. This could be attributed to the continuation over the penalty and deformation parameters. The continuation ensures a gradual progression in the complexity of the problem, thereby avoiding the chance of getting stuck in bad local minima. We however did not evaluate the proposed method with different choices of sampling

patterns. For example, a 1D Cartesian pattern could result in significant overlap artifacts, which may not serve as a good initialization. The radial sampling pattern on the other hand gave benign artifacts, and was still robustly used as an initialization to the algorithm. A natural way to get the algorithm working with any sampling pattern is to initialize it with a basic regularized reconstruction such as the spatial TV reconstruction.

5.4 Materials and Methods

5.4.1 Datasets and $k-t$ sampling

To validate the proposed MC-CS scheme, we initially perform retrospective re-sampling experiments on (i) the Physiologically improved non-uniform cardiac torso (PINCAT) numerical phantom, [119], [18] and (ii) an in-vivo fully sampled myocardial perfusion MRI dataset with motion. We finally show example reconstructions using prospectively undersampled radial data from two subjects imaged during stress and free breathing conditions.

We consider a single slice in the short axis view of the PINCAT phantom whose parameters were set to obtain realistic cardiac perfusion dynamics and contrast variations due to bolus passage, while accounting for respiration with variability in breathing motion. The contrast variations due to bolus passage are realistically modeled in regions of the right ventricle (RV), left ventricle (LV) and the left ventricle myocardium. Here, we assume a temporal resolution of one heart-beat, and the dynamic frames to be acquired during the diastolic phase (where the cardiac motion is minimal). The time series data consists of 35 time frames capturing the first pass passage of bolus through the different regions of the heart. We observe that the predominant motion (due to respiration) is in the superior-inferior direction with a low degree of through plane motion in the anterior-posterior direction. The spatial matrix size is 64×64 , which corresponds to a spatial resolution of $3 \times 3 \text{ mm}^2$.

For the fully sampled in vivo data, we considered a single slice from a satura-

tion recovery FLASH sequence (TR/TE=2.5/1ms, Saturation recovery time =100ms, 3slices). Data was acquired on a Cartesian grid (PE x FE encodes: 90x190, temporal resolution: 1 beat) using a Gadolinium bolus of 0.02mmol/kg under rest conditions. The data contained motion primarily due to breathing and inconsistent gating. Additionally, some integer shifts were added to amplify the motion (see fig 5.7).

For the PINCAT and the fully sampled in-vivo data, resampling experiments were performed by using a radial k-t sampling pattern based on the golden angle ray distribution; i.e, the angle between successive rays was 111.25 degrees. Subsampling was performed by considering (30 to 8) and (30 to 12) rays/frame respectively for the PINCAT and the in-vivo datasets.

For the prospectively undersampled data, we considered data acquired using a perfusion radial FLASH saturation recovery sequence (TR/TE \approx 2.6/1.2 ms, 3 slices per beat, flip angle of 14 degrees, $2.3 \times 2.3 \times 8$ mm voxel size, FOV: 280 mm², Bandwidth 1002 Hz/pixel) on a Siemens 3T Trio scanner [77]. 72 radial rays equally spaced over π radians and with 256 samples per ray were acquired for a given time frame and a given slice. These rays were acquired in an interleaved manner in subsets of 6 rays each. The rays in successive frames were rotated by a uniform angle of $\pi/288$ radians, which corresponded to a period of 4 across time. Data was acquired with the Siemens cardiac coil array with four channels. We considered two stress data sets that were acquired on free breathing subjects where 0.03 mmol/kg of Gd contrast agent was injected after 3 minutes of adenosine infusion. In this work, a PCA based coil compression strategy [85] was used to compress the four coil data set to a single coil principal component data set. With this data, we performed single coil reconstruction comparisons using 21 rays that were chosen to approximately follow the golden angle distribution.

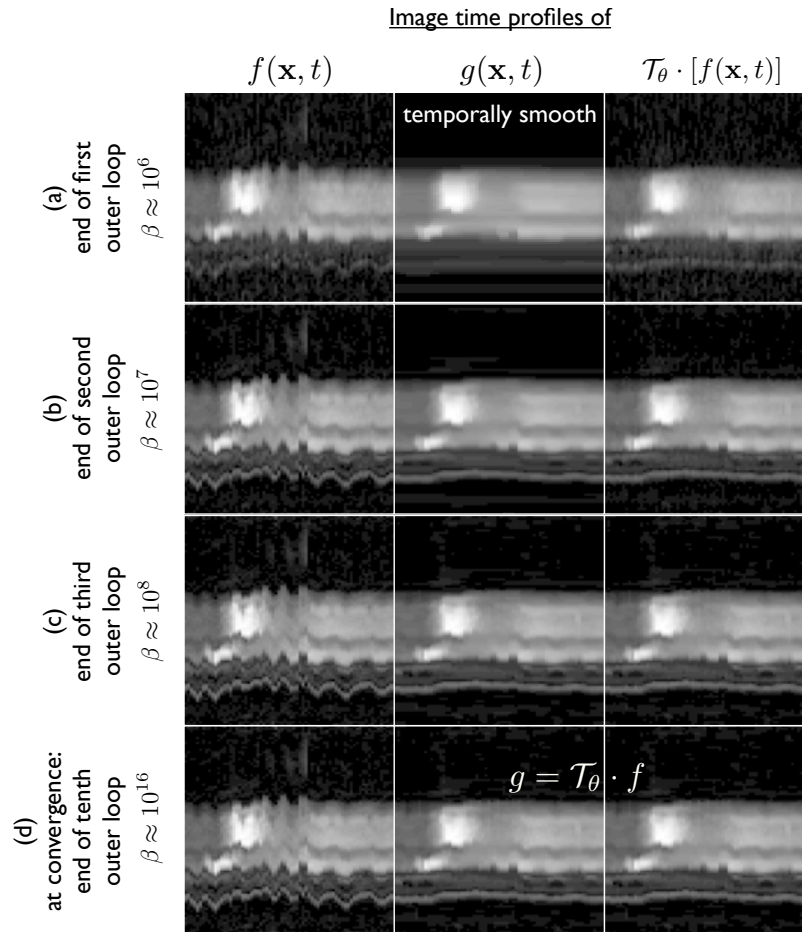


Figure 5.2: Continuation over the penalty parameter β : Here we show the time profiles of $f, g, \mathcal{T}_\theta \cdot f$ as they evolve during the iterations. From the first row, it can be seen that in the initial iterations, the value of β is small that resulted in the temporally smooth reference dataset in g . This dataset is image resgistered frame by frame to the reconstruction f to obtain the motion estimates θ . As the algorithm converge, it can be seen the reconstruction and the motion estimates improve, and at convergence the constraint in eq. (5.4) ($g = \mathcal{T}_\theta \cdot f$) is met.




Figure 5.3: Continuation over α : It can be seen that the continuation over the deformable force strength parameter α results in a faster convergence, and also avoids the cost to diverge in the last iterations.




Figure 5.4: Evolution of the region of interest signal to error ratio as the iterations proceed. The algorithm was found to be more or less robust to the choice of the initial guess of the reconstruction.

5.4.2 Comparison of MC-CS with its CS variants

We compare the proposed MC-CS scheme against its compressed sensing (CS) counterparts for different choices of the temporal sparsifying transform Φ . We consider the choices of the (a) temporal Fourier (x-f), (b) temporal gradient (x-TV), (c) temporal PCA (x-PCA) transforms.

We initialized the MC-CS scheme with the reconstructions obtained from a spatial TV regularized method. We used the initial value of $\alpha = 4$ and used the same continuation rule updates for β , and α in all the three flavors of the MC-CS scheme. The value of Gaussian regularizer in the demons registration algorithm was fixed to a standard deviation of 10 pixels for all the three choices of the transforms. We implemented all the schemes that gave the maximum signal to error ratio (SER) between the reconstructions and the available fully sampled ground truth data:

$$SER_{ROI} = -10 \log_{10} \frac{\sum_{i=1}^N \left(\frac{\|\mathbf{\Gamma}_{recon,i} - \mathbf{\Gamma}_{ideal,i}\|_F^2}{\|\mathbf{\Gamma}_{ideal,i}\|_F^2} \right)}{N}; \quad (5.9)$$

where N is the number of time frames. During this optimization, the SER_{ROI} metric was evaluated only in a field of view that contained regions of the heart. This was motivated by recent findings in [79], and by our own experience in determining a quantitative metric that best describes the accuracy in reproducing the perfusion dynamics in different regions of the heart, and the visual quality in terms of preserving crispness of borders of heart, and minimizing visual artifacts due to reconstructions.

For experiments involving retrospective undersampling, we evaluate the reconstructions in terms of the Signal to error ratio (SER) metric and the High frequency normalized error metric (HFEN). The SER metric is defined in eq. 5.9, while the HFEN metric that gives a measure of image sharpness is defined as:

$$HFEN_{ROI} = \frac{1}{N} \sum_{i=1}^N \frac{\|\text{LoG}(\mathbf{\Gamma}_{recon,i}) - \text{LoG}(\mathbf{\Gamma}_{ideal,i})\|_F^2}{\|\text{LoG}(\mathbf{\Gamma}_{ideal,i})\|_F^2} \quad (5.10)$$

where LoG is a Laplacian of Gaussian filter that capture edges. We use the filter specified by a kernel size of 15×15 pixels, with a standard deviation of 1.5 pixels.

5.5 Results

In figure. 5.5, we show qualitative comparisons of the proposed MC-CS scheme with the temporal Fourier, temporal gradient, and temporal PCA transforms against its CS counterparts. These comparisons are shown after the recovery has been performed using 20 rays/frame. From this figure, we observe that all the three transforms benefit from motion compensation. In specific, the CS methods were sensitive to motion artifacts such as temporal stair casing (with temporal TV), motion blurring (with temporal Fourier), loss in spatio-temporal fidelity resulting in blurring of myocardial borders (with temporal PCA). In contrast, the MC-CS methods were found to be more robust to these artifacts. We observed similar trends over a range of subsampling factors as depicted in the SER and HFEN plots in figure 5.6.

In figure. 5.8, we show the comparisons involving retrospective sampling on the fully sampled in-vivo Cartesian data. The comparisons are shown using 16 radial rays per frame. Similar to the PINCAT phantom observations, we notice superior spatio-temporal fidelity and less motion artifacts with the proposed MC-CS scheme compared to its CS variants. The SER and HFEN plots in figure 5.9 also depict the same trend over a range of subsampling factors.

In figure. 5.10, we show the comparisons using prospectively undersampled radial stress free breathing data. We observe MC-CS to give better reconstructions in terms of minimizing motion blur and artifacts compared to CS. This preliminary result is demonstrated using a single coil with 21 rays, however the performance could be significantly improved by extending to multiple coils and include spatial priors (as demonstrated in chapter 3).

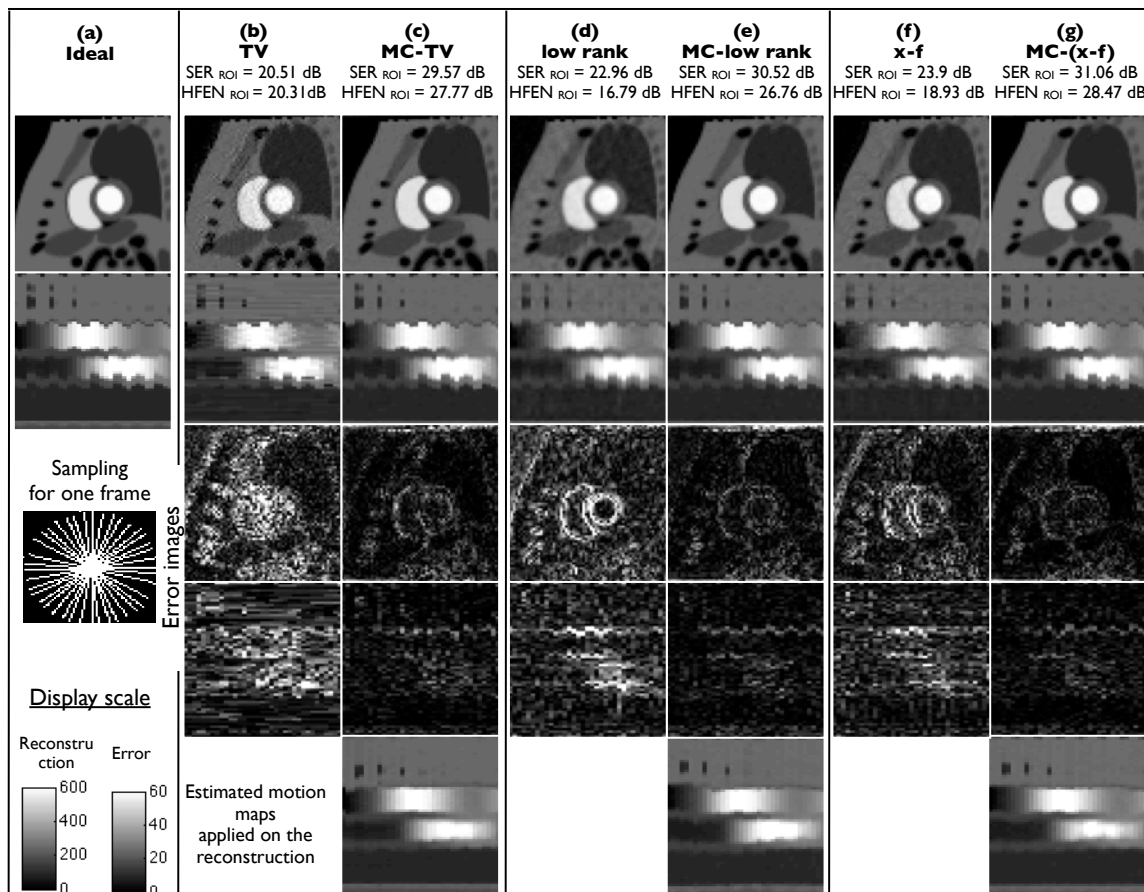


Figure 5.5: Qualitative comparison on the PINCAT phantom data using 20 rays/frame: We show for each of the scheme a spatial frame, its image time profile, and the corresponding error images and error time profiles. The error images are scaled up by a factor of 10 for better visualization. As depicted from these figures, the MC-CS schemes provide superior reconstruction quality in terms on minimizing motion artifacts compared to its CS counterparts. The motion compensated time profiles shown in the bottom row depicts that the proposed algorithm was capable to estimate and compensate most of the inter-frame motion.




Figure 5.6: Quantitative comparisons of different schemes using the signal to error ratio (a), and high frequency normalized error metrics (b) on PINCAT phantom data. These plots demonstrate that the MC-CS schemes significantly outperforms the CS schemes at all sub-sampling factors.




Figure 5.7: The fully sampled myocardial perfusion MRI dataset used as a ground-ruth during the retrospective undersampling experiments. Note the ripples in the image time profile correspond to large inter-frame motion.

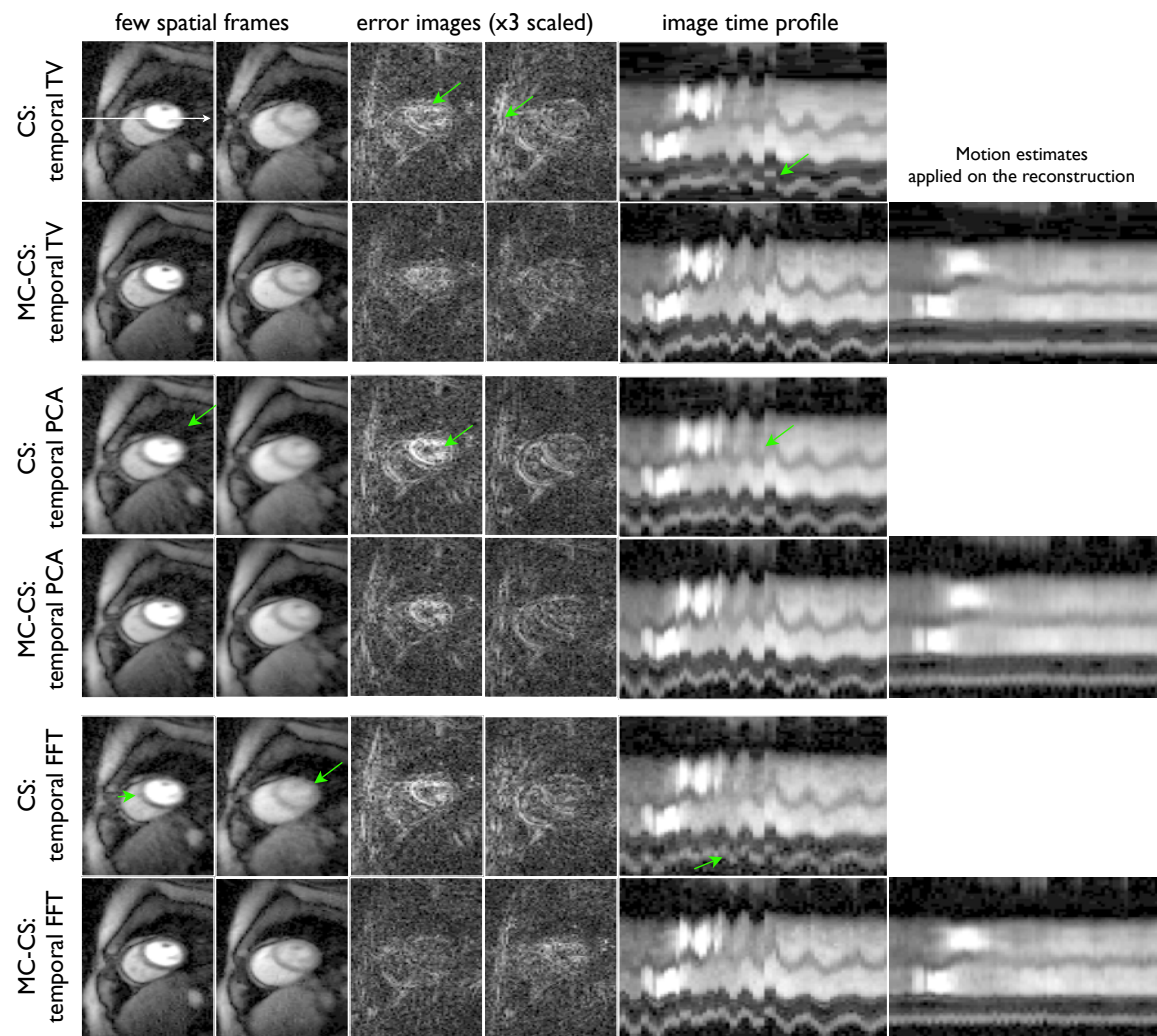


Figure 5.8: Performance evaluation using retrospectively sampling on Cartesian data using 16 radial rays/frame: The compressed sensing (CS) reconstructions depicted motion related artifacts and temporal blurring (see arrows), while the proposed motion compensated CS images (MC-CS) were robust to these compromises.




Figure 5.9: Quantitative comparisons of different schemes using the signal to error ratio (a), and high frequency normalized error metrics (b) on the in vivo myocardial perfusion data. These plots demonstrate that the MC-CS schemes significantly outperforms the CS schemes at all sub-sampling factors.

5.6 Discussion

In this chapter, we proposed a motion compensated compressed sensing framework for improving the reconstruction of free breathing myocardial perfusion MRI data. We developed a variable splitting based optimization algorithm to decouple the problem to multiple well understood subproblems. With the resulting framework, we compared the performances of utilizing motion compensation in improving CS reconstructions that relied on temporal Fourier sparsity, temporal TV sparsity, and temporal PCA regularization (using nuclear norm regularization). We observed that all the methods benefited from motion compensation. However, when compared as to which scheme benefited the most, we observed the temporal TV and the x-f based MC-CS schemes significantly outperformed the x-PCA MC-CS scheme in the experiments with in-vivo data. The x-PCA scheme tended to blur the spatio-temporal dynamics due to contrast changes even when the motion was compensated. This may be due to the usage of the nuclear norm penalty in enforcing the low rank constraint. The performance could be improved by looking at equivalent Schatten-p

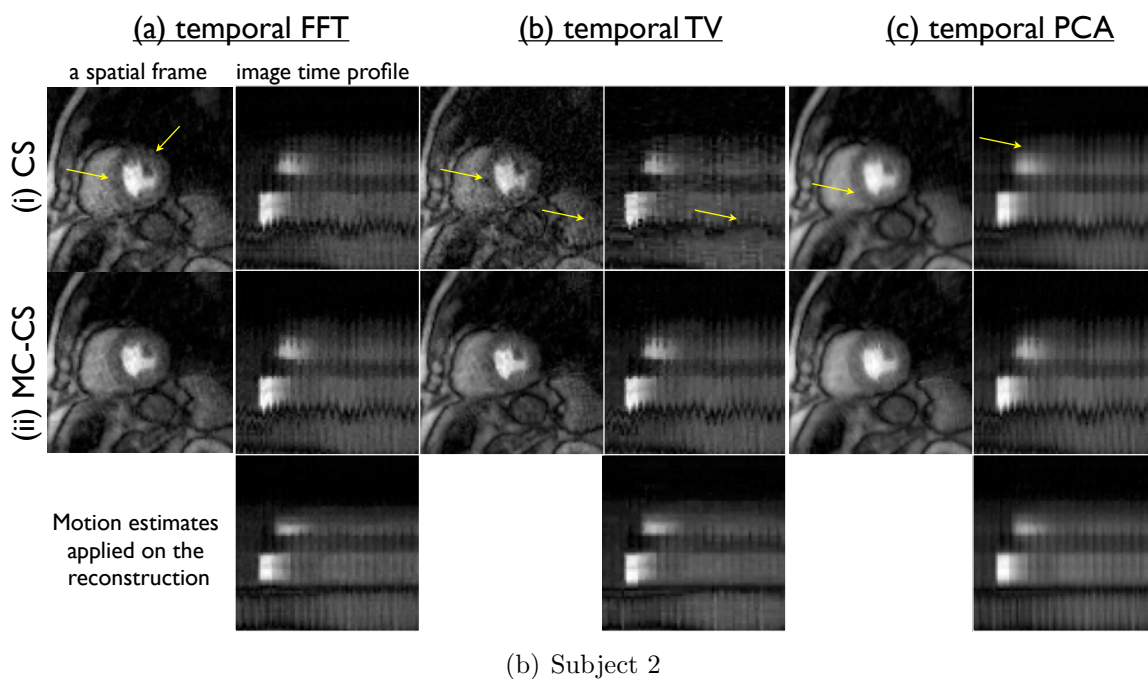
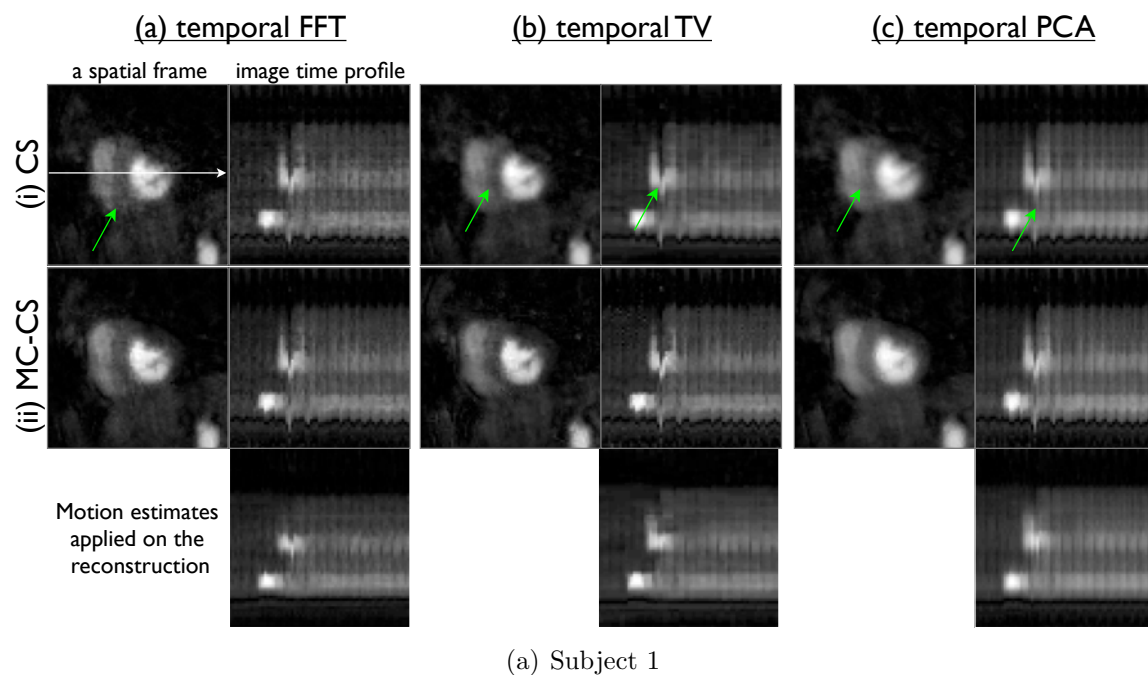


Figure 5.10: Comparison on prospectively undersampled radial data using 21 rays/frame. Here single coil reconstructions are shown on two subjects. It can be seen that MC-CS resulted in lesser motion artifacts than CS schemes in both the subjects (see the arrows that depict motion blur in CS).

norm penalties. However, we realize that for a fair conclusion in terms of comparing the methods after the motion is compensated, a comparison is needed on perfectly breath-hold datasets. Since the proposed framework decouples the sparsity based temporal denoising/de-aliasing as a stand alone problem, it is straightforward to look at multiple extensions of the framework. Extension could include adding spatial regularizers such as Wavelet, spatial total variation, or combine the benefits of different regularizers such as combining the x-PCA and the x-f or x-TV sparse priors. Furthermore, extensions to include ℓ_p norms is also possible by utilizing efficient shrinkage rules during ℓ_p minimization [19]. The reconstruction problem could also be adapted to include information from multiple coils. We look to address some of these extensions systematically in a future work.

In general, we observed the registration model with continuation was robust. In this study, with radial sampling we surprisingly found the performance of the algorithm to not be very sensitive to the choice of initial guess. This may be attributed to the continuation strategies we adapt to update our reconstructions and deformations, and the benign undersampling behavior of the radial pattern. However, a future study is required to evaluate this more on different sampling patterns. The run time of the entire algorithm was about 30 mins for datasets of size $190 \times 90 \times 70$. The implementation was based on performing all the tasks apart from the registration sub-problem on a graphical processing unit. The main bottle neck of the reconstruction time was the deformation algorithm, which was implemented on the CPU in MATLAB. The algorithm could be further optimized by using gpus and other state of the art fast implemented deformable models.

CHAPTER 6 SUMMARY AND FUTURE WORK

6.1 Summary

Myocardial perfusion MRI is a promising tool to noninvasively assess ischemic heart disease. The current commonly employed MRI methods in the clinic have limitations in terms of obtaining adequate image quality due to inherent limitations of the slow MRI acquisition. This dissertation presented novel MRI image reconstruction and acquisition methods that enable dynamic MRI reconstruction from highly accelerated measurements. The developed methods in this thesis are applicable to a wide range of dynamic imaging problems.

The key contributions of the thesis are summarized as follows:

- We have introduced a novel algorithm that exploits low rank and sparse structure of dynamic data for reconstruction from under-sampled k-t space data (k-t SLR). The proposed scheme exploits the correlations in the dynamic imaging dataset by modeling the data to have a compact representation in the Karhunen Louve transform (KLT) domain. The use of the adaptive scheme makes our approach ideally suited for a range of dynamic imaging problems. In contrast to current KLT-based methods that rely on two-step approaches to first estimate the basis functions and then use it for reconstruction, we posed the problem as a regularized matrix recovery problem. The proposed scheme uses both sparsity and spectral priors to significantly improve the recovery rate. Quantitative and qualitative comparisons on numerical phantoms and in-vivo cardiac perfusion MRI data clearly demonstrated a significant improvement in performance over existing methods.
- We have extended k-t SLR by combining the reconstruction with non-Cartesian radial sampling and parallel imaging. A fast augmented Lagrangian (AL) optimization algorithm was introduced to provide fast convergence. We showed

that the resulting developments demonstrates feasible free breathing myocardial MRI reconstruction. Comparisons on myocardial perfusion rest and stress data sets showed that $k - t$ SLR was able to achieve feasible reconstructions using few rays while being robust to artifacts such as spatio-temporal and motion blurring. We have shown preliminary feasibilities that k-t SLR can enable free breathing myocardial perfusion MRI with high spatio-temporal resolution and slice coverage (i 2x2 mm x mm, 1 heart beat, 8 slices).

- We have introduced a novel frame work for blind compressed sensing (BCS) in the context of dynamic imaging. The model represents the dynamic signal as a sparse linear combination of temporal basis functions from a large dictionary. An efficient majorize-minimize algorithm was used to simultaneously estimate the sparse coefficient matrix and the dictionary. The comparisons of the proposed algorithm with alternate BCS implementations demonstrated the computational efficiency, insensitivity to initial guesses, and the benefits of combining ℓ_1 sparsity norm with Frobenius norm dictionary constraints. Our phase transition experiments using simulated dynamic MRI data show that the BCS framework significantly outperforms conventional compressed sensing (CS) methods, and is only marginally worse than the dictionary aware case. This makes the proposed method to be highly useful in dynamic imaging applications where the signal is not sparse in known dictionaries. The validation of the BCS scheme on accelerating free breathing myocardial perfusion MRI show significant improvement over low rank models and compressed sensing schemes. Specifically, the proposed scheme is observed to be robust to spatio-temporal blurring and is efficient in preserving fine structural details.
- We have developed a novel framework that performs joint reconstruction and motion estimation in the context of compressed sensing. The novelties enabled

by this framework are (a) a generalized formulation capable of handling any temporal sparsifying transform (such as temporal Fourier, temporal gradient, temporal PCA), (b) derivation of a reference dataset that is free of motion from the measurements themselves (c) efficient decoupling of the motion estimation problem from the reconstruction problem. Unlike existing MC-CS schemes, the proposed scheme does not require fully sampled prescans or navigators for motion estimation. With validations on numerical phantoms and in vivo free breathing MR-MPI data, we have demonstrated the utility of the proposed scheme in significantly improving compressed sensing reconstructions in terms of minimizing motion artifacts.

6.2 Future work

The dynamic MRI methods developed in this thesis are ultimately targeted towards clinical utility. The proposed methods have been tested on a limited number of subjects in this thesis. To fully evaluate the efficiency and reproducibility of the developed methods in the clinical routine, a systematic study on a large number of patient scans is needed. In this section, we list a few potential improvements of the developed methods, and also list potential directions in terms of clinically validating the methods.

- Extensions to 3D imaging: 3D acquisitions have advantages over multi-slice 2D imaging in terms of providing contiguous spatial coverage, and its capability of accurate estimation of ischemic defect sizes. Extension of the 2D radial sampling patterns to 3D sampling schemes with variable density kz sampling would be more robust to downsampling, and more suitable for acceleration with parallel MRI. Furthermore, due to higher redundancies in the 4D (3D+time) dataset, the joint recovery of the 4D data using the developed methods (k-t SLR, BCS, MC-CS) could provide

better reconstructions than slice by slice (2D+time) reconstructions.

- Improving motion compensated compressed sensing: The proposed MC-CS scheme could be extended in several ways to fully exploit the power of the method. Extensions include combining multiple coils, multiple constraints such as low rank and sparsity constraints, usage of the non-convex ℓ_p norms, inclusion of spatial smoothness regularizers, and improving the computational efficiency of the registration algorithm.
- Clinical evaluation using multiple patient data : A thorough clinical evaluation of the reconstructed image quality using the developed methods is required on a cohort of patient datasets with validations against the gold standard X-ray coronary angiography. Such an analysis is required to characterize the image quality and any typical artifact levels introduced by the developed methods.

REFERENCES

- [1] V. L. Roger, A. S. Go, D. M. Lloyd-Jones, E. J. Benjamin, J. D. Berry, W. B. Borden, D. M. Bravata, S. Dai, E. S. Ford, C. S. Fox *et al.*, “Heart disease and stroke statistics—2012 update a report from the american heart association,” *Circulation*, vol. 125, no. 1, pp. e2–e220, 2012.
- [2] S. D. Fihn, J. M. Gardin, J. Abrams, K. Berra, J. C. Blankenship, P. S. Douglas, J. M. Foody, T. C. Gerber, A. L. Hinderliter, S. B. King *et al.*, “2012 accf/aha/acp/aats/pcna/scai/sts guideline for the diagnosis and management of patients with stable ischemic heart disease: a report of the american college of cardiology foundation/american heart association task force on practice guidelines, and the american college of physicians, american association for thoracic surgery, preventive cardiovascular nurses association, society for cardiovascular angiography and interventions, and society of thoracic surgeons,” *Journal of the American College of Cardiology*, vol. 60, no. 24, pp. e44–e164, 2012.
- [3] E. Di Bella, D. Parker, and A. Sinusas, “On the dark rim artifact in dynamic contrast-enhanced mri myocardial perfusion studies,” *Magnetic resonance in medicine*, vol. 54, no. 5, pp. 1295–1299, 2005.
- [4] M. Salerno and G. A. Beller, “Noninvasive assessment of myocardial perfusion,” *Circulation: Cardiovascular Imaging*, vol. 2, no. 5, pp. 412–424, 2009.
- [5] C. Ham, J. Engels, G. Van de Wiel, and A. Machielsen, “Peripheral nerve stimulation during mri: effects of high gradient amplitudes and switching rates,” *Journal of Magnetic Resonance Imaging*, vol. 7, no. 5, pp. 933–937, 1997.
- [6] J. Reilly, “Peripheral nerve stimulation by induced electric currents: exposure to time-varying magnetic fields,” *Medical and Biological Engineering and Computing*, vol. 27, no. 2, pp. 101–110, 1989.
- [7] P. Kellman and A. E. Arai, “Imaging sequences for first pass perfusion—a review,” *Journal of Cardiovascular Magnetic Resonance*, vol. 9, no. 3, pp. 525–537, 2007.
- [8] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, “Generalized autocalibrating partially parallel acquisitions (grappa),” *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [9] S. Plein, A. Radjenovic, J. P. Ridgway, D. Barmby, J. P. Greenwood, S. G. Ball, and M. U. Sivananthan, “Coronary artery disease: Myocardial perfusion

- mr imaging with sensitivity encoding versus conventional angiography¹,” *Radiology*, vol. 235, no. 2, pp. 423–430, 2005.
- [10] Z. Liang, H. Jiang, C. Hess, and P. Lauterbur, “Dynamic imaging by model estimation,” *International Journal of Imaging Systems and Technology*, vol. 8, no. 6, pp. 551–557, 1997.
- [11] B. Madore, “Using UNFOLD to remove artifacts in parallel imaging and in partial-Fourier imaging,” *Magn Reson Med*, vol. 48, no. 3, pp. 493–501, Sep 2002.
- [12] J. Tsao, P. Boesiger, and K. Pruessmann, “kt BLAST and kt SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations,” *Magnetic Resonance in Medicine*, vol. 50, no. 5, pp. 1031–1042, 2003.
- [13] B. Sharif and Y. Bresler, “Adaptive real-time cardiac mri using PARADISE: Validation by the physiologically improved ncat phantom,” in *ISBI*, 2007.
- [14] B. Sharif, J. A. Derbyshire, A. Z. Faranesh, and Y. Bresler, “Patient-adaptive reconstruction and acquisition in dynamic imaging with sensitivity encoding (PARADISE),” *Magnetic Resonance in Medicine*, pp. 501–513, 2010.
- [15] M. Lustig, J. Santos, D. Donoho, and J. Pauly, “kt SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity,” in *Proceedings of the 13th Annual Meeting of ISMRM, Seattle*. Citeseer, 2006, p. 2420.
- [16] R. Otazo, D. Kim, L. Axel, and D. 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.
- [17] G. Adluru, S. Awate, T. Tasdizen, R. Whitaker, and E. DiBella, “Temporally constrained reconstruction of dynamic cardiac perfusion MRI,” *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1027–1036, 2007.
- [18] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob, “Accelerated dynamic mri exploiting sparsity and low-rank structure: kt slr,” *Medical Imaging, IEEE Transactions on*, vol. 30, no. 5, pp. 1042–1054, 2011.
- [19] Y. Hu, S. G. Lingala, and M. Jacob, “A fast majorize–minimize algorithm for the recovery of sparse and low-rank matrices,” *Image Processing, IEEE Transactions on*, vol. 21, no. 2, pp. 742–753, 2012.

- [20] S. Goud, Y. Hu, and M. Jacob, "Real-time cardiac MRI using low-rank and sparsity penalties," in *Proceedings of the ISBI*.
- [21] S. Goud and M. Jacob, "Free breathing cardiac perfusion mri reconstruction using a sparse and low rank model: Validation with the physiologically improved ncat phantom," in *Communications and Signal Processing (ICCSP), 2011 International Conference on*. IEEE, 2011, pp. 236–240.
- [22] E. M. S.G.Lingala, Y.Hu, "Accelerated first pass cardiac perfusion mri using improved k-t slr," in *Proceedings of the ISBI, 2011*.
- [23] S. G. Lingala, Y. Hu, and M. Jacob, "Blind linear models for the recovery of dynamic mri data," in *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 2011, pp. 81 381V–81 381V.
- [24] E. D. M. J. Sajan Goud Lingala, Yue Hu, "Highly accelerated myocardial perfusion mri using k-t slr with parallel imaging," in *ISMRM*, 2011.
- [25] M. J. Yue Hu, Sajan Goud Lingala, "High resolution structural free breathing cardiac mri enabled by k-t slr," in *ISMRM*, 2011.
- [26] S. G. Lingala, E. DiBella, G. Adluru, C. McGann, and M. Jacob, "Accelerating free breathing myocardial perfusion mri using multi coil radial k- t slr," *Physics in medicine and biology*, vol. 58, no. 20, p. 7309, 2013.
- [27] S. G. Lingala, E. DiBella, M. Jacob *et al.*, "Accelerated imaging of rest and stress myocardial perfusion mri using multi-coil kt slr: a feasibility study," *Journal of Cardiovascular Magnetic Resonance*, vol. 14, no. Suppl 1, p. P239, 2012.
- [28] S. G. Lingala and M. Jacob, "Blind compressive sensing dynamic mri," 2013.
- [29] —, "A blind compressive sensing frame work for accelerated dynamic mri," in *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*. IEEE, 2012, pp. 1060–1063.
- [30] —, "Blind compressed sensing with sparse dictionaries for accelerated dynamic mri," in *International Symposia on Biomedical Imaging*, 2013.
- [31] —, "Accelerated dynamic mri using sparse dictionary learning," in *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 2013, pp. 885 822–885 822.

- [32] S. G. Lingala, M. Nadar, C. Chefd'Hotel, L. Zhang, and M. Jacob, "Unified reconstruction and motion estimation in cardiac perfusion mri," in *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*. IEEE, 2011, pp. 65–68.
- [33] S. G. Lingala, E. DiBella, C. Chefd'hotel, M. Nadar, and M. Jacob, "Motion compensated reconstruction for myocardial perfusion mri," *Journal of Cardiovascular Magnetic Resonance*, vol. 14, pp. 1–2, 2012.
- [34] M. S.G.Lingala, E.DiBella, "Accelerated myocardial perfusion mri using motion compensated compressed sensing (mc-cs)," in *ISMRM*, 2013.
- [35] —, "A novel framework for motion compensated compressed sensing in myocardial perfusion mri," *IEEE Transactions on Medical Imaging*, to be submitted.
- [36] Z. Liang, "Spatiotemporal imaging with partially separable functions," in *ISBI*, 2007, pp. 181–182.
- [37] C. Brinegar, Y. Wu, L. Foley, T. Hitchens, Q. Ye, C. Ho, and Z. Liang, "Real-time cardiac mri without triggering, gating, or breath holding," in *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*. IEEE, 2008, pp. 3381–3384.
- [38] C. Brinegar, H. Zhang, Y. Wu, L. Foley, T. Hitchens, Q. Ye, D. Pocci, F. Lam, C. Ho, and Z. Liang, "Real-time cardiac MRI using prior spatial-spectral information," in *International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 1, 2009, p. 4383.
- [39] H. Pedersen, S. Kozerke, S. Ringgaard, K. Nehrke, and W. Y. Kim, "k-t PCA: temporally constrained k-t BLAST reconstruction using principal component analysis," *Magn Reson Med*, vol. 62, no. 3, pp. 706–716, Sep 2009.
- [40] H. Jung, J. Park, J. Yoo, and J. C. Ye, "Radial k-t FOCUSS for high-resolution cardiac cine MRI," *Magnetic Resonance in Medicine*, Oct 2009.
- [41] E. Candes and B. Recht, "Exact matrix completion via convex optimization," *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
- [42] J. Cai, E. Candes, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *preprint*, 2008.

- [43] E. J. Candes and B. Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational Mathematics*, 2009.
- [44] K. Lee and Y. Bresler, “Admira: Atomic decomposition for minimum rank approximation,” *arXiv*, vol. 905, 2009.
- [45] —, “Guaranteed minimum rank approximation from linear observations by nuclear norm minimization with an ellipsoidal constraint,” *Arxiv preprint arXiv:0903.4742*, 2009.
- [46] R. Chartrand, “Fast algorithms for nonconvex compressive sensing: Mri reconstruction from very few data,” in *IEEE ISBI*, 2009.
- [47] J. Trzasko and A. Manduca, “Highly undersampled magnetic resonance image reconstruction via homotopic ell_0 -minimization,” *IEEE Transactions on Medical Imaging*, vol. 28, no. 1, pp. 106–121, 2009.
- [48] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [49] J. Haldar and Z.-P. Liang, “Spatiotemporal Imaging With Partially Separable Functions: A Matrix Recovery Approach,” in *Proceedings of the ISBI, 2010*.
- [50] K. Toh and S. Yun, “An accelerated proximal gradient method for nuclear norm regularized least squares,” *Preprint*, 2009.
- [51] S. Ma, D. Goldfarb, and L. Chen, “Fixed point and bregman iterative methods for matrix rank minimization,” *Preprint*, 05 2009.
- [52] Y. Wang, J. Yang, W. Yin, and Y. Zhang, “A new alternating minimization algorithm for total variation image reconstruction,” *SIAM J. Imaging Sciences*, vol. 1, no. 3, pp. 248–272, 2008.
- [53] J. Yang, Y. Zhang, and W. Yin, “A fast TVL1-L2 minimization algorithm for signal reconstruction from partial Fourier data,” *IEEE J. Special Topics Signal Processing*, to appear.
- [54] B. Recht, M. Fazel, and P. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *preprint*, 2007.
- [55] J. Trzasko and A. Manduca, “Relaxed conditions for sparse signal recovery with general concave priors,” *IEEE Trans. SP*, vol. 57, no. 11, pp. 4347–4354, 2009.

- [56] D. Wipf and S. Nagarajan, "Iterative reweighted l1 and l2 methods for finding sparse solutions," *SPARS09, Rennes, France*, 2009.
- [57] I. Gorodnitsky and B. Rao, "Sparse signal reconstruction from limited data using focuss: A re-weighted minimum norm algorithm," *IEEE Transactions on Signal Processing*, vol. 45, no. 3, 1997.
- [58] E. Candes, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted 1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 877–905, 2008.
- [59] R. Chartrand, "Exact reconstruction of sparse signals via nonconvex minimization," *IEEE Signal Processing Letters*, vol. 14, no. 10, pp. 707–710, 2007.
- [60] A. Majumdar and R. K. Ward, "An algorithm for sparse mri reconstruction by Schatten l_1/l_2 -norm minimization," *Magnetic Resonance Imaging*, vol. 29, no. 3, pp. 408–417, 2011.
- [61] A. Chambolle and P.-L. Lions, "Image recovery via total variation minimization and related problems," *Numerische Mathematik*, vol. 76, no. 2, pp. 167–188, 1997.
- [62] J. Yang and Y. Zhang, "Alternating direction algorithms for l1-problems in compressive sensing," *preprint*, 2009.
- [63] S. Wright and J. Nocedal, *Numerical optimization*. Springer, 2006.
- [64] "Accelereyes Jacket," 2010. [Online]. Available: <http://www.accelereyes.com/>
- [65] W. Segars and B. Tsui, "Study of the efficacy of respiratory gating in myocardial spect using the new 4-d ncat phantom," *IEEE Transactions on Nuclear Science*, vol. 49, no. 3, p. 675, 2002.
- [66] B. Sharif and Y. Bresler, "Adaptive real-time cardiac MRI using PARADISE: validation by the physiologically improved ncat phantom," in *4th IEEE International Symposium on Biomedical Imaging*, 2007, pp. 1020–1023.
- [67] F. Lin, "B1 simulator: Simulate b1 field for mri rf coils," *electronic recourse: <http://www.nmr.mgh.harvard.edu/~fhlin>*, 2005.
- [68] M. Jacob, "Optimized least-square nonuniform fast fourier transform," *IEEE Transactions on Signal Processing*, vol. 57, no. 6, pp. 2165–2177, 2009.

- [69] S. Ramani, T. Blu, and M. Unser, “Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms,” *IEEE Transactions on Image Processing*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [70] G. Adluru, C. McGann, P. Speier, E. Kholmovski, A. Shaaban, and E. DiBella, “Acquisition and reconstruction of undersampled radial data for myocardial perfusion magnetic resonance imaging,” *Journal of Magnetic Resonance Imaging*, vol. 29, no. 2, pp. 466–473, 2009.
- [71] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, “An augmented lagrangian approach to linear inverse problems with compound regularization,” in *Image Processing (ICIP), 2010 17th IEEE International Conference on*. IEEE, 2010, pp. 4169–4172.
- [72] M. A. Guttman, P. Kellman, A. J. Dick, R. J. Lederman, and E. R. McVeigh, “Real-time accelerated interactive MRI with adaptive TSENSE and UNFOLD,” *Magnetic Resonance in Medicine*, vol. 50, no. 2, pp. 315–21, Aug 2003.
- [73] P. Kellman, F. H. Epstein, and E. R. McVeigh, “Adaptive sensitivity encoding incorporating temporal filtering (tsense),” *Magn Reson Med*, vol. 45, no. 5, pp. 846–52, May 2001.
- [74] Z. Liang, “Spatiotemporal imaging with partially separable functions,” in *Proceedings of the ISBI*. IEEE, 2007, pp. 988–991.
- [75] S. Ramani and J. Fessler, “Parallel mr image reconstruction using augmented lagrangian methods,” *Medical Imaging, IEEE Transactions on*, vol. 30, no. 3, pp. 694–706, 2011.
- [76] J. Trzasko and A. Manduca, “Highly undersampled magnetic resonance image reconstruction via homotopic formula formulatype=,” *Medical Imaging, IEEE Transactions on*, vol. 28, no. 1, pp. 106–121, 2009.
- [77] E. DiBella, J. Fluckiger, L. Chen, T. Kim, N. Pack, B. Matthews, G. Adluru, T. Priester, S. Kuppahally, R. Jiji *et al.*, “The effect of obesity on regadenoson-induced myocardial hyperemia: a quantitative magnetic resonance imaging study,” *The International Journal of Cardiovascular Imaging (formerly Cardiac Imaging)*, pp. 1–10, 2011.
- [78] M. Maggioni, G. Boracchi, A. Foi, and K. Egiazarian, “Video denoising, deblocking and enhancement through separable 4-d nonlocal spatiotemporal transforms,” *IEEE Transactions on Image Processing*, preprint, 2011.

- [79] C. Bilen, I. Selesnick, Y. Wang, R. Otazo, D. Kim, L. Axel, and D. Sodickson, "On compressed sensing in parallel mri of cardiac perfusion using temporal wavelet and tv regularization," *IEEE. Acoustics Speech and Signal Processing (ICASSP)*, 2010 IEEE International Conference on, 2010, pp. 630–633.
- [80] S. Ravishankar and Y. Bresler, "Mr image reconstruction from highly undersampled k-space data by dictionary learning," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 5, pp. 1028–1041, 2011.
- [81] J.-P. Thirion, "Non-rigid matching using demons," in *Computer Vision and Pattern Recognition, 1996. Proceedings CVPR'96, 1996 IEEE Computer Society Conference on*. IEEE, 1996, pp. 245–251.
- [82] S. Lingala, Y. Hu, E. Di Bella, and M. Jacob, "Accelerated dynamic mri exploiting sparsity and low-rank structure: kt slr," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 5, pp. 1042–1054, May 2011.
- [83] R. Chan, E. Ramsay, E. Cheung, and D. Plewes, "The influence of radial under-sampling schemes on compressed sensing reconstruction in breast mri," *Magnetic Resonance in Medicine*, 2011.
- [84] S. Vasanawala, M. Murphy, M. Alley, P. Lai, K. Keutzer, J. Pauly, and M. Lustig, "Practical parallel imaging compressed sensing mri: Summary of two years of experience in accelerating body mri of pediatric patients," in *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*. IEEE, 2011, pp. 1039–1043.
- [85] M. Buehrer, K. Pruessmann, P. Boesiger, and S. Kozerke, "Array compression for mri with large coil arrays," *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1131–1139, 2007.
- [86] G. Adluru, E. DiBella *et al.*, "Compression2: compressed sensing with compressed coil arrays," *Journal of Cardiovascular Magnetic Resonance*, vol. 14, no. Suppl 1, p. P242, 2012.
- [87] M. Blaimer, I. P. Ponce, F. A. Breuer, P. M. Jakob, M. A. Griswold, and P. Kellman, "Temporal filtering effects in dynamic parallel mri," *Magnetic Resonance in Medicine*, vol. 66, no. 1, pp. 192–198, 2011.
- [88] M. A. Lukas, "Robust generalized cross-validation for choosing the regularization parameter," *Inverse Problems*, vol. 22, no. 5, p. 1883, 2006.
- [89] S. Ramani, Z. Liu, J. Rosen, J. Nielsen, and J. A. Fessler, "Regularization pa-

- parameter selection for nonlinear iterative image restoration and mri reconstruction using gcv and sure-based methods,” *Image Processing, IEEE Transactions on*, vol. 21, no. 8, pp. 3659–3672, 2012.
- [90] T. Wech, D. Stäb, J. C. Budich, A. Fischer, J. Tran-Gia, D. Hahn, and H. Köstler, “Resolution evaluation of mr images reconstructed by iterative thresholding algorithms for compressed sensing,” *Medical Physics*, vol. 39, p. 4328, 2012.
- [91] P. Roemer, W. Edelstein, C. Hayes, S. Souza, and O. Mueller, “The nmr phased array,” *Magnetic resonance in medicine*, vol. 16, no. 2, pp. 192–225, 2005.
- [92] L. Ying and J. Sheng, “Joint image reconstruction and sensitivity estimation in sense (jsense),” *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1196–1202, 2007.
- [93] M. Makowski, C. Jansen, I. Webb, A. Chiribiri, E. Nagel, R. Botnar, S. Kozerke, and S. Plein, “First-pass contrast-enhanced myocardial perfusion mri in mice on a 3-t clinical mr scanner,” *Magnetic Resonance in Medicine*, vol. 64, no. 6, pp. 1592–1598, 2010.
- [94] C. Brinegar, S. S. Schmitter, N. N. Mistry, G. A. Johnson, and Z.-P. Liang, “Improving temporal resolution of pulmonary perfusion imaging in rats using the partially separable functions model,” *Magnetic Resonance in Medicine*, vol. 64, no. 4, pp. 1162–1170, 2010.
- [95] B. Zhao, J. P. Haldar, A. G. Christodoulou, and Z.-P. Liang, “Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints.” *IEEE transactions on medical imaging*, 2012.
- [96] J. Trzasko, A. Manduca, and E. Borisch, “Local versus global low-rank promotion in dynamic mri series reconstruction,” in *Proc. Int. Symp. Magn. Reson. Med*, 2011, p. 4371.
- [97] B. Zhao, J. P. Haldar, C. Brinegar, and Z.-P. Liang, “Low rank matrix recovery for real-time cardiac mri,” in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 2010, pp. 996–999.
- [98] U. Gamper, P. Boesiger, and S. Kozerke, “Compressed sensing in dynamic mri,” *Magnetic Resonance in Medicine*, vol. 59, no. 2, pp. 365–373, 2008.
- [99] M. Aharon, M. Elad, and A. Bruckstein, “k-svd: An algorithm for designing

- over-complete dictionaries for sparse representation,” *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [100] M. Elad and M. Aharon, “Image denoising via sparse and redundant representations over learned dictionaries,” *Image Processing, IEEE Transactions on*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [101] S. Gleichman and Y. C. Eldar, “Blind compressed sensing,” *Information Theory, IEEE Transactions on*, vol. 57, no. 10, pp. 6958–6975, 2011.
- [102] M. Lustig, D. Donoho, and J. M. Pauly, “Sparse mri: The application of compressed sensing for rapid mr imaging,” *Magnetic resonance in medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [103] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, “An augmented lagrangian approach to the constrained optimization formulation of imaging inverse problems,” *Image Processing, IEEE Transactions on*, vol. 20, no. 3, pp. 681–695, 2011.
- [104] P. C. Hansen and D. P. O’Leary, “The use of the l-curve in the regularization of discrete ill-posed problems,” *SIAM Journal on Scientific Computing*, vol. 14, no. 6, pp. 1487–1503, 1993.
- [105] S. Ramani, T. Blu, and M. Unser, “Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms,” *Image Processing, IEEE Transactions on*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [106] J. Haldar and Z. Liang, “Spatiotemporal imaging with partially separable functions: A matrix recovery approach,” in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 2010, pp. 716–719.
- [107] C. Plathow, M. Schoebinger, C. Fink, S. Ley, M. Puderbach, M. Eichinger, M. Bock, H.-P. Meinzer, and H.-U. Kauczor, “Evaluation of lung volumetry using dynamic three-dimensional magnetic resonance imaging,” *Investigative radiology*, vol. 40, no. 3, pp. 173–179, 2005.
- [108] S. Lingala, Y. Hu, E. Dibella, and M. Jacob, “Accelerated first pass cardiac perfusion mri using improved k-t slr,” in *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*. IEEE, 2011, pp. 1280–1283.
- [109] M. Protter and M. Elad, “Image sequence denoising via sparse and redundant representations,” *Image Processing, IEEE Transactions on*, vol. 18, no. 1, pp. 27–35, 2009.

- [110] H. Jung and J. C. Ye, “Motion estimated and compensated compressed sensing dynamic magnetic resonance imaging: What we can learn from video compression techniques,” *International Journal of Imaging Systems and Technology*, vol. 20, no. 2, pp. 81–98, 2010.
- [111] H. Pedersen, H. Ólafsdóttir, R. Larsen, and H. B. Larsson, “A unifying model of perfusion and motion applied to reconstruction of sparsely sampled free-breathing myocardial perfusion mri,” in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 2010, pp. 752–755.
- [112] M. Usman, D. Atkinson, F. Odille, C. Kolbitsch, G. Vaillant, T. Schaeffter, P. G. Batchelor, and C. Prieto, “Motion corrected compressed sensing for free-breathing dynamic cardiac mri,” *Magnetic Resonance in Medicine*, 2012.
- [113] R. Otazo, D. Kim, L. Axel, D. K. Sodickson *et al.*, “Combination of compressed sensing and parallel imaging with respiratory motion correction for highly-accelerated cardiac perfusion mri,” *Journal of Cardiovascular Magnetic Resonance*, vol. 13, no. Suppl 1, p. O98, 2011.
- [114] G. Adluru, E. V. DiBella, and M. C. Schabel, “Model-based registration for dynamic cardiac perfusion mri,” *Journal of Magnetic Resonance Imaging*, vol. 24, no. 5, pp. 1062–1070, 2006.
- [115] J. Fessler, “Optimization transfer approach to joint registration/reconstruction for motion-compensated image reconstruction,” in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 2010, pp. 596–599.
- [116] J.-P. Thirion, “Image matching as a diffusion process: an analogy with maxwell’s demons,” *Medical image analysis*, vol. 2, no. 3, pp. 243–260, 1998.
- [117] H. Wang, L. Dong, J. O’Daniel, R. Mohan, A. S. Garden, K. K. Ang, D. A. Kuban, M. Bonnen, J. Y. Chang, and R. Cheung, “Validation of an accelerated ‘demons’ algorithm for deformable image registration in radiation therapy,” *Physics in Medicine and Biology*, vol. 50, no. 12, p. 2887, 2005.
- [118] “Demons deformable image registration,” Sep 2008. [Online]. Available: <http://www.mathworks.com/matlabcentral/fileexchange/21451-multimodality-non-rigid-demon-algorithm-image-registration>
- [119] B. Sharif and Y. Bresler, “Physiologically improved ncat phantom (pincat) enables in-silico study of the effects of beat-to-beat variability on cardiac mr,” in *Proceedings of the Annual Meeting of ISMRM, Berlin*, 2007, p. 3418.