

Fakultät Informatik

Name der promotionsführenden Einrichtung

Advanced Reconstruction Techniques in Free-Breathing Multi-Contrast High-Resolution Cardiac
Magnetic Resonance Imaging

Titel der wissenschaftlichen Abhandlung

Aurelien Bustin

Vorname und Name

Vollständiger Abdruck der von der Fakultät für Informatik

der Technischen Universität München zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften
genehmigten Dissertation.

Vorsitzende/-r: Prof. Dr. Nassir Navab

Prüfende/-r der Dissertation:

1. Prof. Dr. Darius Burschka

2. Prof. Dr. Jacques Felblinger

Die Dissertation wurde am 07.04.2017 bei der Technischen Universität München
eingereicht und durch die Fakultät für Informatik am 28.08.2017 angenommen.

ADVANCED RECONSTRUCTION TECHNIQUES IN
FREE-BREATHING MULTI-CONTRAST HIGH-RESOLUTION
CARDIAC MAGNETIC RESONANCE IMAGING

Aurélien Bustin

November 2017

Abstract

Cardiovascular magnetic resonance (CMR) imaging is a valuable noninvasive tool for high-resolution myocardial structure function, and tissue assessment, providing essential information for clinical diagnosis and treatment decisions in cardiovascular disease without ionizing radiation. The unique ability of CMR to manipulate contrast and morphological information is however challenged by long scan time and physiological motion artifacts, and is strongly influenced by system imperfections leading to noisy acquisition and poor through-plane resolution. In this thesis, we aim to bridge the gap between denoising and motion-correction approaches by developing efficient and reliable reconstruction techniques that jointly reconstruct high-resolution cardiac images free of motion and noise artifacts.

In the first study, a novel algorithm for image enhancement and noise reduction is developed and analyzed. The proposed algorithm utilises the so-called Beltrami regularization to produce high denoising performances for basic two-dimensional (2D) and multi-channel denoising problems. With the proposed mathematical framework and implementation, the benefits of the geometric Beltrami regularizer become available at no extra computational cost, compared to state-of-the-art denoising schemes.

In myocardial T_1 mapping, accurate and precise T_1 maps are typically required to quantify diffuse myocardial diseases. Among all cardiac mapping sequences, saturation-recovery-based myocardial T_1 mapping techniques have shown good accuracy, but suffer from low precision of T_1 values and high level of noise, compared to inversion-recovery-based sequences. The inherent noise penalty in myocardial T_1 mapping can be minimized using reliable denoising algorithms. The technique proposed in this study further improves the precision of myocardial T_1 mapping techniques by

extending the previously developed Beltrami algorithm method to multi-contrast T_1 -weighted denoising. The method is evaluated in simulation, volunteers and patients, leading to improvements in T_1 values precision and map quality.

In the third study, an efficient motion correction technique is developed and implemented for reconstructing a high-resolution, high-signal-to-noise ratio (SNR) image from multiple accelerated single-shot images. 2D single-shot imaging is used to allow motion fields to be easily extracted and employed for motion compensated reconstruction. In the reconstruction step, the highly efficient feature-preserving Beltrami regularization scheme is proposed for recovering sharp details at a low computational cost. The proposed technique yields robust high-SNR motion free reconstruction, and becomes an essential tool when scanning patients with arrhythmia or poor breath-hold compliance.

The balanced steady-state free precession (b-SSFP) sequence is one of the most widely used imaging techniques for cardiac cine imaging. This sequence provides the highest SNR ratio per unit time among all known sequences, with good T_2/T_1 contrast. Moreover, when slice thickness is not too small (i.e. 5 to 10 mm) the 2D b-SSFP does not suffer from severe motion-induced signal dropouts due to its fully balanced gradients. Unlike its three-dimensional (3D) version, the 2D b-SSFP provides excellent contrast between tissues and blood/vessels due to the inflow effect. A multi-slice 2D b-SSFP acquisition is therefore a good candidate, but suffers from poor resolution in the slice direction. One possible way to overcome this limitation is to combine multiple 2D stacks acquired in different orientations using a super-resolution (SR) reconstruction. In this study, a technique combining motion-correction and super-resolution reconstruction is proposed for high-resolution 3D isotropic cine imaging of the heart, and evaluated in volunteer and patient experiments.

Finally, a novel technique that combines geometrical and statistical SR using a 3D self-similarity learning framework is proposed to reconstruct high-resolution and high-SNR isotropic 3D MR images. The technique aims to achieve two goals: (i) merge information from multiple 3D anisotropic volumes acquired in different orientations into one single isotropic volume and (ii) recover sharp edges and thin anatomical structures using the self-similarity prior, which is hypothesized to be more robust in

the presence of high noise levels and artifacts. The method is evaluated in numerical simulations, and in actual MRI experiments with volunteer and patient data. Preliminary results in clinical delayed enhancement CMR on patients with myocardial infarction show that the proposed isotropic 3D reconstruction provides improved information to the cardiologist at no cost for assessing the location and transmural extent of myocardial fibrosis.

In conclusion, the set of reconstruction methods described in this thesis, as well as their applications, address several challenges in typical image acquisition that compose routine cardiac MR examination. These new advances might enable the acceleration of MRI scan, reduction of motion and noise artifact, better visualization of cardiac images, and ultimately improve clinical diagnosis and patient comfort.

Acknowledgments

This adventure actually started four years ago when I landed at the University of California in Los Angeles. Back then, I was striving to learn more about the field of compressed sensing and, with some help, was able to integrate the prestigious department of Mathematics at UCLA. This experience, although challenging and stimulating, has not only given me an abundance of scientific knowledge, but a great new look deep within myself and the way I look at science.

I am greatly indebted to Dominique Zosso, who has been a great mentor, colleague and friend. Today a professor at the Montana State University, he has helped me to look at science from a different perspective, sharing his experience in research, and teaching me the fine art of applied Mathematics. When you are collaborating with mathematicians like Dominique, Luminita Vese, and Stanley Osher at UCLA, you have really been given a chance to perform in the best Orchestra at the best Opera.

In April 2014, I integrated the prestigious Marie Curie European Network *BERTI: Biomedical Imaging and Informatics*. This experience enables me to extend my mathematical tools to the field of cardiovascular magnetic resonance by establishing a unique interdisciplinary connection with scientists, industrial partners, and clinicians. The three years I spent as an early stage researcher within BERTI gave me a tremendous amount of opportunities, challenges, and support from many people.

My deepest gratitude goes to my supervisor Freddy Odille. The support, expertise and attention Freddy gives to his students go far beyond what any young scientist can dream of. Thank you very much Freddy for everything, I feel very fortunate to have worked with you during the last three years, you have been a huge source of inspirations, and I very much look forward to working with you again.

I would like to thank the IADI team at Brabois Hospital in Nancy for hosting me for one year. In particular Jacques, Pierre-Andre, Marine, Gabriela, Emilien, Damien, Julie, and Pauline who helped me conduct my research in a very productive environment. Thank you to radiologists and cardiologists Damien Mandry, Laurent Bonnemains, Damien Voilliot and Andrei Codreanu who provided clinical support and guidance to my research.

I am greatly indebted to many colleagues at GE Global Research in Munich for their discussions and encouragement, in particular to Martin, Andre, Anabea, Peng, Piero and Guido. But I can't mention the GE team without expressing my deepest thanks to Anne Menini, my industrial supervisor. Thank you very much for your tremendous guidance, support, and mentorship. It has been a real privilege to work with you. I would like to thank Anja Brau and Timo Schirmer for giving me the opportunity to be part of the cardiac team, GE is the best environment a PhD student can learn from.

I would like to thank all the BERTI partners who provided three years of medical trainings, solid knowledge in biomedical imaging and also business competences. I am particularly grateful to UnternehmerTUM who helped me developed my own start-up, believing in golidayz.com from the very beginning.

I would like to acknowledge Petra Dorfner, Katharina Lang and Andrea Glogger who helped me with all administrative matters during the last three years, and all the BERTI ESRs: Ming, Pedro, Yash, Teresa, Jaber, Pedro, Suat, Shufang, Miguel, Marta, Saeed, Fatih, Marwan, Bea and Xin who made this PhD a beautiful adventure.

The exceptional interdisciplinary and international networks that BERTI offers, led me to travel and spend several months abroad. I am particularly grateful to Greg Hager and Dorothee Heisenberg for hosting me at the prestigious Johns Hopkins University as well as Daniel Herzka for inviting me at the Medical School.

Thank you to the team in Stanford, John Pauly, Joseph Cheng and Tao Zhang, for giving me the great opportunity to spend some time in San Francisco and meet their fabulous team.

I would also like to thank TUM, the European Commission (Grant Number

605162) and Darius Burschka for their research support, without which this thesis would not have been possible.

Last but not least, I would like to express my deepest thanks to my parents Michelle and Christophe, my brother Thomas and sister Audrey for always being supportive of my career and choice to travel the world, and for their unconditional love. Research is a collaborative but also very solitary environment which needs a lot of support. The unending support of my family and friends is what helped me complete my work. Talking about friends, I want to thank my precious ones: Baptiste, Manu, Jo, Pierre, Ken, Arin, and Berenice.

I dedicate this thesis to my grand-fathers, models of respect, love, positivity, and intelligence. You have been traveling for too long now.

Contents

Abstract	iii
Acknowledgments	vi
1 Introduction	1
1.1 Motivation	1
1.2 Outline	3
2 MRI Fundamentals	7
2.1 Spin Nuclei Behavior	8
2.1.1 T_1 relaxation	10
2.1.2 T_2 relaxation	10
2.2 Magnetic Resonance Imaging (MRI)	11
2.2.1 Slice Selection	12
2.2.2 Phase Encoding	13
2.2.3 Frequency Encoding	14
2.2.4 Signal Generation	14
2.3 Imaging Sequences	17
2.3.1 2D single-shot late gadolinium enhancement	17
2.3.2 3D late gadolinium enhancement	20
2.3.3 2D myocardial T_1 mapping with SMART ₁ Map	21
3 MRI Reconstruction	24
3.1 Inverse Problems	24

3.1.1	SENSE reconstruction	26
3.2	Sparsity and Regularization	29
3.2.1	Compressed sensing	29
3.2.2	Wavelet sparsity	31
3.2.3	Total variation regularization	32
3.2.4	Beltrami regularization	33
3.2.5	Patch-Based regularization	36
3.3	Reconstruction methods	37
3.3.1	Nonlinear Conjugate Gradient (NLCG)	38
3.3.2	Primal-Dual Gradient Descent (PDGD)	39
3.3.3	Iterative Soft/Hard Thresholding	40
4	Primal-Dual Optimization with Beltrami	42
4.1	Introduction	43
4.2	Primal-Dual Projected Gradients	46
4.3	Primal-Dual Beltrami Regularization	48
4.3.1	Primal-Dual approach	48
4.3.2	Algorithm	49
4.3.3	Extension to multichannel and color images	51
4.3.4	Simplified multichannel Beltrami functional	52
4.4	Denoising Results	53
4.5	Conclusions	59
5	Myocardial T_1 mapping denoising	61
5.1	Introduction	62
5.2	Materials and Methods	63
5.2.1	Single-Frame Beltrami Denoising	63
5.2.2	Vectorized Beltrami Denoising	65
5.2.3	Anisotropic Beltrami	65
5.2.4	Denoising Implementation	66
5.2.5	T_1 Map Reconstruction	67
5.2.6	Imaging	67

5.2.7	Simulations	68
5.2.8	Ex Vivo Imaging	69
5.2.9	In Vivo Imaging	70
5.3	Results	72
5.3.1	Simulations	72
5.3.2	Ex Vivo Imaging	75
5.3.3	In Vivo Imaging	75
5.4	Discussion	80
5.5	Sub-Study A: T_1 mapping denoising and MOLLI	84
6	Motion-Correction in Free-Breathing	90
6.1	Introduction	91
6.2	Theory	93
6.2.1	General Motion Compensation Framework	93
6.2.2	Respiratory Binning for Raw Data Clustering	94
6.2.3	Beltrami-Regularized SENSE	97
6.2.4	Motion Estimation	98
6.2.5	Motion Compensated Reconstruction with Preserved-Features	99
6.3	Methods	100
6.3.1	Offline Simulation on Synthetic Data	100
6.3.2	Phantom Imaging	101
6.3.3	In-Vivo Validation Experiment with Self-Navigation	101
6.4	Results	104
6.4.1	Computing Resources	104
6.4.2	Offline Simulation and Phantom Imaging	104
6.4.3	In-Vivo Validation Experiment with Self-Navigation	106
6.5	Discussion	109
6.6	Appendix A: Processing of the self-navigation data	111
6.7	Sub-Study A: Free-Breathing 2D Radial Cardiac MRI	113
7	Super-Resolution in Cardiac Cine MRI	119
7.1	Introduction	120

7.2	Image Reconstruction	121
7.2.1	Motion-Compensated 3D Cine Reconstructions	121
7.2.2	Super-Resolution Reconstruction	122
7.3	MRI Experiments	123
7.4	Validation	124
7.5	Results	125
7.6	Discussion and Conclusion	128
7.7	Sub-Study A: Super-Resolution MRI and 3D Printing	130
8	3D Isotropic Reconstruction	134
8.1	Introduction	135
8.2	Theory	137
8.2.1	Geometrical Super-Resolution	137
8.2.2	Dictionary Learning	138
8.2.3	3D Patch Representation	139
8.2.4	Proposed Joint Denoising and SR Cost Functional	140
8.2.5	Optimization	142
8.3	Methods	146
8.3.1	Implementation	146
8.3.2	Experiments on Numerical Phantom Data	147
8.3.3	Experiments on Healthy Volunteer (Brain Imaging)	149
8.3.4	Experiments on Patients (Cardiac Imaging)	150
8.4	Results	151
8.4.1	Parameter Optimization	151
8.4.2	Experiments on Phantom Data	157
8.4.3	Experiments on Healthy Volunteer (Brain Imaging)	160
8.4.4	Experiments on Patients (Cardiac Imaging)	160
8.4.5	Discussion	163
9	Summary and Future Work	167
9.1	Summary of Contributions	167
9.2	Future Work	169

9.2.1	Primal-Dual Optimization with Beltrami Regularization	170
9.2.2	Myocardial T_1 Mapping Denoising	170
9.2.3	Motion-Correction in Free-Breathing Single-Shot MRI	172
9.2.4	Super-Resolution Cardiac Cine MRI	173
9.2.5	3D Patch-Based Isotropic MR Reconstruction	173

Bibliography		179
---------------------	--	------------

List of Tables

4.1	Beltrami and TV denoising applied to a large set of images	55
5.1	Recommended 2D cardiac T_1 Mapping using SMART ₁ Map on a GE system.	71
5.2	Comparison of PSNR (P), RMSE (R) and Sharpness index (S) for the denoising of the numerical phantom (Simulation study) with different values of noise.	74
5.3	Comparison of denoising time, accuracy (μ) and precision (SD) of the reconstructed numerical phantom map. Values are reported as the average over all noise levels.	75
5.4	Mean and standard deviation (SD) of T_1 values within the myocardium ROI in the healthy volunteer population (N=16) using SMART ₁ Map with and without the proposed denoising technique.	78
5.5	Impact of the proposed noise correction technique on T_1 values on a patient with myocardial infarction (see Figure 5.5).	80
6.1	Parameters used for the different experiments. The acquisition matrix size was 192×256	101
6.2	Recommended single-shot LGE protocol on a GE system.	102
8.1	Recommended 3D LGE protocol on a GE system.	152
8.2	Comparison of PSNR (P), RMSE (R) and Sharpness index (S) for the denoising of the numerical phantom (Simulation study) with different values of noise.	158

List of Figures

2.1	Spin precession and tip of the magnetization	9
2.2	T_1 and T_2 relaxation curves	11
2.3	Signal detection in magnetic resonance imaging	12
2.4	Slice selection by application of a linear magnetic field gradient	13
2.5	Duality between k -space and image space	15
2.6	Effect of k -space sampling on image reconstruction	16
2.7	Timing diagram of segmented and single-shot LGE pulse sequences	18
2.8	Timing diagram of the LGE pulse sequence	18
2.9	Single-shot 2D LGE images	19
2.10	3D LGE on a patient with myocardial infarction	20
2.11	Modified Look-Locker (MOLLI) acquisition diagram	22
2.12	SMART ₁ Map acquisition diagram	23
2.13	Example of a T_1 map using SMART ₁ Map	23
3.1	Description of the encoding operator E for SENSE reconstruction	27
3.2	SENSE reconstruction results on phantom	28
3.3	Wavelet sparsity in cardiac MRI	32
3.4	Total variation sparsity in cardiac MRI	34
3.5	Staircasing visualization on a 1D signal	35
3.6	Denoising methods comparison in cardiac MRI	36
4.1	Comparison of image denoising with different regularity priors	44
4.2	Feature preservation versus staircasing	53
4.3	Beltrami and TV denoising applied to a large set of images	56

4.4	Parameter choice and sensitivity	57
4.5	Simplified Color-Beltrami denoising results	58
5.1	Denoised T_1 maps: simulation study	73
5.2	Ex vivo reconstructions	76
5.3	T_1 maps on a healthy volunteer	77
5.4	Comparison of T_1 maps denoising methods	78
5.5	T_1 map reconstruction on a patient with myocardial infarction	79
5.6	Visual comparison of the standard and corrected SMART $_1$ Map	85
5.7	Qualitative analysis in volunteers	86
5.8	Pre-contrast SMART $_1$ Map with and without denoising	87
5.9	Screenshot of the developed T_1 mapping software	88
5.10	Screenshot of the integrated automatic denoising tool	89
6.1	Comparison between 2D segmented LGE and 2D single-shot LGE	92
6.2	Schematic illustration of the proposed reconstruction	95
6.3	Schematic illustration of the binning method	96
6.4	Visual comparison of the RMSE and residual error	99
6.5	Cardiac short axis reconstruction on synthetic dataset	105
6.6	Quality comparison of the motion compensated reconstruction	106
6.7	Motion-compensated reconstructions on phantom	107
6.8	Proposed motion-compensated reconstruction on volunteer	108
6.9	Automatic respiratory signal extraction for liver motion tracking	112
6.10	Scheme of the proposed motion compensated reconstruction	114
6.11	Derived self-gating signals	115
6.12	Motion compensated reconstruction results (volunteer 1)	117
6.13	Motion compensated reconstruction results (volunteer 2)	118
7.1	2D multi-slice SSFP vs. 3D SSFP	121
7.2	Motion-correction with GRICS reconstruction	122
7.3	Super-resolution scheme in MRI	124
7.4	Example 3D cine datasets from a patient	126

7.5	Example intensity profiles from two volunteers	127
7.6	Directional information contained in the native and SR images	128
7.7	Super-resolved MR reconstruction of the ex vivo human heart	131
7.8	Volume rendering and slice selection prior to 3D printing	132
7.9	Physical 3D model of the ex vivo human heart and patient	133
8.1	Flowchart of the proposed Super-Resolution framework	141
8.2	Effect of patch size	153
8.3	Effect of similarity threshold	155
8.4	Effect overlap size	156
8.5	Effect number of outer iterations	157
8.6	Isotropic 3D reconstruction comparisons	159
8.7	Visual quality comparison of isotropic reconstruction on Brain imaging	161
8.8	Comparison of the isotropic 3D reconstruction of the Brain ventricle .	161
8.9	Comparison of the isotropic 3D reconstruction of the Brain cerebellum	162
8.10	3D isotropic reconstruction on a Duchenne patient	164
8.11	3D isotropic reconstruction on a patient with myocardial infarction .	165

Chapter 1

Introduction

1.1 Motivation

Cardiovascular magnetic resonance (CMR) imaging is a valuable noninvasive tool for high-resolution myocardial structure function, and tissue assessment, providing essential information for clinical diagnosis and treatment decisions in cardiovascular disease without ionizing radiation. The unique ability of CMR to manipulate contrast and morphological information is however challenged by long acquisitions and motion artifacts, and is strongly influenced by system imperfections leading to noisy images with poor through-plane resolution.

In myocardial T_1 mapping, accurate and precise T_1 values are typically required to quantify diffuse myocardial diseases. While saturation-recovery-based T_1 mapping sequences have shown good T_1 accuracy, they particularly suffer from low precision of T_1 values resulting in maps with low visual quality, compared to inversion-recovery-based sequences. Image denoising techniques can be used to improve precision of T_1 values and reduced the inherent noise penalty in myocardial T_1 mapping.

In addition to their high-level of noise, CMR acquisitions often suffer from cardiac and breathing motion due to the long scan duration. Consequently, two-dimensional (2D) single-shot imaging sequences have been widely used for rapid image acquisition of a whole slice within a single-shot, greatly reducing scan time and motion

artifacts, compared to standard segmented acquisitions. However, this motion robustness comes at the expense of lower spatial resolution and signal-to-noise ratio (SNR). Recent techniques can be used to enhance the SNR of single-shot methods by motion correcting and averaging multiple single-shot images acquired in free-breathing. These techniques however proved to be inadequate for high scan acceleration, and may not provide optimal image quality. Novel methods suitable for reconstructing a high-resolution, high-SNR image from multiple accelerated single-shot images are therefore needed.

In functional CMR, the balanced steady-state free precession (b-SSFP) sequence has shown to provide the highest SNR per unit time among all known sequence, with good T_1/T_2 contrast. Additionally, unlike its three-dimensional (3D) version, the 2D b-SSFP provides excellent contrast between tissues and blood/vessels due to the inflow effect. A multi-slice 2D b-SSFP acquisition is therefore a good candidate but suffers from poor resolution in the slice direction. Techniques that improves through-plane resolution need to be proposed for high-resolution 3D isotropic cine imaging of the heart.

In 3D delayed enhancement imaging, acquisition noise is a very important factor affecting image quality and myocardial scar assessment. Additionally, the poor through-plane resolution of 3D late gadolinium enhancement (LGE) sequences poses a significant challenge to quantify myocardial scar volume in patient with ischemic cardiomyopathy. Migration towards isotropic high-resolution, 3D LGE techniques are today required in clinical practice.

This thesis presents efficient and robust image denoising and reconstruction techniques that aim to reduce the inherent acquisition noise, improve image resolution in both 2D and 3D MRI sequences, and address the problem of motion in typical image acquisitions that compose routine cardiac MR examination. Superior image quality are expected without compromising the important image features, such as edges, anatomical structures and image contrast. These new advances might ultimately improve patient comfort as well as facilitate the visualization and quantification of cardiac MRI data, which are essential for a better clinical diagnosis.

1.2 Outline

In the following we give an overview of the different studies presented in this dissertation, starting with the basic motivations and then providing an outline of its contributions. The thesis is organized as follows:

Chapter 2: MRI Fundamentals

This chapter gives an introduction to the basic principles of MR signal generation and acquisition, followed by an overview of three CMR sequences that are particularly relevant to the rest of the dissertation, such as single-shot imaging, myocardial T_1 mapping, and 3D late gadolinium enhancement imaging.

Chapter 3: MRI Reconstruction

The fundamental concepts of inverse problem in MRI is presented in this chapter, setting the stage for sparsity in MR images and regularized optimizations. Some fundamental sparsity-based transforms are presented, and potential applications, such as image denoising and image reconstruction, are discussed.

Chapter 4: Primal-Dual Optimization with Beltrami

The purpose of this chapter is to present a new regularization technique for image denoising. Inverse problems in imaging and computer vision are typically addressed as data-fidelity optimization problems, where data-regularizers such as H^1 or Total Variation (TV) are included to render the optimization problem well-posed. However, while H^1 regularization is known to produce overly smooth reconstructions, the TV model is feature-preserving but introduces staircasing artifacts. The geometrically derived Beltrami framework, introduced by Sochen, Kimmel and Malladi (1998) offers an ideal compromise between feature preservation and avoidance of staircasing

artifacts. Until now, one of the main limiting factor of the Beltrami regularizers was the lack of really efficient optimization schemes. In this chapter, we start from one of the most efficient TV-optimization methods, primal-dual projected gradients, and apply it to the Beltrami functional.

The main contribution in this work lies in the introduction and description of an efficient and constructive Beltrami algorithm for image enhancement. Comparison to other state-of-the-art image enhancement techniques are shown. Doing so, we achieve better performance than TV denoising for the basic grey-scale and color image denoising problems. With the proposed optimization algorithm, the benefits of the geometric Beltrami regularizer become available at no extra computational cost.

Chapter 5: Myocardial T_1 Mapping Denoising

This chapter concentrates on the development of an improved version of the Beltrami denoising algorithm presented in chapter 4, with application to myocardial T_1 mapping. We introduce a fast and robust denoising method developed for magnetic resonance T_1 mapping. The technique exploits correlations in the native spatio-temporal space. The proposed approach was assessed in simulations, ex vivo data and in vivo imaging on a cohort of sixteen healthy volunteers both in pre- and post-contrast injection. The method was evaluated in myocardial T_1 mapping at 3T with a saturation-recovery technique that is accurate but sensitive to noise. Simulations on synthetic phantom showed SNR and sharpness improvement with the proposed method in comparison with conventional denoising. In vivo results demonstrated that our method preserves accuracy, as no difference in mean T_1 values was observed in the myocardium. Meanwhile, precision was improved with standard deviations of T_1 values being significantly decreased.

The main contribution in this work is to propose an efficient denoising method which preserves accuracy and improves precision in myocardial T_1 mapping, with the potential to offer cardiologists better map visualization and analysis.

Chapter 6: Motion-Correction in Free-Breathing

In this chapter, we introduce a joint image reconstruction and motion correction method consisting of several steps including a non-rigid motion extraction and a motion-compensated reconstruction. The reconstruction includes a denoising with the Beltrami regularization, which offers an ideal compromise between feature preservation and staircasing reduction. We assessed results in simulation, phantom and volunteer experiments in terms of image quality and robustness to high acceleration factors. The proposed joint image reconstruction and motion correction method exhibits visible quality improvement over previous methods while reconstructing sharper edges. Moreover, when the acceleration factor increases, standard methods show blurry results while the proposed method preserves image quality. The presented technique enables accurate motion corrected reconstruction from single-shot images with higher spatial resolution than conventional methods, with the potential to offer high-quality delayed enhancement scans in challenging patients.

The main contribution is to develop a new efficient motion-compensated reconstruction technique for free-breathing cardiac MRI that allows high-quality and sharp images to be reconstructed from multiple undersampled single-shot acquisitions using the Beltrami regularization introduced in chapter 4.

Chapter 7: Super-Resolution in Cardiac Cine MRI

Cardiac cine MRI with 3D isotropic resolution is challenging as it requires efficient data acquisition and motion management. It is proposed to use a 2D balanced steady-state free precession (SSFP) sequence rather than its 3D version as it provides better contrast between blood and tissue. In order to obtain 3D isotropic images, 2D multi-slice datasets are acquired in different orientations (short axis, horizontal long axis and vertical long axis) while the patient is breathing freely. Image reconstruction is performed in two steps: (i) a motion-compensated reconstruction of each image stack corrects for nonrigid cardiac and respiratory motion; (ii) a super-resolution (SR) algorithm combines the three motion-corrected volumes (with low resolution in the slice direction) into a single volume with isotropic resolution. The SR reconstruction was

implemented with two regularization schemes including the conventional Tikhonov and the feature-preserving Beltrami regularization.

The method was validated in eight volunteers and ten patients with breathing difficulties. Image sharpness, as assessed by intensity profiles and by objective metrics based on the structure tensor, was improved with both SR techniques. The Beltrami constraint provided efficient denoising without altering the effective resolution.

Chapter 8: 3D Isotropic Reconstruction

Isotropic three-dimensional (3D) acquisition is a challenging task in Magnetic Resonance Imaging (MRI). Particularly in cardiac MRI, due to hardware and time limitation, current 3D acquisitions are limited by low-resolution, especially in the through-plane direction, leading to poor image quality in that dimension. To overcome this problem, super-resolution (SR) techniques have been proposed to reconstruct a single isotropic 3D volume from multiple anisotropic acquisitions. In this study, inspired by the recent progress in dictionary learning, a novel isotropic 3D reconstruction scheme is proposed that integrates non-local and self-similarity information from 3D patch neighborhoods. By grouping 3D patches with similar structures, the natural sparsity of MR images is enforced, and can be expressed by a low rank structure, leading to robust image reconstruction with high-SNR efficiency.

Experimental results in simulations, brain imaging and clinical cardiac MRI, demonstrate that the proposed joint SR and self-similarity learning framework outperforms current state-of-the-art methods. The proposed reconstruction of isotropic 3D volumes may be particularly useful for cardiac applications such as fibrosis assessment by late gadolinium enhancement MRI.

Chapter 9: Summary of Contributions and Future Work

In this chapter, a summary of the contributions presented in this thesis are outlined. Particularly, a selection of relevant future directions are introduced and discussed.

Chapter 2

MRI Fundamentals

In this first chapter, we present the fundamental principle of Nuclear Magnetic Resonance (NMR) and Magnetic Resonance Imaging (MRI). We will focus on the most important characteristics of signal generation and image acquisition in MRI, with a particular emphasis on pulse sequences later employed in the manuscript. For more practical and theoretical description of MRI, we refer the reader to [52, 13].

The principle of NMR was first introduced by Rabi in 1939 [94] and was derived from the work by Otto Stern and Walther Gerlach in the early 1900s which established that particles have intrinsic quantum properties. The first NMR experiments were simultaneously performed by Bloch and Purcell from Stanford University and MIT respectively, who both later were awarded the Physics Nobel prize in 1952. Although these studies were successfully applied later in spectroscopy [92] and tissue characterization [35], it is only in 1973, with the work of Lauterbur from New York University [69], that NMR was successfully used as an imaging modality to represent a transversal slice of water tubes. The main idea was relatively simple: since the resonance frequency of a proton is proportional to the strength of the magnetic field, by spatially varying the resonance frequency, one can spatially localize the measured NMR signal. Finally, the first human experiments were conducted by Mansfield in 1976 [75] and in 1980, NMR imaging was introduced to the medical community and is still in common use today as a reference imaging modality.

2.1 Spin Nuclei Behavior

NMR phenomenon is based on the magnetic properties of nuclei, and more precisely to their nuclear angular spin momentum, acting as a small magnetic dipole. Macroscopically, the vector sum of all magnetic moments of the nuclei, also called net magnetization, are observed. The spin \vec{S} of a particle, introduced by Pauli in 1924 [87], can be interpreted as its rotation about itself. The positively charged nuclei induces a force parallel to the axis of rotation, which is called the magnetic moment $\vec{\mu}$:

$$\vec{\mu} = \gamma \vec{S} \quad (2.1)$$

where γ is called the gyromagnetic ratio of the nucleus. In most of the MRI applications, the hydrogen atom ^1H is the most used nuclei since it is highly abundant in the human body (60-80% of the human mass is composed of water), even though other nucleus such as phosphorus (^{31}P) or sodium (^{23}Na) can also be measured. Therefore, protons from water are the main sources of signal in most MRI applications.

In the absence of an external magnetic field, the magnetic moments are randomly distributed and the resulting net magnetization is null. However, when an external static magnetic field \vec{B}_0 is applied in a specific direction (z -axis), the nuclei will align to the direction of this field. The angular momentum of the spin is given by the Bloch equation:

$$\frac{d\vec{\mu}}{dt} = \gamma \vec{\mu} \times \vec{B}_0 \quad (2.2)$$

where the solution of this equation is the precession of the magnetization vector $\vec{\mu}$ about the direction of the field \vec{B}_0 . More precisely, the nuclei precession will align in the parallel or anti-parallel state, with a slightly higher number of spins in parallel to the magnetic field (the difference is approximately six protons per million at room temperature). The precessional frequency ω of the protons is found to be proportional to the external magnetic field and is given by the Larmor equation:




Figure 2.1: (a) Spin precession and (b) Tip of the magnetization through 90° into the transverse plane $x - y$ using a RF pulse. (Courtesy of Freddy Odille)

$$\omega_0 = 2\pi f_0 = \gamma B_0 \quad (2.3)$$

The magnetic resonance phenomenon happens when an additional field $\vec{B}_1(t)$ tip the magnetization through 90° into the transverse plane $x-y$ using a RF pulse. Under these conditions, the flip angle is given by $\Delta\theta = \gamma B_1 \tau$, where B_1 is the strength of the RF magnetic field and τ is the duration of the pulse (Figure 2.1).

As soon as the protons have been excited and flipped into the transverse plane, the RF pulse is switched off and the protons start to interact between them and relax back to the equilibrium position. These interactions are the main features of the MRI signal and include: a realignment along the z axis of the magnetic field \vec{B}_0 , also called T_1 -relaxation, and a dephasing of the spins (or loss of phase coherence), describing the T_2 -relaxation.

2.1.1 T_1 relaxation

As soon as the RF pulse is switched off, the spins gradually return to the equilibrium so that the magnetization along the z axis returns back to the original state M_0 in an exponential fashion. This effect is called the spin-lattice interaction and is characterized by the T_1 relaxation time. T_1 has the property to increase with the magnetic field strength. In this case, the signal $M_z(t)$ of the magnetization vector \vec{M}_0 is given at time t by:

$$M_z(t) = |\vec{M}_0|(1 - e^{-t/T_1}) \quad (2.4)$$

where \vec{M}_0 is the magnetization vector in the equilibrium state and T_1 is the relaxation time in milli seconds. As an example, at 3T, the T_1 relaxation time is typically 800 ms for the myocardium, 250 ms for the fat, and 1200 ms for the blood pool.

2.1.2 T_2 relaxation

Apart from T_1 relaxation, spins interact between them within tissues, causing a loss of phase coherence of the precessing nuclei. This interaction, also called spin-spin relaxation, is characterized by the relaxation time T_2 , which contrary to the T_1 relaxation, is independent of the magnetic field strength in which the object is imaged.

The transverse signal component $M_{xy}(t)$ of the magnetization vector \vec{M}_0 is given at time t by the following equation

$$M_{xy}(t) = M_{xy}(0)e^{(-t/T_2)} \quad (2.5)$$

where $M_{xy}(0)$ represents the transverse magnetization signal \vec{M}_0 after the RF pulse and T_2 is the relaxation time in milli seconds (Figure 2.2). As an example, at 3T, the T_2 relaxation time is typically 30 ms for the myocardium, 60 ms for the fat, and 250 ms for the blood pool.

Considering both spin-spin and spin-lattice relaxations, the change in the magnetization during RF excitation and relaxation can be described by the following Bloch equation:

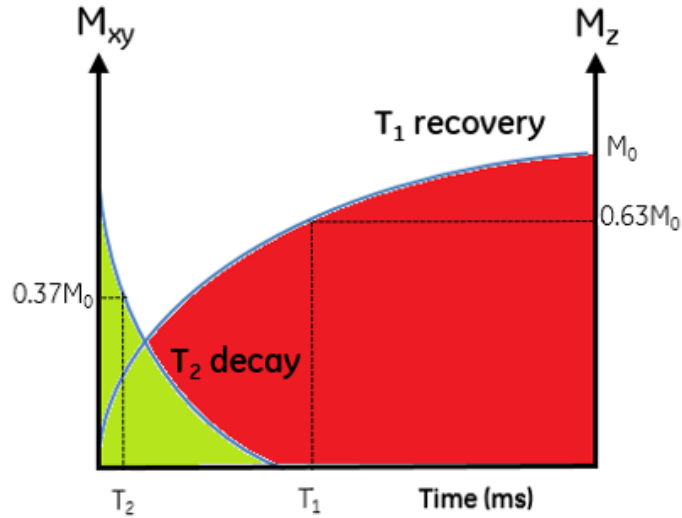


Figure 2.2: T_1 and T_2 relaxation curves. T_1 is the time taken for the magnetization to recover 63% of its equilibrium state.

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B} + \frac{1}{T_1}(M_0 - M_z)\vec{z} - \frac{1}{T_2}\vec{M}_{xy} \quad (2.6)$$

The well-known Bloch equation (2.6) is the fundamental equation describing the behavior of any MRI processes and gives a better understanding of spins evolution during RF pulse application. Once the magnetization reaches the transverse $x - y$ plane, it precesses about the z axis and generates a measurable signal which can be subsequently detected with an RF coil placed around the imaging object (Figure 2.3).

2.2 Magnetic Resonance Imaging (MRI)

In the previous Section 2.1, the magnetization signal was measured and the relaxation times were defined. In order to get the final image, it is important to spatially localize and encode the resonance signal, which is the most crucial step of MRI.




Figure 2.3: Signal detection in magnetic resonance imaging. The received FID after application of a 90° RF pulse. Due to relaxation, the received NMR signal describes exponential decay. (Courtesy of Freddy Odille)

2.2.1 Slice Selection

Since the resonance frequency f_0 is proportional to the main static magnetic field \vec{B}_0 , a simple way to localize each individual spin, with regards to their position, is to apply an additional static magnetic field, also called magnetic field gradient, linearly dependent on the spatial position. Applying an additional gradient enables the location of spins with varying resonance frequencies.

The first magnetic gradient G_z is used along the direction \vec{z} to select the slice to image by limiting the RF excitation to the slice of interest. This selection is accomplished by introducing a dependance between the static magnetic field and the spatial localization of the MR signal in the \vec{z} direction:

$$\vec{B}(z) = \vec{B}_0 + G_z \vec{z} \quad (2.7)$$

Therefore, applying an RF pulse \vec{B}_1 with frequency f_0 and bandwidth $2\Delta f$, will




Figure 2.4: Slice selection by application of a linear magnetic field gradient.

only excite the spins in the selected slice $[z_0 - \Delta z, z_0 + \Delta z]$, as shown in Figure 2.4. The gradient strength and timing thus controls the localization and size of the slice: a strong gradient or narrow bandwidth will result in a thinner slice.

2.2.2 Phase Encoding

A second magnetic field gradient G_y is then applied along the y axis after the RF pulse excitation, and therefore the precession frequency of the spins, in a specific slice, is affected for a short time. As a result, spins at different spatial position will precess at different frequencies. Once the magnetic gradient G_y has been switched off, the spins in the selected slice precess at the original frequency but the phase $\Delta\phi$ has changed linearly along the y axis. This step is referred to as the phase encoding and can be mathematically formulated into the following equation

$$\Delta\phi(y) = e^{-i\gamma G_y \tau_y y} \quad (2.8)$$

where γ is the gyromagnetic ratio, and τ_y denotes the duration of the gradient G_y along the y dimension.

2.2.3 Frequency Encoding

Finally, a third magnetic field gradient G_x is applied along the x dimension. This gradient is applied directly after the slice selection gradient and RF pulse excitation, and enables the spins to precess at different frequencies, linearly depending on their position along the x axis. This gradient is called "frequency encoding gradient" and is also referred to as the readout gradient because the MRI signal is acquired during its application. This step can be formulated into the equation:

$$\Delta\phi(x) = e^{-i\gamma G_x \tau_x x} \quad (2.9)$$

where the spins phase change at time τ_x and position x .

2.2.4 Signal Generation

In order to describe the MRI acquisition process, we introduce the following notations: $s(t)$ represents the MR signal measured during readout encoding with a RF coil, and $\rho(x, y)$ is the proton density (i.e. the transverse magnetization) in the selected slice, with $(x, y) \in \Omega$ (where the support Ω represents the non-zero entries). The received RF signal is the superposition of all the precessing magnetization (i.e. sum of harmonic functions with varying phases and frequencies) and can be described by the following equation:

$$s(G_x, G_y, t) = \int \int_{\Omega} \rho(x, y) e^{-i2\pi(k_x x + k_y y)} dx dy \quad (2.10)$$

where

$$\begin{cases} k_x = \frac{\gamma}{2\pi} \int_0^t G_x(t') dt' = \frac{\gamma}{2\pi} G_x \tau_x \\ k_y = \frac{\gamma}{2\pi} \int_0^t G_y(t') dt' = \frac{\gamma}{2\pi} G_y \tau_y \end{cases} \quad (2.11)$$

The equation (2.10) connects the proton density (i.e. the MRI signal) of a given voxel with a two-dimensional Fourier transform, also referred to as the so-called k -space (Figure 2.5). To reconstruct the image ρ from the measured signal s , it is necessary to apply a succession of experiments by varying the gradient amplitudes

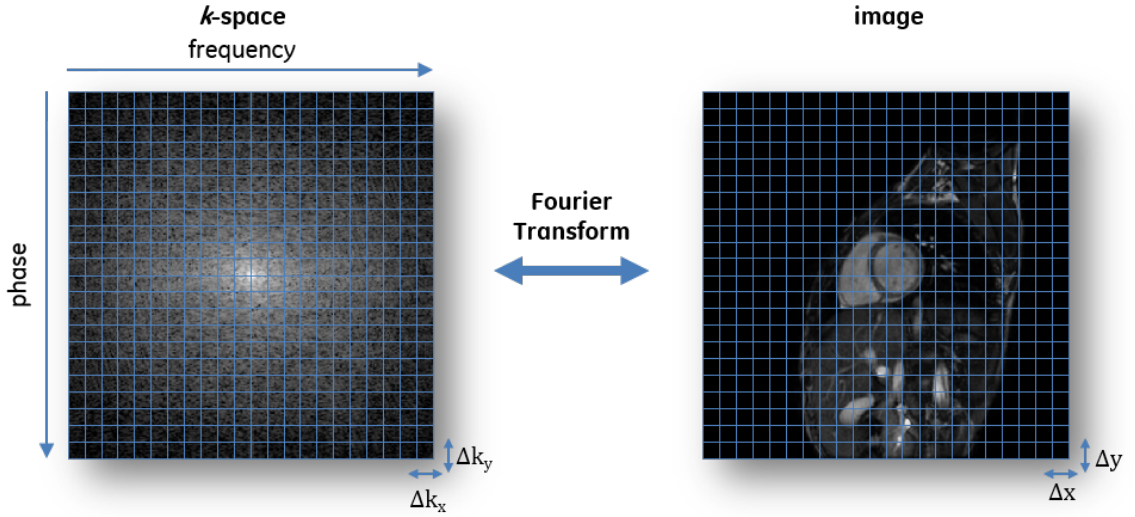


Figure 2.5: Duality between k -space and image space. (Courtesy of Freddy Odille)

G_x and G_y , in order to sample a suitable grid of (k_x, k_y) values.

Properties of a signal and its related k -space are well-known in the field of image processing. In particular, since the MR acquisition is done in a sequential and discrete fashion, it is known that for an acquired matrix of size $N_x \times N_y$ in the Fourier domain, the related image in the image space will have same dimension $N_x \times N_y$. Moreover, we have the following connection between image resolution and size of the field-of-view (FOV):

$$\begin{cases} \Delta x = \frac{1}{FOV_{k_x}} \text{ and } \Delta y = \frac{1}{FOV_{k_y}} \\ \Delta k_x = \frac{1}{FOV_x} \text{ and } \Delta k_y = \frac{1}{FOV_y} \end{cases} \quad (2.12)$$

As described previously, the k -space is sequentially filled by sampling the RF signal while varying k_x and k_y . As shown in Equation (2.10), MRI measurements are based on the Fourier transform, meaning that one can choose to fill in the k -space in a particular way in order to achieve special properties, for instance higher resolution, better contrast, or fast acquisition. The effect of k -space sampling is illustrated in Figure 2.6. Throwing away the high-frequencies (i.e. low-pass filtering) results in a blurry reconstructed image with still a good contrast (Figure 2.6b). On the other

hand, when the low-frequencies are removed from the k -space (i.e. high-pass filtering), most of the k -space energy is cut off and the recovered image only shows edges and noise (Figure 2.6c).

Selective acquisition is one of the unique features of MR imaging; strategies to design acquisition with missing complex k -space values or sparse and pseudo random undersampling has been a very active research area during the past decade, and is the main idea behind Compressed Sensing (CS) as we will see in Section 3.2.1.

In a typical clinical cardiac MR study, several imaging sequences are commonly used, each covering a specific task, e.g. functional imaging, anatomical imaging or late enhancement imaging and differ in the timing of the gradients used, the chosen excitation pulse, the class of readout or the way to acquire the complex data.




Figure 2.6: Effect of k -space sampling on image reconstruction.

2.3 Imaging Sequences

A brief overview of the CMR sequences used through the different studies is provided in this section, with a special attention given to 2D and 3D delayed enhancement imaging as well as myocardial T_1 mapping sequences. Details of typical MR scan parameters are later outlined in each sub-study.

2.3.1 2D single-shot late gadolinium enhancement

The Late Gadolinium Enhancement (LGE) sequence is designed to null signal from tissue with a particular T_1 relaxation while showing signal from tissues from tissues with other T_1 characteristics. This sequence is used to assess the viability of myocardium tissue in patients with ischemic heart disease. The presence of hyper-enhancement after the administration of gadolinium-based contrast agent indicates the presence of myocardial infarction. Visualization of the scar tissue can be accomplished using an inversion RF pulse followed by FGRE, SPGR or FIESTA readout. A cine inversion-recovery (IR) scout scan is used to determine the appropriate inversion time for signal nulling. The LGE image then contains no signal for the tissue corresponding to the selected inversion time: normal myocardium appears dark and focal scarring appears bright. The main characteristics of this sequence are as follow

1. **IR prep pulse:** Here the intention is to achieve improved robustness against spatially inhomogeneous B_1 and B_0 fields. The inversion pulse is designed to have a broad bandwidth and to be less sensitive to changes in B_1 which can be caused by calibration errors and dielectric effects at high field.
2. **Single-shot:** The single-shot LGE pulse sequence is a modification of the standard segmented LGE sequence that is used for cardiac viability imaging. Single-shot sequences acquire all image data in a single acquisition window, enabling the entire myocardium to be imaged in either a single breath-hold or during free-breathing in a single heart-beat (Figure 2.7).
3. **Fat saturation:** Two fat-selective inversion pulses are usually designed to null signal from fat. The pulse sequence takes into account the heart rate and

different levels of signal recovery due to heart rate changes. The double fat-selective inversion RF pulse scheme allows optimization of fat suppression while minimally perturbing the suppression of normal myocardium [47].




Figure 2.7: Timing diagram of segmented and single-shot LGE pulse sequences.

Figure 2.8 shows the timing diagram of the pulse sequence and a possible M_z signal evolution. At the center of k_y readout M_z of normal myocardium (green) and fat (blue) is intended to be nulled. The tissue with shortened T_1 relaxation, in this case myocardial infarct, creates a signal in the image. Fat saturation can be obtained by using a 180° fat-selective inversion pulse which flips fat back up and a second 180° fat-selective inversion pulse which inverts fat.




Figure 2.8: Timing diagram of the LGE pulse sequence with fat saturation and possible M_z signal evolution.

With single-shot LGE, images covering the entire myocardium can be obtained

in 15 to 30 seconds of free-breathing imaging (with a 2RR acquisition), greatly improving exam workflow and providing more consistent image contrast for subjects who have difficulty holding their breath or who are arrhythmic. The total readout duration ("Temporal Resolution") should remain below 250 ms.

From a clinical point of view, LGE sequences in CMR imaging can enable a clinician to visualize irreversible myocardial injury, quantify fibrosis, and allow accurate assessment of myocardial viability in patients with ischemic heart disease. However, segmented LGE acquisitions are susceptible to artifact due to patient factors such as difficulty holding their breath or an irregular heart rhythm due to frequent ectopic beats or atrial fibrillation. The use of a single-shot techniques enable to acquire diagnostic-quality images and address an essential clinical question about myocardial viability in a critically ill patient. An example of single-shot 2D LGE acquisitions in a patient with sub-endocardial delayed gadolinium enhancement is shown in Figure 2.9.




Figure 2.9: Single-shot 2D LGE images in a patient with sub-endocardial delayed gadolinium enhancement, indicative of previous myocardial infarction

2.3.2 3D late gadolinium enhancement

Three-dimensional imaging of the myocardium can provide additional information for quantification of the volume and location of the ischemic myocardial scar. As for the 2D version, the 3D LGE acquisition requires the appropriate inversion time to achieve optimal myocardium nulling, which is determined by running a cine IR scout scan. This sequence is based on the 3D enhanced fast gradient echo (EFGRE) sequence with IR preparation and cardiac gating. The 3D LGE volume then is intended to contain no signal for the tissue corresponding to the selected inversion time. The 3D LGE scan can be performed within a breath hold or with respiratory navigation, and the acquisition window that determines the temporal resolution in 3D LGE varies proportionally with the repetition time (TR), the phase encoding matrix and the phase FOV values. An example of a 3D LGE acquisition on a patient with myocardial infarction is shown in Figure 2.10.




Figure 2.10: 3D LGE whole heart coverage in one breath-hold on a patient with myocardial infarction.

2.3.3 2D myocardial T_1 mapping with SMART $_1$ Map

Many cardiac pulse sequences are currently used in the CMR community to directly quantify T_1 values for each voxel in the myocardium. The earliest attempts to measure myocardial T_1 values used the Look Locker (TI scout) method to characterize diseased tissue [58, 5]. Look Locker sequences image the heart at different inversion times and estimate the final T_1 map by fitting an exponential model through the corresponding pixels. While widely used, such a technique acquires the slices at different cardiac phases, and thus the cardiac shape can vary between the slices, resulting in inaccurate T_1 maps.

Myocardial T_1 mapping with the Modified Look-Locker (MOLLI) sequence was proposed in 2004 by Messroghli [79] as a variant of Look Locker acquisitions. This sequence consists of a single-shot TrueFISP image acquisition gated at end-diastole, allowing for the accurate reconstruction of T_1 map by merging multiple IR experiments according to their inversion times. The standard MOLLI protocol provides accurate T_1 maps with high-SNR over a wide range of T_1 s that cover the myocardial signal curve (e.g. 11 T_1 -weighted images are usually acquired over 17 heartbeats) and can be used in both pre- and post-contrast administration. An example of myocardial T_1 mapping with MOLLI is illustrated in Figure 2.11.

A major disadvantage of MOLLI is that the curve fitting yields an "apparent" T_1 of the tissue, rather than the true T_1 . The apparent T_1 , also known as T_1^* , is a function of true T_1 , heart rate, and other imaging parameters, such as flip angle, views per segment, and TR. T_1^* is also sensitive to non-ideal slice profiles and B_1 inhomogeneities. Consequently, when using IR methods, T_1^* is always shorter than the true T_1 .

The gold standard for T_1 measurement is the single-point approach, where an image at a single delay time is acquired after each magnetization preparation. The magnetization preparation may be inversion-recovery or saturation-recovery. Single-point imaging has the advantages to be independent of imaging parameters, insensitive to non-ideal slice profiles and B_1 error, and directly measures the true T_1 , therefore no T_1 correction methods are necessary.

The SMART $_1$ Map pulse sequence uses a single-point [108, 106], saturation-recovery

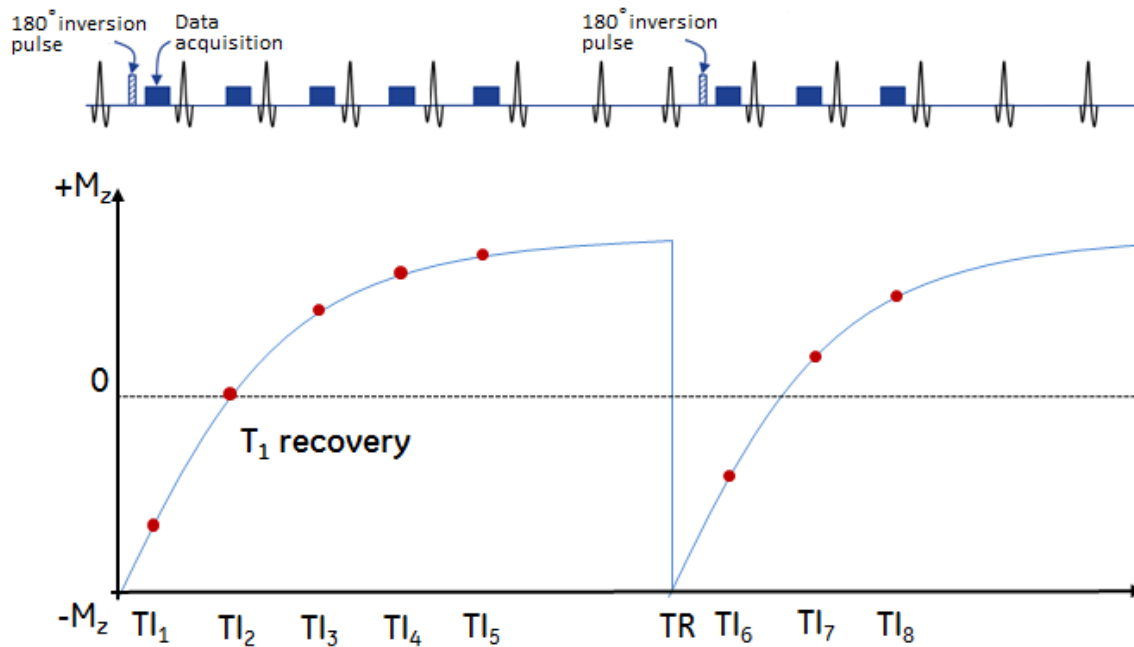


Figure 2.11: Modified Look-Locker (MOLLI) acquisition diagram.

(SR) FIESTA acquisition to measure the T_1 of myocardium and allows accurate determination of long magnetization recovery times after saturation, a feature that has not been possible with cardiac scans until now. The pulse sequence (Figure 2.12) uses a series of SR experiments, each consisting of:

1. a non-slice-selective saturation pulse
2. a saturation recovery time (TS) during which free T_1 relaxation occurs
3. a FIESTA data acquisition period.

Short TSs ($< T_{RR}$) are acquired within a single heartbeat while longer TSs ($> T_{RR}$), which have previously been unachievable in cardiac imaging, are performed across heartbeats and are therefore heart-rate dependent. In order to accurately quantify the long TSs, the pulse sequence measures every heartbeat during non-imaging portions of each cardiac cycle (Figure 2.12). The TS for each image is recorded and can be used for T_1 curve fitting during post-processing.

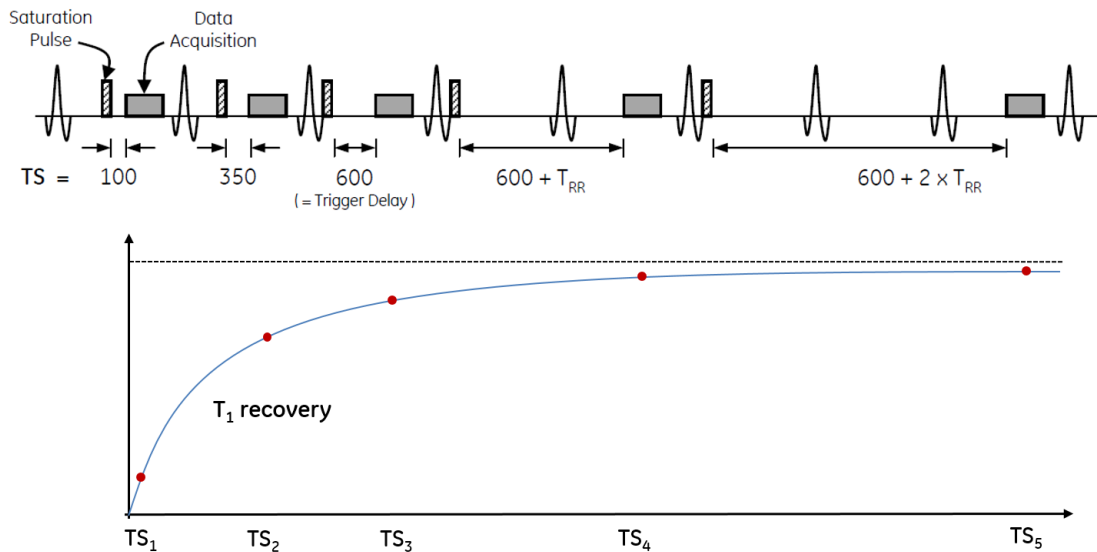


Figure 2.12: SMART₁Map acquisition diagram.

For single-shot scans, one additional heartbeat is added to the beginning of the scan in order to acquire an image with an effective infinite saturation delay time. This is achieved by not playing out the saturation pulse prior to data acquisition for this image only. This data point is important for reliable curve fitting. The additional image is annotated with a TI of about 20 sec. An example of T₁ map acquired on a healthy volunteer using SMART₁Map is shown in Figure 2.13.




Figure 2.13: Example of a T₁ map on a healthy volunteer using SMART₁Map.

Chapter 3

MRI Reconstruction

This chapter provides a succinct introduction to the notion of MR image reconstruction with emphasis on inverse problems, image sparsity and regularization. Then, we provide some examples of transform domains which exploit the sparsity of MR images to achieve accurate image reconstruction and denoising.

3.1 Inverse Problems

The previous chapter 2 derived the general framework of MR signal generation and detection, which can be written as a Fourier transform of the object of interest:

$$s(k) = \int \int_{\Omega} \rho(x, y) e^{-i2\pi(k_x x + k_y y)} dx dy \quad (3.1)$$

A discretized image can be obtained using an inverse Fourier transformation of the signal $s(k)$. Modern MR scanners acquire the MR signal, generated by the system, with a built-in body coil as well as several surface RF coils that are usually placed on the patient around the region-of-interest. The construction of such phased-array coils has been a major step and innovation in the field of MR, and led to the well-known *Parallel Imaging* field. The advantages of using parallel signal reception are mainly the acceleration of the image acquisition and a significant gain in signal-to-noise-ratio, mainly due to the reduced distance to the imaged organ compared to the body coil

integrated into the scanner bore. Indeed, the use of multiple receiver coils allows to reduce the number of phase encoding steps during the image acquisition and thus the total acquisition time.

Now considering N_γ receiver coils with local sensitivities σ_γ , the general MR signal Equation (3.1) can be rewritten for multi-coils imaging as:

$$s_\gamma(k) = \epsilon(k) \int_{\Omega} \rho(r) \sigma_\gamma(r) e^{-i2\pi k \cdot r} dr \quad (3.2)$$

where

$$r = [x \ y]^T \text{ (} r = [x \ y \ z]^T \text{ for three-dimensional imaging)}$$

$$k = [k_x \ k_y]^T \text{ (} k = [k_x \ k_y \ k_z]^T \text{ respectively)}$$

$$\epsilon(k) \text{ (is the acquisition grid of the } k\text{-space (} \epsilon(k) \in \{0, 1\} \text{))}$$

Now considering the coils $s = [s_1 \ s_2 \ \dots \ s_{N_\gamma}]$, we can rewrite the MR data acquisition problem into a simple system of linear equation by generalizing the problem in matrix form:

$$s = E\rho \quad (3.3)$$

where E is now the encoding operator comprising the coil sensitivity weighting and the Fourier sampling, and ρ is the image to recover and s the acquired k -spaces. The image ρ is thus reconstructed by solving a large linear system:

$$s_k = \sum_{j=1}^{N_\gamma} E_{k,j} \rho_j \quad \forall k \in \mathbb{K} = \{k_{x_{min}}, \dots, k_{x_{max}}\} \quad (3.4)$$

Retrieving a quantitatively accurate estimate of ρ from the operator E and data s can prove to be a challenging task. The main reasons which can potentially compromise this task includes: i) the large size of the problem, ii) the natural imperfection of the coils sensitivities, iii) the reduced amount of acquired lines (undersampling) and iv) the acquisition noise which corrupt the acquired data. These different factors make the MR reconstruction model an *ill-posed* problem, i.e., not all of the following

well-posedness conditions are true:

- *Existence of the solution:* For each s there exists a solution ρ for which $s = E\rho$
- *Uniqueness of the solution:* The image ρ is unique
- *The solution depends smoothly on the data:* Small perturbations in s do not result in large perturbations in ρ

In order to treat the ill-posedness of the the problem, MR image reconstruction is typically cast as a regularized optimization problem:

$$\underset{\rho}{\operatorname{argmin}} = \|E\rho - s\|_2^2 + \lambda R(\rho) \quad (3.5)$$

where the discretized image ρ is found by balancing the trade-off between data consistency $\|E\rho - s\|_2^2 < \epsilon$ and reduction of measurement artifacts such as noise and aliasing artifact. The regularization term $R(\rho)$ is incorporated as an additional constraint on the reconstructed image, trading-off between reduction of noise and preservation of image structures and edges. The parameter λ controls the degree of regularization in the reconstructed image and is adapted to the noise level. The performance of a specific reconstruction technique thus depends on the good choice of R as well as a robust optimization method. In particular, incorporating sparsity and statistical properties of images into the reconstruction problem have shown great potential to enhance image quality and sharpness.

3.1.1 SENSE reconstruction

In this section, we describe a technique used in MRI to accelerate the acquisition. Sensitivity Encoding (SENSE) is a widely used technique employed in clinical MRI to decrease scan time by acquiring only a reduced number of phase lines in the k -space [93]. The technique uses multiple phase-array coils to record the information, each coil having different spatial sensitivities. Information from each coil is used to unfold the aliased image caused by undersampling acquisition.

SENSE reconstruction can be written as the following inverse problem:

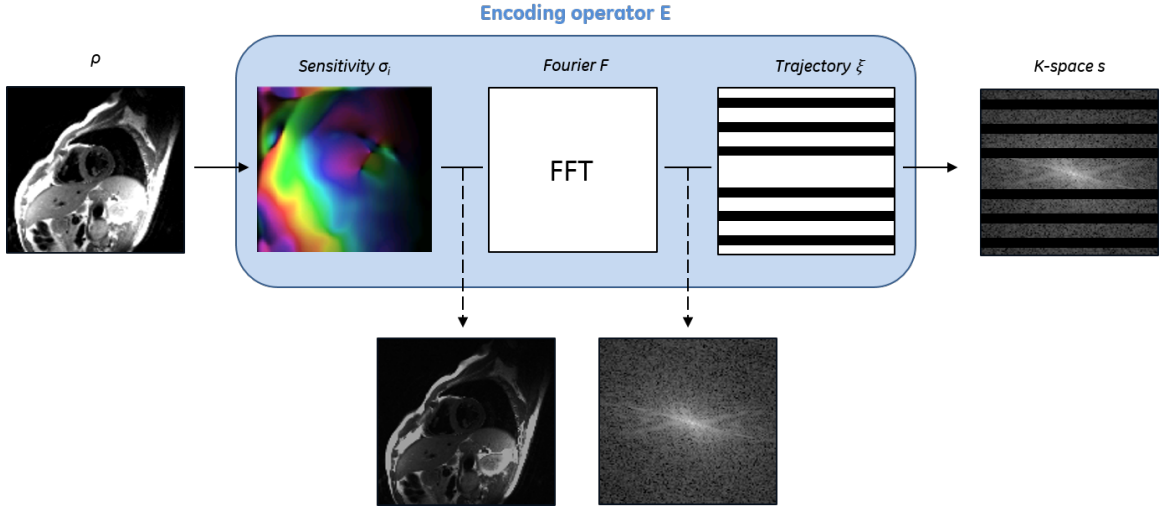


Figure 3.1: Description of the encoding operator E for SENSE reconstruction. (Courtesy of Anne Menini)

$$\underset{\rho}{\operatorname{argmin}} = \|E\rho - s\|_2^2 \quad (3.6)$$

where ρ is the image to recover, which is vectorized with size $(N_x \cdot N_y) \times 1$, the acquired k -space s is a vector of size $(N_\gamma \cdot N_x \cdot N_y) \times 1$ and the operator E is a matrix of size $(N_\gamma \cdot N_x \cdot N_y) \times (N_x \cdot N_y)$ (typically, $N_x = N_y = 256$).

$$E = \begin{bmatrix} \xi F \sigma_1 \\ \xi F \sigma_2 \\ \vdots \\ \xi F \sigma_{N_\gamma} \end{bmatrix} \quad (3.7)$$

Here F is the Fourier operator, ξ of size $(N_x \cdot N_y) \times (N_x \cdot N_y)$ is the sampling trajectory, and σ_γ of size $(N_x \cdot N_y) \times (N_x \cdot N_y)$ are the coils sensitivities. In the case of a cartesian sampling of the k -space, ξ is a diagonal matrix composed of 1 and 0 values.

The solution of the Equation (3.6) can be determined by the least square fit:

$$\rho = (E^H \Psi^{-1} E)^{-1} E^H \Psi^{-1} s \quad (3.8)$$




Figure 3.2: SENSE reconstruction results on phantom for the case of three acceleration factors with eight receiver coils.

where Ψ is the noise correlation matrix between the coil elements, E^H indicates the transposed complex conjugate matrix of E , and therefore, $E^H E$ is a symmetric matrix. Contrary to direct techniques that assume the matrix E to be known, here it is only required to know the function $x \mapsto E^H E x$. Figure 3.1 illustrates the encoding operator E used in SENSE reconstruction.

An example of SENSE reconstruction is shown in Figure 3.2. In this example, phantom data were acquired with different acceleration factors using an eight-channel array. The full k -space center, also called auto-calibration lines, was acquired in order to extract the sensitivity profiles of the coils. High image quality can be obtained from SENSE with a speed of about two, compared to standard reconstruction. It can be seen that the performance of SENSE usually drops dramatically as the acceleration

factor increases, with the appearance of noise and aliasing artifacts. When the linear equation (3.6) is ill-conditioned, for instance due to high-reduction factor, the data noise can be amplified, leading to poor reconstruction and therefore regularizations should be considered and included in the reconstruction problem.

3.2 Sparsity and Regularization

Sparsity representation has long been studied and recognized as a powerful tool in many signal and image processing problems, including image enhancement, image restoration, image segmentation, and signal compression with the recent advancements in Compressed Sensing (CS). Most medical images, and in particular MR images, are characterized by smooth and textured structures with few prominent edges and can be represented as sparse vector in specific transform domains. In this section, we give some examples of transform domains which exploit the sparsity of MR images and achieve accurate image reconstruction and denoising.

3.2.1 Compressed sensing

Sparsity of images has been intensively exploited in image compression with the recent discovery of CS, which states that a measured signal can be accurately recovered from few samples under the following three assumptions:

1. the measured samples are randomly chosen (incoherent measurements in the image or transform domain)
2. the signal is sparse in some basis (transform sparsity)
3. the reconstruction must be performed using nonlinear techniques (L_1 optimization)

A very important property of the above assumptions is the notion of sparsity in some basis, which refers here to the number of nonzero (or significant) coefficients in a specific transform domain Φ . This transform can be the identity ($\Phi = Id$) if the

image is sparse in the image domain, the Fourier transform ($\Phi = \mathcal{F}$) if the sparsity is analyzed in the frequency domain, or any other sparse transformations.

The first assumption states that only few random measurements, even obtained at a lower rate than the Nyquist condition, need to be acquired to perfectly recover the image. Unfortunately, the MRI acquisition process is generally not purely random, and searching for a sampling that makes the acquired k -space incoherent in a sparsity transform is a difficult task, and only pseudo-random sampling can be achieved in practice.

Finally, a nonlinear reconstruction has to be performed to enforce the sparsity of the image in the transform domain and the consistency of the reconstruction with the acquired MR data. This is usually accomplished by replacing the regularization term in Equation (3.5) with the l_0 norm of the sparse coefficients [71]. Thus, the general MR reconstruction problem is given by

$$\operatorname{argmin}_{\rho} = \|E\rho - s\|_2^2 + \lambda\|\Phi\rho\|_0 \quad (3.9)$$

where $\|\cdot\|_0$ counts the number of non zero components and promotes the sparsity of the reconstructed image ρ , and the data consistency term $\|E\rho - s\|_2^2$ relates to the expected noise level. The recovered image is here supposed to have a sparse representation in the basis Φ . The pseudo-random sampling pattern of the acquired k -space is explicitly included in the operator E . Despite the fact that the l_0 regularization term measures the sparsity in the image, this is not a norm in the strict sense, as it does not satisfy the convexity property of a norm:

$$\|\lambda v + (1 - \lambda)w\| \leq \lambda\|v\| + (1 - \lambda)\|w\| \quad \forall \lambda \in [0, 1] \quad (3.10)$$

The optimization problem is often called an NP-hard combinatorial problem, and approximate solutions are considered instead. A common practice is to convexify the problem posed in (3.9) by replacing the l_0 term by the l_1 -norm [26]:

$$\operatorname{argmin}_{\rho} = \|E\rho - s\|_2^2 + \lambda\|\Phi\rho\|_1 \quad (3.11)$$

The MR reconstruction problem now turns out to be convex and a solution can be found using a variety of iterative optimization techniques, as described in Section 3.3. In cardiac MRI, CS has been successfully applied in many applications, such as parameter mapping [39], delayed enhancement imaging [2], and dynamic imaging [115].

Choosing the appropriate transformation basis Φ is an important step in CS reconstruction, and can provide strong prior information for image reconstruction and efficient sparse representation. In what follows, we discuss four representative sparsity transforms.

3.2.2 Wavelet sparsity

The use of wavelets in image processing has grown considerably during the last two decades. Their performances in several image compression and denoising problems have been demonstrated through the last standard of numerical image compression with JPEG 2000. The first application of wavelets in medical imaging goes back to 1992 [121] and since then is a common practice in medical image denoising [3].

Wavelet transforms consist of generating a sparse representation by decomposing an image into its low and high frequency coefficients. An example of such a decomposition is shown in Figure 3.3, where a typical noisy cardiac late gadolinium enhancement image is transformed into the wavelet space. We see that most wavelet coefficients are very small and only few coefficients have high values. Therefore, we can extract a good sparse approximation of the cardiac image by thresholding or setting the small coefficients to zero. Figure 3.3 shows the reconstructed image when only 7% of the higher coefficients have been kept before applying the inverse wavelet transform, resulting in a uniformly better image quality with sharp edges and structures. While wavelet transforms yield higher perceptual image quality, the choice of which coefficients to keep will depend on the image itself and therefore robust nonlinear optimization techniques need to be developed (see Section 3.3).

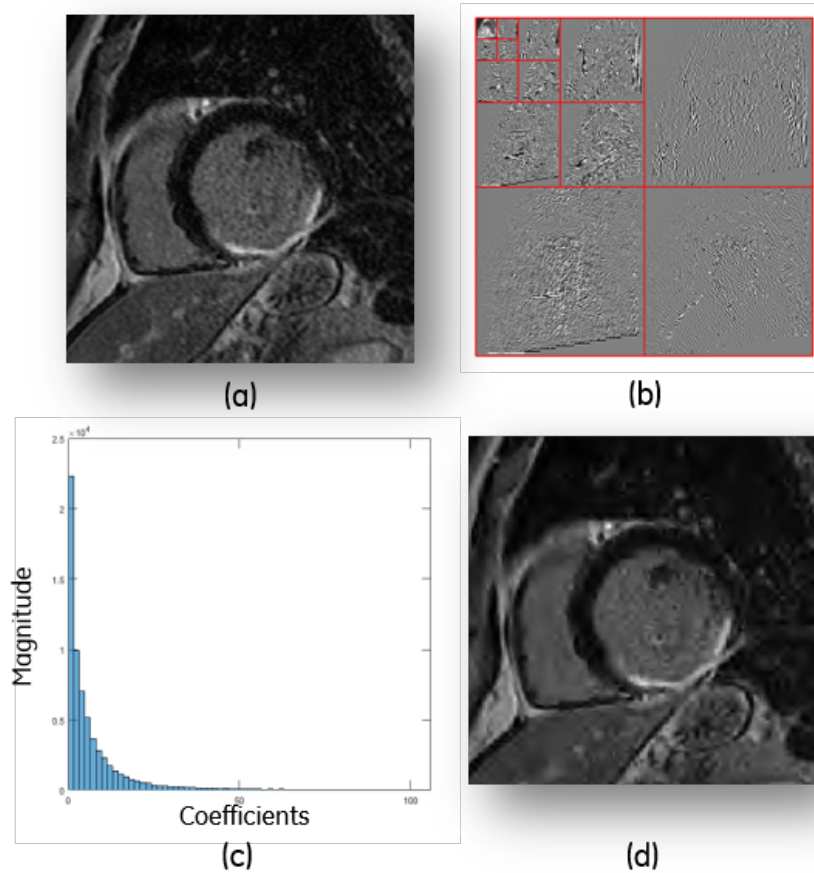


Figure 3.3: Wavelet sparsity in cardiac MRI. (a) Original noisy image, (b) its sparse wavelet coefficients of the original image, (c) histogram of the coefficients and (d) the denoised image obtained after hard thresholding the coefficients. Large wavelet coefficients are represented by white pixels, while small coefficients are represented by grey and dark pixels.

3.2.3 Total variation regularization

A typical approach for local regularization and sparsity enhancement is achieved by manipulating the derivatives of the recovered image. The simplest prior model incorporates the discretized Sobolev norm in the optimization problem, assuming the solution ρ to be smooth:

$$\operatorname{argmin}_{\rho} = \|E\rho - s\|_2^2 + \lambda\|\nabla\rho\|_2^2 \quad (3.12)$$

where ∇ represents the finite difference approximation of the gradient of the image ρ . This class of quadratic regularity, also called H^1 regularizers, have gained popularity due to computational ease, i.e. fast analytical global solution can be obtained, but suffer from overly smooth solution and loss of features preservation.

The total variation (TV) regularization has been introduced in 1992 by Rudin, Osher and Fatemi as a spatial sparsity measure [100], enabling the recovery of images with sharp edges and preserved structures. TV measures the total amplitude of pixels variations and therefore exhibits a spatially sparse gradient in the recovered image:

$$\begin{aligned} R(\rho) &= \|\nabla\rho\|_1 \\ &= \sum_{x,y} \sqrt{(\rho(x+1,y) - \rho(x,y))^2 + (\rho(x,y+1) - \rho(x,y))^2} \end{aligned} \quad (3.13)$$

The sparse representation of the gradient transform is illustrated in Figure 3.4, where histogram of pixels in the native and gradient image are represented (Figure 3.4a-c). We see that the cardiac image has a very sparse representation; only few coefficients of the absolute gradient have a large signal, compared to the signal intensity in the image space. Therefore, this sparsity in the gradient domain will lead to a very low total variation.

Besides the advantage of reconstructing a denoising image with sharp edges, TV-regularization suffers from the appearance of staircasing artifacts as shown for 1D-signal denoising in Figure 3.5.

3.2.4 Beltrami regularization

Another edge-aware feature-preserving diffusion flow stems from the Beltrami embedding [63, 107]. Most simply :

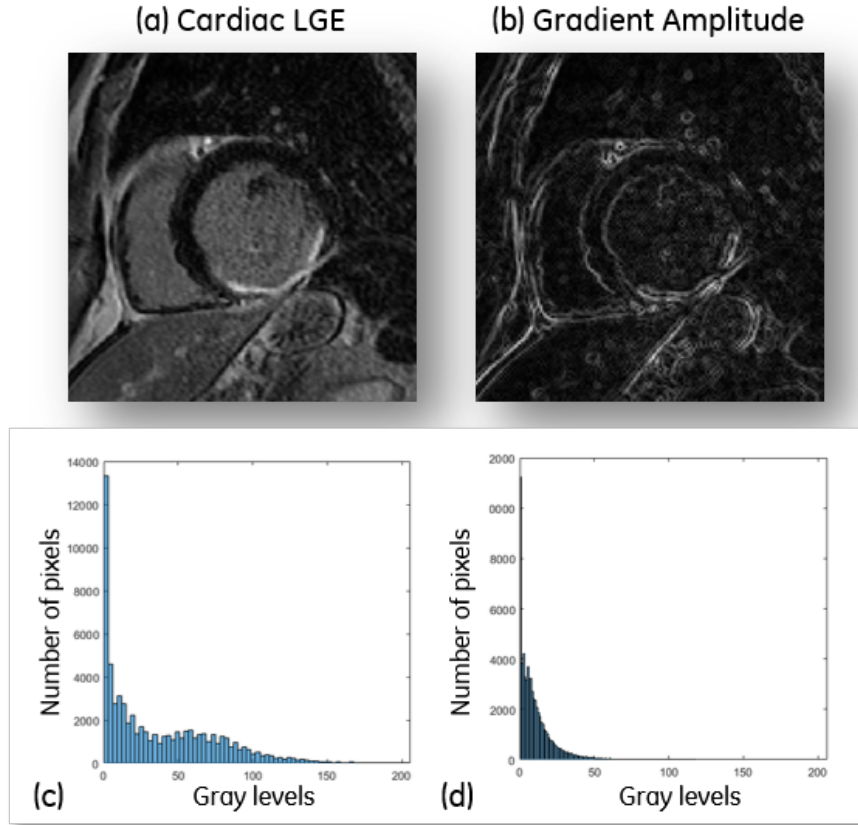


Figure 3.4: Total variation sparsity in cardiac MRI. (a) Original noisy image, (b) gradient amplitude of the image, (c) the histogram of the original image and (d) of its gradient transform.

$$\begin{aligned}
 R(\rho) &= \sqrt{1 + \beta^2 \|\nabla \rho\|^2} \\
 &= \sum_{x,y} \sqrt{1 + \beta^2 (\rho(x+1, y) - \rho(x, y))^2 + \beta^2 (\rho(x, y+1) - \rho(x, y))^2} \quad (3.14)
 \end{aligned}$$

More generally, this model is based on the Polyakov model introduced in string theory for physics [91]. The Polyakov model represents strings as harmonic maps in high-dimensional and curved spaces defined by Riemannian manifolds. Adopting this pure geometric point of view amounts to seeing objects such as images, shapes, or vector fields as geodesics or harmonic maps, much like two-dimensional topographic

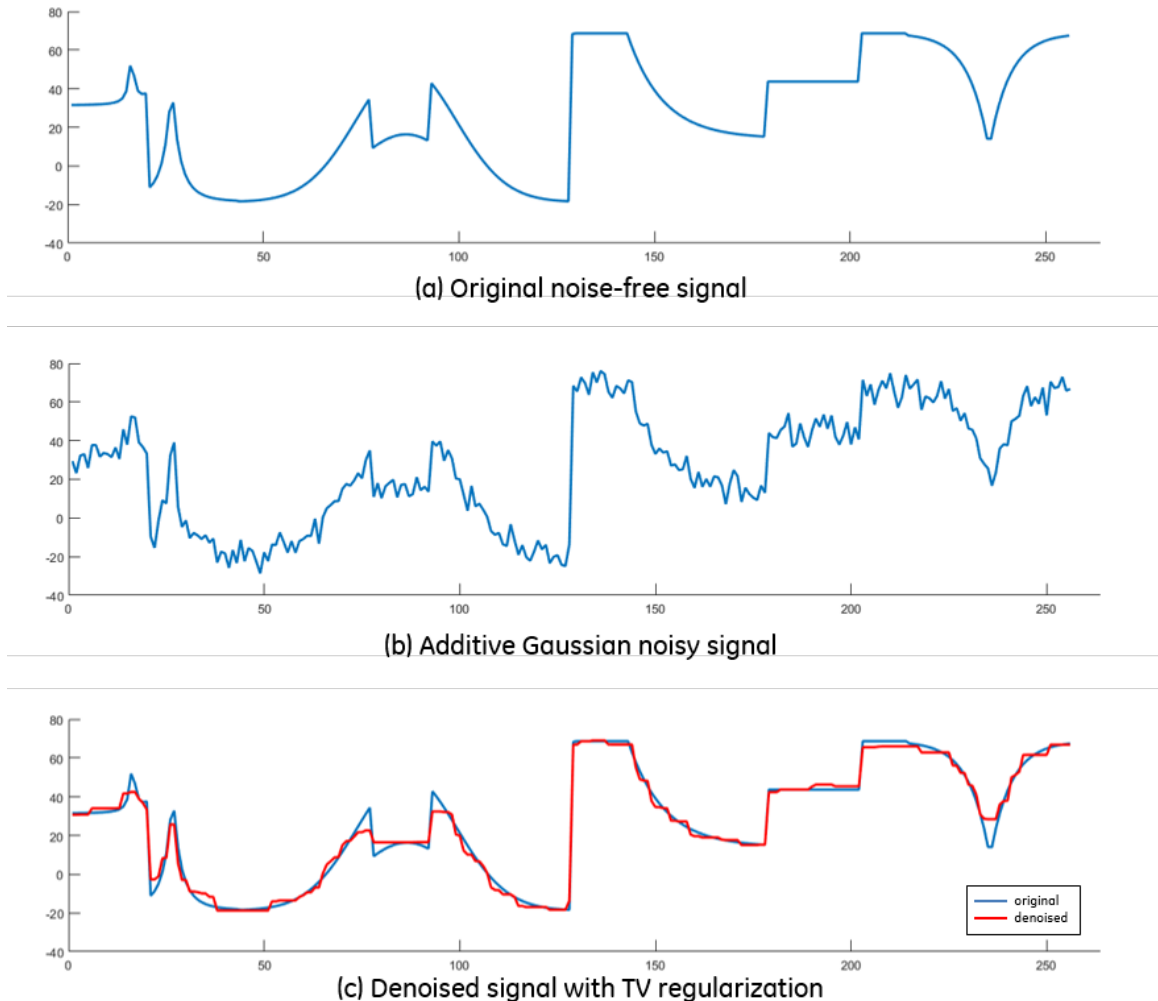


Figure 3.5: Staircasing visualization on a 1D signal denoised with TV minimization.

map corresponds to a 3D surface in the real world.

This geometric framework can be chosen such that the Polyakov energy corresponds to an interpolation between quadratic H^1 or TV gradient penalty. In Figure 3.6, we illustrate results for denoising using H^1 and Beltrami regularization on a 2D cardiac image. We can see the oversmoothing effect associated to the H^1 regularization, while Beltrami is able to combine smooth gradients with feature preservation. One of the main limiting factors of the Beltrami regularizers is the lack of efficient




Figure 3.6: A late gadolinium enhancement MR image (a), on a patient with myocardial infarction, denoised with (b) H^1 and (c) Beltrami regularizations. One can see the oversmoothing associated to the H^1 regularization. Beltrami regularization combines smooth gradients with feature preservation.

optimization schemes as developed for the TV norm and is the subject of the main study presented in Chapter 4.

3.2.5 Patch-Based regularization

Additionally to the above mentioned non-linear local gradient and wavelet-based filtering techniques, non-local denoising techniques have emerged as robust methods to preserve image structures, sharpness and efficiently reduce noise, especially when the amount of noise is significantly high. Non-local denoising techniques exploit sparsity among small similar patches, considering the repetitiveness of patterns on the whole image, e.g. structures and textures, as an indication of sparseness.

This idea is behind the success of the well-known non-local means filter (NLM), which recovers each pixel as a weighted linear combination of all the pixels in the noisy image [24]. The technique is motivated by the assumption that each pixel in the image has similar pixels in its neighborhood or in other areas of the image, not necessarily located around the pixel of interest. The NLM denoising algorithm is defined by the following equation:

$$NL(\rho)(x_i) = \sum_{x_j \in W_i} \frac{1}{Z_i} e^{-\frac{\|\rho(N_i) - \rho(N_j)\|_2^2}{h^2}} \rho(x_j) \quad (3.15)$$

where the restored intensity $NL(\rho)(x_i)$ is a weighted average of the pixel intensities $\rho(x_i)$ in the search window W_i . The local neighborhoods N_i and N_j are centered on pixels x_i and x_j . The parameter h can be seen as the regularization parameter, controlling the decay of the exponential term, and $Z_i = \sum_{x_j \in W_i} e^{-\frac{\|\rho(N_i) - \rho(N_j)\|_2^2}{h^2}}$ is a normalizing factor. The denoised estimation of $\rho(x_i)$ is thus computed as a weighted average of similar pixels in a particular neighborhood.

The NLM filter has been successfully applied to MR images and demonstrated good performance for high level of noise [74, 32]. Despite its tremendous interest in recent years, NLM is however not applicable to MRI reconstruction problems since the acquired k -spaces are usually undersampled and therefore the patch-similarity search as well as the weighted estimation Z_i would be highly corrupted by aliasing artifacts.

An extension and improvement of the NLM filter, called block-matching 3D (BM3D) has been proposed by Dabov et al [34] and is today considered as the state-of-the-art denoising technique. The method exploit the nonlocal similarity of patches by using sparse representation in a 3D linear transform domain and a shrinkage of the transform spectrum coefficients. The idea of grouping patches and 3D sparsifying similar 3D blocks considerably enhances the sparsity and shows outstanding performance in image denoising.

3.3 Reconstruction methods

When a particular sparsity basis is chosen, advanced and robust optimization techniques such as conjugate gradient descent [56], primal-dual projected gradient [29] or iterative soft/hard thresholding [40, 41, 42], can be used to reconstruct the final MR image. In this section, we briefly introduce some optimization techniques later used in the different studies.

3.3.1 Nonlinear Conjugate Gradient (NLCG)

The nonlinear conjugate gradient [53, 71] approach minimizes the cost function

$$\operatorname{argmin}_{\rho} f(\rho) = \|E\rho - s\|_2^2 + \lambda\|\Phi\rho\|_1 \quad (3.16)$$

where a backtracking line search is performed in the conjugate gradient direction of the derived cost function. The conjugate gradient requires the computation of the derivative of the objective function $f(\rho)$ which is,

$$\nabla f(\rho) = 2E^H(E\rho - s) + \lambda\nabla\|\Phi\|_1 \quad (3.17)$$

The derivative of $\|\cdot\|_1$ being not defined at zero, it is common practice to approximate the absolute value norm with a smooth function $|\rho| = \sqrt{\rho^H\rho + \epsilon}$ and define the diagonal matrix W with the diagonal elements $w_i = \sqrt{(\Phi\rho)_i^H(\Phi\rho)_i + \epsilon}$ where $\epsilon > 0$ is a smoothing factor, and Φ represents the sparse transform.

Once the gradient is computed using the previous equation, it is subtracted from the reconstructed image in a backtracking line search process, and the final image is obtained by iterating the previous step until the maximum number of iterations reached. The convergence speed of the nonlinear conjugate gradient technique is very sensitive to the regularization parameter, the step size, and varies with the imaging setup.

The nonlinear conjugate gradient procedure for MR image reconstruction is summarized in algorithm 1.

Algorithm 1 Nonlinear conjugate gradient

Initialize $\rho_0, \Delta\rho_0 = -\nabla f(\rho_0)$

repeat

 Backtracking line search for $t = \operatorname{argmin}_{t>0} f(\rho_i + t\Delta\rho_i)$

 Update final image $\rho_{i+1} = \rho_i + t\Delta\rho_i$

 Update gradient $\Delta\rho_{i+1} = -\nabla f(\rho_{i+1}) + \frac{\|\nabla f(\rho_{i+1})\|_2^2}{\|\nabla f(\rho_i)\|_2^2} \Delta\rho_i$

until convergence

3.3.2 Primal-Dual Gradient Descent (PDGD)

Another optimization technique that has successfully been used in MR reconstruction is the use of *primal-dual optimization* [64]. Primal-dual projected gradients descent algorithms are based on the convex conjugate (a.k.a. Legendre-Fenchel transform) of the main cost functional $f(\rho)$. The Legendre-Fenchel transform is used to transform the original primal problem Equation (3.5), with primal variable ρ , into an equivalent primal-dual problem. The primal-dual method has proven to be very useful for a wide variety of problems in the field of image processing and is very well adapted to gradient-based sparsity transform, e.g. TV or Beltrami.

Considering the minimization in Equation (3.5) with the sparse TV norm, the MR image reconstruction problem is given by

$$\operatorname{argmin}_{\rho} f(\rho) = \|E\rho - s\|_2^2 + \lambda \|\nabla\rho\|_1 \quad (3.18)$$

This problem can be viewed as the primal approach to solving the TV-based MR reconstruction problem, where ρ is the primal variable. It has been shown that the discrete TV of the image ρ is the Legendre-Fenchel (see Chapter 4 for the definition of the Legendre-Fenchel transform) conjugate of the characteristic function δ defined as:

$$\delta_A(d) = \begin{cases} 0 & \text{if } d \in A \\ +\infty & \text{if } d \notin A \end{cases} \quad (3.19)$$

where $A \equiv \{d : \|d\|_{\infty} \leq 1\}$. Therefore, we can get the equivalent primal-dual problem:

$$\operatorname{argmin}_{\rho} \operatorname{argmax}_{d \in A} f(\rho, d) = \lambda \|E\rho - s\|_2^2 - \langle \nabla\rho, d \rangle \quad (3.20)$$

The equation (3.20) is solved using a projection type methods by iteratively optimizing the dual d and primal ρ variables respectively, and by considering the fact that the adjoint of the gradient operator ∇ is the divergence operator $\operatorname{div} = -\nabla^H$. The primal-dual gradient procedure for MR image reconstruction is summarized in algorithm 2.

Algorithm 2 Primal-dual gradient descent

Initialize ρ_0 , d_0 , step size r_1 and r_2 **repeat**

$$\bar{d}_{k+1} = d_k - r_1 \nabla \rho_k$$

$$d_{k+1} = \frac{\bar{d}_{k+1}}{\max(|\bar{d}_{k+1}|, 1)}$$

$$\rho_{k+1} = \rho_k + r_2 \operatorname{div} d_{k+1} - \lambda r_2 E^H(E\rho_k - s)$$

until convergence

3.3.3 Iterative Soft/Hard Thresholding

An alternative formulation of the problem (3.16) consists of directly search for the optimal sparse coefficients that characterize the k -space measurements. The optimization is thus written as the following unconstrained minimization problem:

$$\operatorname{argmin}_{\phi} \|A\phi - s\|_2^2 + \lambda \|\phi\|_0 \quad (3.21)$$

where A is the matrix mapping the sparse coefficients ϕ to the measured k -space s . This formulation of the MR reconstruction problem is also known as the synthesis model for orthonormal sparse transformation $A = E\Phi$.

To solve this optimization problem, we can derive the following iterative shrinkage algorithm

$$\phi_{i+1} = \operatorname{Thresh}_{\lambda}(\phi_i + A^H(s - A\phi)) \quad (3.22)$$

where the operator *Thresh* is a thresholding operator and is either defined as the soft-thresholding operator:

$$S_{\lambda}(x_i) = \begin{cases} \operatorname{sign}(x_i)(|x_i| - \lambda) & \text{if } |x_i| \geq \lambda \\ 0 & \text{otherwise} \end{cases} \quad (3.23)$$

or the hard-thresholding operator:

$$H_{\lambda}(x_i) = \begin{cases} x_i & \text{if } |x_i| \geq \lambda \\ 0 & \text{otherwise} \end{cases} \quad (3.24)$$

Both algorithms are guaranteed to converge to a local minimum, under specific assumptions. Hard-thresholding is often found to be more effective than soft-thresholding for removing aliasing artifacts. A current review on iterative shrinkage methods can be found in [131].

Chapter 4

Primal-Dual Optimization with Beltrami Regularization¹

¹Published in: Zosso D., Bustin A. A Primal-Dual Projected Gradient Algorithm for Efficient Beltrami Regularization. UCLA Cam Report, June 2014 and presented in part at the Gordon Research Conference on Image Science, Easton, MA, USA, 2014

4.1 Introduction

Modern day imaging is mainly challenged by so-called inverse problems, namely to find an underlying image I , given derived measurements I_0 through the system K , affected by noise n , e.g. :

$$\text{find } I: I_0 = KI + n \quad (4.1)$$

Typical tasks are image restoration, reconstruction, segmentation, registration, classification, and many others. These problems are ill-posed, i.e. not all of the following conditions for well-posedness are true: (1) a solution exists, (2) the solution is unique, (3) the solution depends smoothly on the data [14]. One classical resolution scheme consists in minimizing a quadratic energy E in terms of the forward problem, while adding a quadratic regularity penalty to render the problem well-posed [112]. Formally:

$$\min_I \left\{ E = \|KI - I_0\|_2^2 + \lambda \|\Gamma I\|_2^2 \right\} \quad (4.2)$$

where typically λ is a regularization parameter and $\Gamma = \nabla$, requiring the solution to be smooth.

Although widely used today thanks to its simplicity, this resolution scheme has some important shortcomings. Despite their popularity due to computational ease (fast analytical global solution), it is well known that these H^1 regularizers are not feature-preserving and that the obtained solutions are typically overly smooth. Different anisotropic, feature-preserving regularization schemes have therefore been researched [4]. The Rudin-Osher-Fatemi (ROF) or Total Variation (TV) model [100] has particularly gained in importance, where the following minimization problem using the TV-norm is proposed. Formally:

$$\min_I \left\{ \int_{\Omega} |\nabla I| + \frac{\lambda}{2} \int_{\Omega} (I - I_0)^2 \right\}, \quad (4.3)$$

for a given balancing parameter $\lambda > 0$. Here, $\int_{\Omega} (I - I_0)^2$ is a fidelity term w.r.t. the observed, noisy image I_0 , and $\int_{\Omega} |\nabla I|$ is a total variation regularizing term, where




Figure 4.1: An image with values in $[0, 256]$ with Gaussian noise $\sigma = 10$ denoised with three different regularity priors. One can see the over-smoothing associated to the H^1 regularization, and the staircasing effect of TV, while Beltrami combines smooth gradients with feature preservation.

we suppose $I \in BV(\Omega)$, the space of functions of bounded variation, see also [117]. Soon after its introduction, the TV norm was locally weighted, in order to make the diffusion even more feature-aware and spatially adaptive [110]. The weighted TV was later rediscovered in the context of global solutions to active contours [22]. One main problem with TV-regularization is the appearance of staircasing artifacts as shown in figure 4.1.

Another, edge-aware feature-preserving diffusion flow stems from the Beltrami embedding [63, 107]. Most simply:

$$\min_I \left\{ \int_{\Omega} \sqrt{1 + \beta^2 |\nabla I|^2} + \frac{\lambda}{2} \int_{\Omega} (I - I_0)^2 \right\}. \quad (4.4)$$

A particular choice of gradient-descent preconditioner allows for a geometric Laplace-Beltrami flow, which is parametrization invariant, and which successfully reduces diffusion close to features. More generally, this model is based on the Polyakov model [91] introduced in string theory for physics. The Polyakov model represents strings as harmonic maps in high-dimensional and curved spaces defined by Riemannian manifolds. Adopting this pure geometric point of view amounts to seeing objects such as images, shapes, or vector fields as geodesics or harmonic maps, much like a 2D

topographic map corresponds to a three-dimensional surface in the real world.

The potential of this geometric framework lies in the general definition of the space-feature manifold and the choice of its metric. In particular, the metric can be chosen such that the Polyakov energy corresponds to an interpolation between quadratic or total variation gradient penalty. The features are not restricted to scalar values but include vector features encountered in color, texture or multispectral image analysis [62]. Similarly, the embedding is not limited to 2-dimensional image surfaces and generalizes naturally to n -dimensional manifolds associated to volumetric or time varying images or videos. Moreover, the choice of the metric enables the study of complex geometries inherent to scale-space methods [23] and non-flat images generated, e.g by catadioptric or omnidirectional cameras [19]. Local, pixel-based Beltrami-regularization was already successfully used for image denoising [95, 97, 96].

The direct way of implementing the Beltrami flow is obtained by using an Euler explicit forward scheme of gradient descent [12]. This scheme, however, is heavily limited by the stiffness constraint [109] and the time step has to be chosen carefully. An upper bound that ensures stability of such a scheme is provided in [37]. In order to make the scheme more robust, semi-implicit schemes have been used instead, e.g. [36, 132], largely based on work by [122]. An entirely different approach makes use of decoupling schemes that have been known in computer vision for quite some time [95, 97]. Equivalent splitting schemes have previously been applied to TV-regularized image denoising with great success [111].

Another technique that has successfully been used with TV regularization problems is the use of *primal-dual optimization* [29, 27]. These algorithms are simple and easy to implement and are potentially faster than primal gradient-descent methods. While Chambolle’s primal-dual TV algorithm has clearly led the way, more recent, very similar projected-gradient type methods were shown to be even more efficient [129, 130].

Here, we propose to use such a *primal-dual projected gradient* approach to design an efficient recovery algorithm for gray-scale images by employing the Beltrami energy as regularity prior. We present a simple yet efficient primal-dual projected gradient algorithms for Beltrami minimization with applications in image denoising.

Beyond, we present a convex simplification of the color-version of the Beltrami energy, and thereby render our algorithm applicable to multichannel images. Numerical results are presented on different images to show the performance of the primal-dual projected gradient algorithm using the Beltrami framework. We compare our results against the similarly implemented ROF model. We can show that the proposed Beltrami algorithms converge faster, and slightly improve the quality of results w.r.t. TV: Beltrami regularization maintains TV's feature preservation but reduces the staircasing effect in restored images.

The contribution of this study is not primarily a new denoising scheme, since more advanced (non-local) models exist that clearly outperform any local approach. We use the denoising problem as a simple illustration for the application of efficient Beltrami regularization through primal-dual projected gradients.

The rest of this work is organized as follows. First, we provide in Section 4.2 a short introduction of primal-dual optimization and then describe the primal-dual projected gradient algorithm as used for Beltrami regularization in Section 4.3. The method is applied to denoising problems, including a simplified generalization to color images. Denoising results and a short comparison to the related TV/ROF implementation are shown in Section 4.4. Then we conclude the study in Section 4.5. A MATLAB implementation of the proposed algorithms is available at <https://www.berti.tum.de/network-structure/esrs/aurelien-bustin/> and <http://www.math.ucla.edu/~zosso/code.html>.

4.2 Primal-Dual Projected Gradients

In this section, we provide a short description of primal-dual projected gradient optimization. We recall the main results existing in the literature to find the primal-dual model of an initial convex problem, and how to solve this new model by using a gradient projection type method. We only recall some core principal results, and refer the interested reader to, e.g. [129, 130, 44] for further details, analysis and improvements of this general method.

We start with a definition of the convex conjugate (a.k.a. Legendre-Fenchel transform) of a function f defined by

$$f^*(s) = \sup_{x \in \mathbb{R}^n} \left\{ \langle s, x \rangle - f(x) \right\}, \quad s \in \mathbb{R}^n. \quad (4.5)$$

The Legendre-Fenchel transform is used in the following primal-dual equivalence: Let $F : W \rightarrow \mathbb{R}$ be a proper closed and convex functional on the set W , G a proper closed and convex functional on the set V and let $K : V \rightarrow W$ be a continuous linear operator. Then we have the following equivalence:

$$\begin{aligned} \underbrace{\min_{x \in V} \left\{ F(Kx) + G(x) \right\}}_{\text{Primal}} &= \underbrace{\min_{x \in V} \max_{\varphi \in W^*} \left\{ \langle Kx, \varphi \rangle - F^*(\varphi) + G(x) \right\}}_{\text{Primal-Dual}} \\ &= \min_{x \in V} \max_{\varphi \in W^*} \left\{ f(x, \varphi) \right\} \end{aligned} \quad (4.6)$$

where x and φ are the primal and dual variables, respectively, F^* is the convex conjugate of F , and W^* is the dual space of W . For a proof, see [44]. This equation can be solved using a simple fixed-point method: Given $z_0 \in X := V \times W^*$ compute the solution at step $n + 1$ by iterating the scheme

$$z_{n+1} = P_X(z_n - rH(z_n)), \quad (4.7)$$

where

$$z = \begin{pmatrix} x \\ \varphi \end{pmatrix} \quad \text{and} \quad H(z) = \begin{pmatrix} \partial_x f(x, \varphi) \\ -\partial_\varphi f(x, \varphi) \end{pmatrix}. \quad (4.8)$$

until convergence (see [15]).

4.3 Primal-Dual Beltrami Regularization

Let us now return to the Beltrami regularization problem, and apply the primal-dual projected gradient method for its optimization. We first consider the simple case of denoising. Once the primal-dual projected gradient method has been described for this starting case, it will be easily adapted to multichannel images.

4.3.1 Primal-Dual approach

Computationally, the discrete Beltrami denoising model

$$\min_I \left\{ \sum_{\Omega} \sqrt{1 + \beta^2 |\nabla I|^2} + \frac{\lambda}{2} (I - I_0)^2 \right\}. \quad (4.9)$$

is usually solved by its formal Euler-Lagrange equations. This would be viewed as the primal approach to solving the Beltrami denoising problem, and I is the primal variable. Based on the results of the previous section, we can now formulate the primal-dual problem. By comparison of our Beltrami-denoising optimization problem with the primal-dual model in Section 4.2, we get the following notation:

$$F(KI) := \sum_{ij} \sqrt{1 + \beta^2 |(\nabla I)(i, j)|^2} \quad \text{and} \quad (4.10)$$

$$G(I) := \sum_{ij} \frac{\lambda}{2} (I(i, j) - I_0(i, j))^2 \quad (4.11)$$

with the reflexive spaces $V = \mathbb{R}^{m \times n}$ and $W = \mathbb{R}^{(m \times n) \times 2}$. Now, we define the functionals:

$$\begin{aligned} K: V &\rightarrow W & \text{and} & & F: W &\rightarrow \mathbb{R} \\ I &\mapsto \nabla I & & & x &\mapsto \sum_{ij} f(x(i, j)), \end{aligned} \quad (4.12)$$

with $f: \mathbb{R}^2 \rightarrow \mathbb{R}$, $f(s) = \sqrt{1 + \beta^2 |s|^2}$. The Legendre-Fenchel transform f^* of the function f is found as:

$$f^*: \mathbb{R}^2 \rightarrow \mathbb{R} \\ s \mapsto \begin{cases} -\frac{\sqrt{\beta^2 - |s|^2}}{\beta} & \text{if } \beta^2 \geq |s|^2 \\ +\infty & \text{otherwise,} \end{cases} \quad (4.13)$$

and therefore the convex conjugate F^* of F is

$$F^*: W \rightarrow \mathbb{R} \\ \varphi \mapsto \sum_{ij} f^*(\varphi(i, j)) \quad (4.14)$$

It is easy to prove that F is convex. Thus, we now get the equivalent primal-dual problem:

$$\min_I \max_{\varphi \in X} \left\{ \sum_{ij} \langle \nabla I, \varphi \rangle + \frac{\sqrt{\beta^2 - |\varphi|^2}}{\beta} + \frac{\lambda}{2} (I - I_0)^2 \right\} \quad (4.15)$$

where $X = W^* = \left\{ \varphi \in \mathbb{R}^{(m \times n) \times 2} \mid |\varphi(i, j)|^2 \leq \beta^2 \right\}$. Further, using $\nabla^* = -\text{div}$ by construction of the discretized differential operators, (4.16) can be rewritten:

$$\min_I \max_{\varphi \in X} \left\{ \sum_{ij} -I \text{div} \varphi + \frac{\sqrt{\beta^2 - |\varphi|^2}}{\beta} + \frac{\lambda}{2} (I - I_0)^2 \right\} \quad (4.16)$$

4.3.2 Algorithm

Having reformulated the Beltrami denoising problem in a primal-dual form, we can now make use of the projected gradient method for optimization. Using the notation in Section 4.2, we get for each (i, j)

$$H(I, \varphi) = \begin{pmatrix} -\text{div} \varphi + \lambda(I - I_0) \\ -\beta \nabla I \sqrt{\beta^2 - |\varphi|^2} + \varphi \end{pmatrix} \quad (4.17)$$

In order to avoid a division by zero, we have multiplied the second term of H by $\beta\sqrt{\beta^2 - |\varphi|^2}$.

Given a current estimate (I^k, φ^k) at iteration step k , the projection algorithm (4.7) then yields the following update formula (for each (i, j)):

1. $\varphi^{k+1} = P_X \left(\varphi^k - r_1 \left(-\beta \nabla I^k \sqrt{\beta^2 - |\varphi^k|^2} + \varphi^k \right) \right)$
2. $I^{k+1} = I^k - r_2 \left(-\operatorname{div} \varphi^{k+1} + \lambda (I^k - I_0) \right)$

where r_1, r_2 are positive constants and the projection on the $X = W^*$ is simply given by

$$P_X(s) = \frac{\beta s}{\max(|s|, \beta)}. \quad (4.18)$$

Indeed,

$$|P_X(s)|^2 = \begin{cases} |s|^2 & \text{if } \beta \geq |s| \\ \beta^2 & \text{otherwise} \end{cases} \quad (4.19)$$

as required. This primal-dual projected gradient algorithm for Beltrami-regularized image denoising is summarized in algorithm 3.

Algorithm 3 Beltrami primal-dual denoising

Initialize $I^0 = I_0, \varphi^0 = 0$.

repeat

$$\begin{cases} \bar{\varphi}^{k+1} \leftarrow (1 - r_1) \varphi^k + \beta r_1 \nabla I^k \sqrt{\beta^2 - |\varphi^k|^2} \\ \varphi^{k+1} \leftarrow \frac{\beta \bar{\varphi}^{k+1}}{\max(|\bar{\varphi}^{k+1}|, \beta)} \\ I^{k+1} \leftarrow (1 - \lambda r_2) I^k + r_2 (\operatorname{div} \varphi^{k+1} + \lambda I_0) \end{cases}$$

until convergence

It is easy to see that the proposed algorithm has the structure of the classical Arrow-Hurwicz-Uzawa algorithm [6]. Different strategies for choosing r are discussed in [130]. Adaptive strategies have been proposed in [49]. Here, for simplicity we'll pick $r = 0.2$ as a constant. Indeed, we believe that a detailed convergence analysis of

this particular algorithm is beyond the scope of this work. Due to the tight relations with Chambolle and Pock's first order primal-dual algorithm [28], one can expect convergence for step-sizes $r_1 r_2 \leq \rho^{-2}$, where $\rho^2 = 4$ is the spectral radius (operator norm) of the discrete Laplacian, $\Delta := \text{div} \circ \nabla$. Beyond, we note that the primal-dual approach intrinsically relates to augmented Lagrangian and split Bregman schemes [66]. The proposed algorithm therefore naturally relates to implementations based on augmented Lagrangian formulations, such as [95, 97].

4.3.3 Extension to multichannel and color images

So far, we have dealt with gray-scale image reconstruction problems, only. However, this is just a particular case, and more generally images have higher co-dimension. Such multichannel images occur for example with color, texture, multimodal or hyperspectral features. In particular, RGB-color images have co-dimension 3.

We consider a d -channel image as a function \mathbf{u} defined as

$$\begin{aligned} \mathbf{u}: \Omega \subset \mathbb{R}^n &\rightarrow \mathbb{R}^d \\ x &\mapsto \mathbf{u}(x) := (u_1(x), \dots, u_d(x)). \end{aligned} \tag{4.20}$$

Thus, d -channel images are simply a collection of d gray-scale images u_1, \dots, u_d .

To work with such images, we now need to extend our definitions of gradient, divergence and norms to two-dimensional d -channel images.

$$\begin{aligned} \nabla: \mathbb{R}^d &\rightarrow \mathbb{R}^{d \times 2} \\ \nabla \mathbf{u} &\mapsto (\nabla u_1, \dots, \nabla u_d). \end{aligned} \tag{4.21}$$

The d -channel gradient is thus a collection of the d gradients of the d respective image channels. Similarly, the d -channel divergence reduces d vector fields to d real numbers: We define the d -channel divergence operator on d -channel vector fields as in [21] by

$$\begin{aligned} \text{div}: \mathbb{R}^{d \times 2} &\rightarrow \mathbb{R}^d \\ \text{div} \mathbf{p} &:= (\text{div} p_1, \dots, \text{div} p_d). \end{aligned} \tag{4.22}$$

$$|\mathbf{u}| := \sqrt{\sum_{i \in \{R, G, B\}} |u_i|^2}$$

4.3.4 Simplified multichannel Beltrami functional

The Beltrami functional derived in differential geometry terms now leads to particular gradient coupling terms in the determinant [107, 62]. However, these gradient terms render the functional non-convex, and are thus a computational obstacle in the present context. Here, we thus formally define a simpler d -channel Beltrami energy as

$$E(\mathbf{u}) := \int_{\Omega} \sqrt{1 + \beta^2 \sum_{s=1}^d |\nabla u_s|^2}. \quad (4.23)$$

With this simplified d -channel Beltrami regularizer discretized as before, we rewrite the initial denoising model using Beltrami regularization and L^2 fidelity term w.r.t. to given data \mathbf{u}_0 :

$$\min_{\mathbf{u}} \left\{ \int_{\Omega} \sqrt{1 + \beta^2 \sum_{s=1}^d |\nabla u_s|^2} + \frac{\lambda}{2} \sum_{s=1}^d (u_s - u_{0,s})^2 \right\} \quad (4.24)$$

Without doubt, this simplified model will not be entirely as powerful as the full multichannel Beltrami functional involving cross-channel terms, as in, e.g. [95, 97].

Using the same method as in the scalar case, we can proceed to the definition of the following equivalent primal-dual variational model:

$$\min_{\mathbf{u}} \max_{\sum_{s=1}^d |\varphi_s|^2 \leq \beta^2} \left\{ \int_{\Omega} - \sum_{s=1}^d u_s \operatorname{div} \varphi_s + \frac{\sqrt{\beta^2 - \sum_{s=1}^d |\varphi_s|^2}}{\beta} + \frac{\lambda}{2} \sum_{s=1}^d (u_s - u_{0,s})^2 \right\} \quad (4.25)$$

The modified variational inequality problem can again be solved using a projection

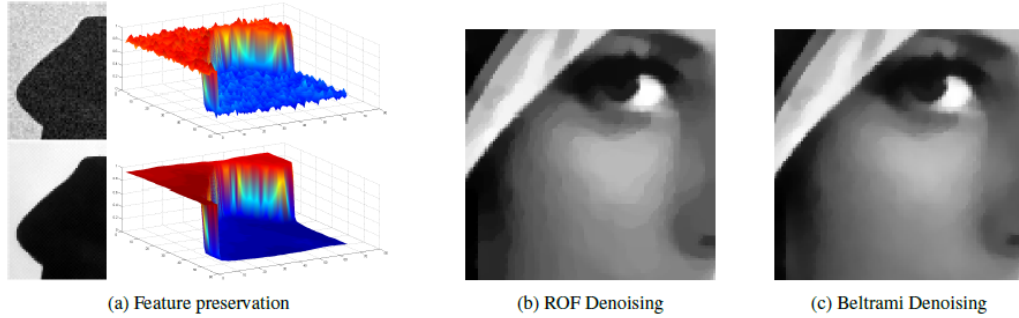


Figure 4.2: Feature preservation versus staircasing. (a) A surface plot of a section of the cameraman image before and after Beltrami denoising illustrates the feature preserving nature of the denoising model: The step edge is preserved (feature), while the little oscillations (noise) are removed. (b)–(c) The TV model is also feature preserving, but has a strong tendency to introduce piecewise constant patches separated by artificial steps (staircasing). This effect is much less prominent in the Beltrami model, where soft gradients remain smooth.

type method:

$$u_s^{k+1} = u_s^k - r_1 \left(-\operatorname{div} \varphi_s^{k+1} + \lambda (u_s^k - u_{s,0}) \right) \quad (4.26)$$

and

$$\varphi_s^{k+1} = P_X \left(\varphi_s^k - r_2 \left(-\beta \nabla u_s^k \sqrt{\beta^2 - \sum_{l=1}^d |\varphi_l|^2 + \varphi_s^k} \right) \right). \quad (4.27)$$

Here, the projection on the ball $X = \{x \mid |x|^2 \leq \beta^2\}$ is simply given by

$$P_X(p_s) = \frac{\beta p_s}{\max \left(\beta, \sqrt{\sum_{l=1}^d |p_l|^2} \right)}. \quad (4.28)$$

4.4 Denoising Results

We have implemented the proposed primal-dual projected gradient Beltrami-denoising algorithm. Here, we want to demonstrate its use, and compare our Beltrami algorithm with a simple primal-dual formulation of the ROF model [129, 130], known to be efficient. The methods are very comparable in structure. Indeed, the only difference between the proposed Beltrami denoising and the primal-dual ROF model is in

the projection and the update of the dual variable. Here, the ROF algorithm (see Algorithm 4) uses

$$\varphi^{k+1} = P_X(\varphi^k - r_1 \nabla I^k) \text{ with } P_X(s) = \frac{s}{\max(|s|, 1)}. \quad (4.29)$$

It has been shown that this primal-dual ROF algorithm converges very fast and gives a better convergence rate, compared to Chambolle's earlier projection method [27].

Algorithm 4 ROF primal-dual denoising

Initialize $I^0 = I_0, \varphi^0 = 0$.

repeat

$$\left\{ \begin{array}{l} \bar{\varphi}^{k+1} \leftarrow \varphi^k - r_1 \nabla I^k \\ \varphi^{k+1} \leftarrow \frac{\bar{\varphi}^{k+1}}{\max(|\bar{\varphi}^{k+1}|, 1)} \\ I^{k+1} \leftarrow (1 - \lambda r_2) I^k + r_2 (\operatorname{div} \varphi^{k+1} + \lambda I_0) \end{array} \right.$$

until convergence

In figure 4.3 we illustrate results for denoising a series of test images using $\lambda = 0.075$ and $\beta = 1$. The images have been degraded by additive Gaussian noise of standard deviation $\sigma = 10$ (compared to an image intensity range of $[0, 256]$). In all examples we use the same descent parameter $r = 0.2$. We choose to terminate the algorithm when the primal relative error improvement, $(P(u^{k+1}) - P(u^k)) / P(u^0)$, reaches a desired precision, here $\epsilon = 2 \cdot 10^{-5}$. Both methods perform almost equally; only subtle differences are visible in image locations with smoothly varying intensities, such as Lenna's cheek.

Indeed, this is the major difference between the two methods. While the ROF model is feature preserving, it also has a tendency to produce artificial intensity jumps (staircasing) in the output images. The Beltrami method is also feature preserving, but has greatly reduced staircasing. This difference is best appreciated in Figure 4.2.

To measure more quantitatively the similarity between the recovered image and the original (supposedly noise-free) image, we use both, the signal-to-noise ratio, and

Image	Beltrami				TV			
	SNR	SSIM	Iter.	CPU [s]	SNR	SSIM	Iter.	CPU [s]
cameraman*	25.632	0.888	33	0.624	25.574	0.886	74	1.248
circuit*	24.562	0.877	34	0.764	24.401	0.870	59	1.108
coins*	25.968	0.930	37	0.733	25.980	0.932	77	1.451
concord	23.605	0.828	35	0.546	23.531	0.823	55	0.655
liftingbody	32.106	0.951	47	4.228	31.981	0.948	92	6.505
mri	23.027	0.910	27	0.109	23.090	0.941	63	0.328
rice*	23.838	0.811	31	0.546	23.728	0.801	59	0.764
snowflakes	15.660	0.921	41	0.328	15.518	0.917	80	0.530
text	24.574	0.901	33	0.515	24.913	0.950	70	0.967
tire	22.880	0.916	41	0.406	22.803	0.914	67	0.593
bag	21.984	0.912	31	0.296	21.938	0.910	49	0.484
boat*	26.291	0.942	37	3.089	26.174	0.939	69	5.132
lenna*	25.084	0.948	40	3.448	24.940	0.945	78	5.710
lincoln*	28.519	0.845	46	2.012	28.306	0.834	78	2.839

Table 4.1: Beltrami and TV denoising applied to a larger set of test images. Asterisk denotes images shown in figure 4.3. Signal-to-noise ratio and structural similarity Index are in favor of the Beltrami model (bold). Also, the Beltrami model converges in a little more than half the iterations required for the TV model. Each iteration is only slightly more computationally complex than the TV algorithm, since both implementations are based on a very similar primal-dual projected gradient method and thus structurally identical. Convergence criterion is based on the decrease of the primal problem energy.



Figure 4.3: Beltrami and TV denoising applied to a larger set of test images. Asterisk denotes images shown in figure 4.3. Signal-to-noise ratio and structural similarity Index are in favor of the Beltrami model (bold). Also, the Beltrami model converges in a little more than half the iterations required for the TV model. Each iteration is only slightly more computationally complex than the TV algorithm, since both implementations are based on a very similar primal-dual projected gradient method and thus structurally identical. Convergence criterion is based on the decrease of the primal problem energy.

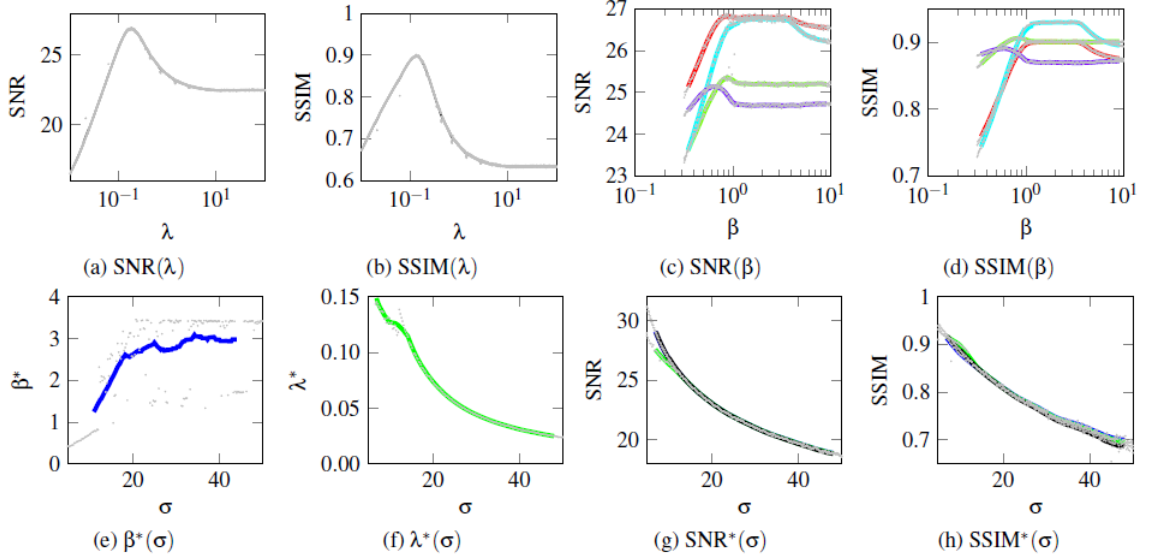


Figure 4.4: Parameter choice and sensitivity. (a)–(b) For a given image and a fixed noise level, the balancing parameter λ has a prominent peak in the resulting signal-to-noise ratio (SNR) and the structural similarity index (SSIM). (c)–(d) The optimal aspect ratio β is image dependent, but has also exhibits a flatter plateau of near-optimality. (e)–(f) the locations of the optimal parameters β^* and λ^* as determined by numerical optimization. For significant noise levels ($\sigma > 10$), the optimal β is almost noise independent, and exhibits large variation. Due to the flat plateau, almost any value between 1 and 3.5 is “good”. In contrast, the optimal balancing parameter λ is well-defined by an inverse proportional relationship with the noise-level. (g)–(h) SNR/SSIM at SNR-optimal parameters β^* and λ^* . (blue: optimize both β and λ , green: keep $\beta = 2$ and only optimize λ , black: TV with optimal λ). Optimally tuned Beltrami model slightly outperforms TV denoising.

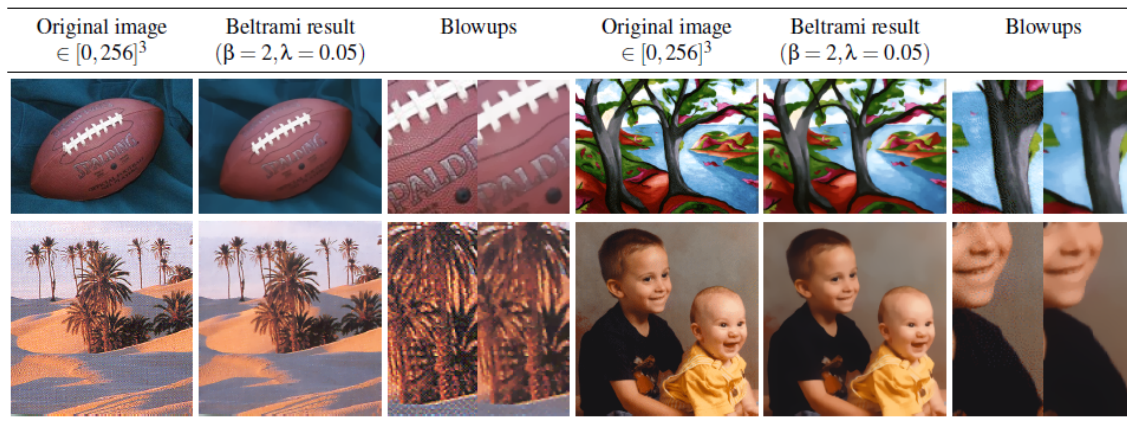


Figure 4.5: Simplified Color-Beltrami denoising results. 4 color test images with strong inherent noise (object texture, half toning). The simplified color-Beltrami model can be seen to effectively remove color granularity (noise, texture), while preserving features such as lines and edges. Color gradients are smooth and do not exhibit artificial staircasing.

the structural similarity index (SSIM), the latter being well known to better reflect perceived visual quality [119]. Comparative numbers are reported for both methods and over a series of 16 test images, in Table 4.1. It can be seen that the quality difference between ROF and Beltrami is not extreme, but almost consistently in favor of the proposed Beltrami model. It is also to note that given the primal-energy decrease criterion, the Beltrami model converges in roughly half as many iterations as the TV model.

While the ROF model has only one balancing parameter, which is essentially governed by the noise level, the Beltrami model involves the aspect ratio β , in addition. We study the influence of parameter selection (β and λ) in terms of SNR and SSIM. The curves are shown in Figure 4.4(a)–(d). In addition, we have performed numerical optimization to search for optimal β^* and λ^* over a broad range of noise levels, σ , see Figure 4.4(e)–(f). The variation on β^* is large, and above a certain noise level the optimal choice seems largely independent of σ . In contrast, the optimal balancing parameter λ is clearly anti-proportional to the noise-level. In conclusion, as with the ROF model, the balancing parameter λ is governed by the noise level, is largely image independent and exhibits a prominent peak. Conversely, the aspect ratio β is mostly

image dependent, and typically has a broader plateau. In a typical application, the balancing parameter λ can be inferred from the estimated noise level, while the optimal β should be learned from a set of representative training images. In our case, the curves suggest optimal values in a typical range of 1 – 3 (and less in low-noise regimes). Finally, we have traced the SNR and SSIM achieved at optimal parameter choices, for both the Beltrami and the ROF model, see Figure 4.4(g)–(h). The curves suggest that the Beltrami model typically outperforms ROF in both SNR and SSIM, albeit only marginally so.

We demonstrate the denoising power of the proposed simplified color-Beltrami model on 4 different color images. Images and denoising results are shown in Figure 4.5. We have chosen parameters such that the denoising is rather excessive, in order to better highlight the nature of the denoising. The simplified color-Beltrami model can be seen to effectively remove color granularity (noise, texture), while preserving features such as lines and edges. Color gradients are smooth and do not exhibit artificial staircasing.

4.5 Conclusions

In this study, we have proposed a primal-dual projected gradient method to efficiently solve the Beltrami regularization based image denoising. This algorithm is inspired by very successful schemes used with the ROF model, proposed by Zhu and Chan in [129, 130]. Beltrami regularization offers an interesting compromise between the smooth regularization of H^1 priors, and the feature preservation of TV regularization. Until now, this compromise came at a cost in the form of increased computational complexity and was thus rarely used in practice. Here, we present a primal-dual Beltrami-based denoising algorithm and compare its performance to the corresponding TV-based algorithm. We can show that increase denoising quality can be achieved at comparable computational complexity. Beyond the simple denoising problem, the primal-dual projected gradient Beltrami regularization algorithm can be extended to a larger family of inverse image processing problems, such as in-painting, deconvolution and compressed sensing. All these cases can be tackled in a very similar way.

Future work will also focus on convex relaxations of the true multichannel-Beltrami model to not have to make such simplifications.

Chapter 5

Myocardial T_1 mapping denoising¹

¹Submitted to: Bustin A, Ferry P, Codreanu A, Beaumont M, Liu S, Burschka D, Felblinger J, Brau A, Menini A, Odille F. Impact of a Vectorial Anisotropic Denoising Technique on Precision and Accuracy of Myocardial T_1 Mapping. *Journal of Magnetic Resonance Imaging*, 2017, and presented in part at the 24th Annual Meeting of ISMRM. Singapore; O587, 2016.

5.1 Introduction

Quantitative Magnetic Resonance Imaging (QMRI) is a valuable tool that intrinsically studies tissue properties in the human body, providing additional information for clinical diagnosis and treatment decisions [20]. Quantitative MR as assessed by parameter mapping, can measure changes of properties such as the longitudinal relaxation time (T_1), transverse relaxation time (T_2) and proton density (PD). Quantitative parameter mapping requires the acquisition of multiple measurements with different acquisition parameters, e.g. inversion time (TI), saturation time (TS), echo time (TE) or flip angle (FA), and often results in long scan times, thereby limiting its use as a clinical application in some situations.

Several approaches have been proposed to overcome this limitation, including parallel imaging [93], compressed sensing [39, 57] and sparsity-based reconstructions [127, 126, 128], providing a way to enable the acquisitions to be performed in a single breath hold, e.g. in myocardium mapping [123, 78], and few minutes, e.g. in neuroimaging [120]. Acquiring multiple images with different contrast has been made possible in cardiac magnetic resonance (CMR) with single-shot imaging and ECG-gated acquisitions [78, 30, 106]. However, such benefits often come at the cost of reduced signal-to-noise ratio (SNR), lower spatial resolution and more generally increase of noise in the collected images.

Among all cardiac mapping sequences, inversion-recovery (IR) techniques for T_1 assessment, such as MOLLI [79], have received most clinical interest due to their high precision (i.e. reproducibility) and satisfactory SNR [59]. Despite its wide acceptance, MOLLI suffers from the underestimation of myocardial T_1 values, leading to apparent T_1^* values. On the other hand, saturation-recovery (SR) based myocardial T_1 mapping techniques, such as SASHA [30] or SMART T_1 Map [106], have shown better accuracy as they acquire only one image after each saturation pulse and are independent of the previous spin history [108]. However, those techniques have lower precision of T_1 values compared to IR-based sequences. The low precision can partly be explained by the low dynamic range of the relaxation curve, making SR acquisitions more sensitive to acquisition noise [59].

Additionally, appropriate choice of the fitting parameter is of utmost importance and a trade-off between accuracy and precision has to be done when considering whether two or three-parameter fitting is necessary. In particular, three-parameter fitting of the form is advised in some situations, especially at 3T, where the impact of B_1 inhomogeneities must not be neglected. SR-based methods have been shown to provide excellent precision when combined with a two-parameter fitting, at the expense of a loss of accuracy. In contrast, a three-parameter fitting has shown highly accurate T_1 values but was more sensitive to noise, greatly reducing T_1 values precision [59, 61].

The inherent noise penalty in myocardial T_1 mapping, directly related to the mapping sequence and the choice of parameter fitting, can be minimized using reliable denoising algorithms. Popular regularization approaches from the image processing literature, such as Tikhonov [112], Total Variation (TV) [100] or Beltrami [63] can be applied in a straight-forward manner for image restoration and noise removal. However, as previously reported in diffusion-weighted imaging [54, 68], such techniques need to be adapted to exploit the natural redundancy and correlation observed in the parameter dimension.

In this study, we propose to further improve the precision of myocardial T_1 mapping techniques by employing a novel denoising method which exploits the spatio-temporal correlations in the space of T_1 -weighted images while preserving sharp anatomical details. Validation was performed by investigating the impact of the proposed denoising method on accuracy and precision of myocardial T_1 mapping on simulation, ex vivo and in vivo experiments.

5.2 Materials and Methods

5.2.1 Single-Frame Beltrami Denoising

To describe a typical MR T_1 mapping acquisition, we consider a discrete dataset S comprised of n images, representing the T_1 -weighted images. In magnetic resonance imaging, denoising techniques aim to solve the following inverse problem: find the

underlying images ρ free of noise, given derived measurements S , affected by noise ϵ : $S = \rho + \epsilon$, where ρ and $S \in \mathbb{R}^{n_x \times n_y \times n}$ with $(n_x \times n_y)$ the size of each frame and ϵ represents the noise in the measurements, which follows a Rician distribution in the magnitude image [51]. It should be noted that the bias introduced by noise with non-zero mean - as it is the case with Rician noise - can be dealt with in the curve fitting stage (see Section 5.2.5 hereafter). The denoising discussed here is a preprocessing stage to reduce the standard variation of noise by exploiting the structure of the images. The aim of the method is to reconstruct the images ρ free of noise, from the acquired data $S = (S_k)_{k=1, \dots, n}$. Images ρ , can be found by solving the following regularized least squares optimization problem frame-by-frame:

$$\hat{\rho}_i = \underset{\rho_i}{\operatorname{argmin}} \left\{ \|\rho_i - S_i\|_2^2 + \lambda R(\rho_i) \right\} \quad (5.1)$$

Here R is the chosen regularization function and λ is the corresponding nonnegative regularization parameter. The least square term in (5.1) forces the reconstructed image $\hat{\rho}_i$ to preserve a degree of fidelity with respect to the acquired data, whereas the regularization term discourages aberrant values from appearing in the denoised image. Previously studied regularization schemes include: Tikhonov, typically defined by $R_{TK}(\rho_i) = \|\rho_i\|^2$; edge-preserving regularizers such as TV regularization, where $R_{TV}(\rho_i) = \|\nabla \rho_i\|$, or Beltrami, defined as $R_{BL}(\rho_i) = \sqrt{1 + \beta^2 \|\nabla \rho_i\|^2}$ where β is the Beltrami constant, which offers a compromise between the smooth regularization of priors and the feature preservation of TV regularization. In Chapter 4, we showed that the recently proposed Beltrami regularization is able to maintain the advantage of TV, i.e. edges preserving and noise reduction, as well as reducing the effect of staircasing usually observed in the reconstructed image.

We propose to take advantage of the later regularization by adding two main contributions. First, we extend the problem to a multi-contrast Beltrami regularization by introducing a coupling between the T_1 -weighted images, thus improving the denoising model by penalizing across a common edge (T_1 encoding direction) for all samples. We further improve the model by replacing the isotropic regularization by an anisotropic term that considers the directional structure within the images.

5.2.2 Vectorized Beltrami Denoising

The classic Beltrami functional was previously defined as a simple 1-channel energy which is applied independently on each frame without exploiting correlations in the spatio-parameter space. Considering this, we define the vectorized Beltrami regularization for multi-dimension images as $R_{BL-V}(x) = \sqrt{1 + \beta^2 \sum_{i=1}^n |\nabla x_i|^2}$. As in the 1-channel case, this vectorial definition of Beltrami allows discontinuities in each frame by preserving edges in a common direction. Indeed, common structures in different frames will result in a stronger weighting and therefore a better edge preservation, while local intensity variation specific to each channel will be treated as noise. Therefore, edges/details in the images will be better enhanced if they are present in multiple frames. Note that $n = 1$ implies $R_{BL-V} \equiv R_{BL}$, i.e. we recover the classic Beltrami norm for a 1-channel image.

5.2.3 Anisotropic Beltrami

We further improve the model by replacing the isotropic regularization by an anisotropic term that considers the directional structure within the images. Hereby, the regularization is set to $R_{BL-V-ANI} = \sqrt{1 + \beta^2 \sum_{i=1}^n |\nabla^w x_i|^2}$. This new formulation of the vectorial Beltrami regularization includes a weighting function w , inspired by the recent non-local means filtering (25), that amplifies gradient information, and can be seen as an edge-aware pre-conditioner. We define the weighting gradient at the i^{th} iteration using the following notation: $\nabla^w x_i^{(t)} = w(x_i^{(t-1)}) \nabla x_i^{(t)}$ with the corresponding weight defined as $w(x_i^{(t-1)}) = e^{-\frac{(\nabla x_i^{(t-1)})^2}{h^2}}$. The geometry of structures and directional knowledge in the images at the previous iteration are thus incorporated into the denoising problem in order to decrease the local smoothing in the directions where sharp transitions are observed. From the above expression, we see that if a pixel is located on an edge curve, then the weight w will be small because the gradient ∇x_i will be relatively large and therefore major edge features will be preserved. The parameter h controls the strength of the edge amplification and therefore impact the resolution of the reconstructed map. The lower the parameter h is, the sharper the

denoised map is. On the other hand, a high value of h almost leads to the result of the vectorized isotropic Beltrami. The final vectorized anisotropic Beltrami denoising problem can be expressed as the following optimization problem:

$$\hat{\rho} = \underset{\rho}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \|\rho_i - S_i\|_2^2 + \lambda \sqrt{1 + \beta \sum_{i=1}^n g(|\nabla^w \rho_i|^2)} \right\} \quad (5.2)$$

where β is the Beltrami constant and λ is the regularization parameter controlling the desired noise reduction. ∇^w is the 2D weighted-gradient transformation applied to the 2D multi-frame cardiac T_1 -weighted images. The summation in the Beltrami norm now provides edge-preserving smoothing and coupling between the parameter-weighted images. The weighting term involves a direction smoothness constraint on the solution, by penalizing the gradient depending on its local orientation. This scheme is equivalent to the vectorial isotropic algorithm on homogenous areas (small gradient) but focuses on reducing blur and preserving local structures (e.g. small vessels, myocardial walls, papillary muscles and more generally on edges and corners) where the gradient is large.

5.2.4 Denoising Implementation

Denoised images ρ are reconstructed by solving the optimization in Eq. (5.2) using a recently introduced primal-dual projected gradient approach [129]. The fundamental idea is to formulate an equivalent optimization problem involving a dual variable to the primal variable (images ρ), and where the Beltrami-regularizer is replaced by the convex conjugate of this dual variable. The algorithm is simple to implement and has shown to be efficient, in terms of speed and robustness (see 4). The regularization weight λ was optimized in the simulation study to achieve the best SNR, and was set to be 0.05 for all in vivo and ex vivo experiments. This parameter was empirically estimated in simulation and MR experimental studies to balance between good SNR, smoothness and reconstruction fidelity of the recovered map. A weight of 1.0 was employed for the Beltrami constant as indicated. The sharpness of the reconstructed map was found to be optimal when $h = 18$. The same value was used for simulation,

ex vivo and in vivo experiments.

5.2.5 T_1 Map Reconstruction

T_1 maps are then generated by applying a three-parameter least-squares curve-fitting to the magnitude signal intensity using a Levenberg-Marquardt fitting algorithm [76]: $\rho(t) = A - Be^{-\frac{t}{T_1}}$ where ρ is the signal intensity, t the saturation time and T_1 the longitudinal relaxation time. As opposed to a two-parameter model of the form $\rho(t) = A \left(1 - e^{-\frac{t}{T_1}}\right)$, imperfect saturation due to B_1 transmit field inhomogeneities is intrinsically compensated. Additionally, the systematic bias introduced by Rician noise (because it has nonzero mean) results in a small constant overestimation of each image of the series. In the three-parameter model this constant term will be reflected in the A and B coefficients and will not affect T_1 , while in the two-parameter model A and T_1 would be affected.

5.2.6 Imaging

All imaging studies were performed on a 3T GE Signa HDxt system (GE Healthcare, Milwaukee, WI). T_1 mapping was performed using the SR-based sequence called SMART1Map [106]. Cardiac studies used an eight-channel phased array coil. The human studies (patients and injected volunteers) were approved by the local Institutional Review Board. Written informed consent was obtained from all volunteers and patients before study participation. Cardiac images from patients and volunteers were collected from the scanner in DICOM format and were loaded into a computer for offline image processing and analysis. Simulations, image processing, and statistical analysis were implemented in MATLAB (MathWorks, Natick, Massachusetts, USA). When needed, respiratory induced motion, due to imperfect breath holding, was corrected using a nonrigid image registration framework which was validated for T_2 mapping in a large patient database [81]. The registration is applied prior to the denoising step.

5.2.7 Simulations

Simulations were performed to assess the impact of the proposed T_1 mapping denoising method on precision and accuracy of the T_1 values. A synthetic cardiac phantom, containing straight and curved edges, was designed to represent a post-contrast short-axis acquisition, consisting of a left and right ventricle ($T_1 = 500$ ms), and blood pool ($T_1 = 230$ ms), representing typical T_1 values seen in clinical practice. A thin blood structure was added to simulate the left anterior descending artery with T_1 equal to 230 ms. Eight images were simulated with the numerical phantom using the three-parameter SR equation. Saturation times were chosen equal to [250, 420, 590, 1590, 2590, 3590, 4590] ms, plus an infinite delay time T_∞ corresponding to the acquisition when no saturation pulse is applied. Prior to T_1 map estimations, random Rician noise was generated and added to each TS image at a level of $\sigma = (.1, 3, 5, 7, 9)\%$ the maximum intensity. For the given noise, the denoising parameter λ in the algorithm was optimized to achieve the highest SNR. In all experiments, the proposed BL-V-ANI denoising was compared to that obtained with conventional BL and BL-V denoising as defined in the previous section. Additionally, we compare the proposed technique with the state-of-the-art frame-by-frame NLM denoising recently adapted to Rician noise [114], which reduces noise by exploiting the redundancy of information within the image itself. The noise variance, in the NLM algorithm, was tuned to achieve the best SNR while keeping the other default parameters fixed. Results were compared quantitatively with the T_1 values obtained from the ground truth. The impact of the denoising on image quality and SNR was studied, and optimal results were compared to the ground-truth T_1 map, free of noise. Several measures were used to quantify the performances of the method. The first one is the intensity root mean square error (RMSE) defined as:

$$RMSE(T_1^{ref}, T_1^{den}) = \sqrt{\frac{1}{n_x \times n_y} \|T_1^{ref} - T_1^{den}\|_2^2} \quad (5.3)$$

where $\|\cdot\|_2$ stands for the usual l_2 norm. Additionally, the peak-signal-to-noise ratio (PSNR) between the reference T_1 map and the denoised map, evaluated in decibels, was computed as:

$$PSNR(T_1^{ref}, T_1^{den}) = 20 \log_{10} \frac{\max(T_1^{den})}{RMSE(T_1^{ref}, T_1^{den})} \quad (5.4)$$

Besides PSNR, which can accurately measure the local intensity difference between two volumes, we also employed another image quality measurement based on image sharpness, called Sharpness Index (SI) [18]. SI has been shown to provide a sensitivity to image sharpness with crucial information about the image geometry and image edges and therefore provides an interesting indicator of image sharpness:

$$SI(T_1^{den}) = -\log_{10} \phi \left(\frac{m - TV(T_1^{den})}{v} \right) \quad (5.5)$$

where $m = \mathbb{E}[TV(T_1^{den})]$ is the expectation of the total variation of the denoised T_1 map, and $v = \text{Var}[TV(T_1^{den})]$ is the corresponding variance. This index achieves high values when the reconstructed map exhibits sharp and enhanced edges.

5.2.8 Ex Vivo Imaging

To evaluate the performance of the proposed T_1 mapping denoising technique, an ex vivo human heart was imaged with the SMART $_1$ Map sequence [105] with the following parameters: FOV = 220 x 220 mm $_2$, slice thickness = 8 mm, flip angle = 45 degrees, TR = 3.74 ms, TE = 1.63 ms, acquiring eight samples on the T_1 recovery curve at $TS = [250, 420, 590, 1590, 2590, 3590, 4590, T_\infty]$ ms. The acquisition was performed using a simulated ECG signal with a heart rate of 100 bpm. Reference T_1 values of the ex vivo heart were determined by repeating the k-space acquisition by setting up eight number of excitations (NEX). T_1 values were assessed within the myocardium using a three-parameter fit. The same region of interest (ROI) was used for all sequences. The average T_1 estimation was compared between the Multi-NEX gold standard and the anisotropic denoised map. Precision was assessed as the standard deviation (SD) in the left ventricle for the given T_1 map. The correlations between the uncorrected, corrected and reference T_1 maps were studied using Bland-Altman plots, and linear regression analysis was performed against the Multi-NEX reference.

5.2.9 In Vivo Imaging

A cohort of sixteen ($N = 16$) healthy volunteers (12 males, average age 39 ± 8 years, 62 ± 9 bpm) were included in the study over a period of two months. Volunteers were recruited for the study according to several inclusion criteria: no competitive-level sport activity, no allergy history, no exertion symptoms, normal physical exam and normal rest electrocardiogram.

Additionally, imaging was performed in one patient with stable ischemic heart disease and history of old (> 2 years) myocardial infarction, referred for MRI examination with Gadolinium-based contrast injection. Similar imaging parameters were used as in the previous section except $TS = [100, 220, 340, 460, 1744, 3034, 4250, T_\infty]$ ms (typical values, depending on heart rate) with an acquisition matrix of 224×192 and slice thickness of 8 mm. Both the T_1 mapping sequence and LGE technique were applied between 0 and 15 minutes after gadolinium-based contrast injection to evaluate the feasibility of the denoising on a wide range of T_1 values. Images corresponding to all TS were acquired in a single breath hold (12-28 s depending on heart rate) with prospective ECG gating.

The image sets were subsequently used to generate T_1 maps with and without the proposed anisotropic correction. Segmentation was conducted using dedicated MR Analytical Software System (MASS, Medis, Leiden, the Netherlands). For quantitative analysis of T_1 values, endo- and epicardial contours were drawn manually by an experienced reader on each T_1 map separately. ROIs in the myocardium and left-ventricle blood pool were analyzed. Mean T_1 value and SD of T_1 value were extracted and compared between the different techniques. Statistical analysis was based on Student t-tests with statistical significance for p-values of less than 0.05.

Table 5.1: Recommended 2D cardiac T_1 Mapping using SMART₁Map on a GE system.

Patient Position	
Patient position	supine
Patient entry	feet first
Coil	8-channel cardiac coil
Series description	2D SMART ₁ Map
Scanning Range	
Scan plan	Oblique
Frequency FOV	34-40 cm
Phase FOV	0.75
Slice thickness	5-8 mm
Spacing	to select
Imaging Parameters	
Plane	Oblique
Mode	2D
Family	Gradient echo
Pulse sequence	FIESTA
Imaging options	Cardiac gating, IR prep
Details	
TE	Minimum Full
Flip angle	≥ 45
Prep time	250-300 (Non-contrast) 100 (Post contrast, approx. 10 min)
NEX	0.5
Frequency	160-256
Phase	128
Bandwidth	125
Shim	auto
Cardiac Gating	
Trigger window	10
Trigger delay	End-systole or mid-diastole
View per segment	64

5.3 Results

All T_1 mapping tests were performed on a desktop computer (2.0 Ghz Intel Core with 8 GB RAM). A fixed number of 50 iterations were used for the anisotropic denoising method. The algorithm was found to converge to this number of iterations in all experiments.

5.3.1 Simulations

Results of simulations with the synthetic phantom for a noise level of $\sigma = 5\%$ are shown in Figure 5.1. Interestingly, NLM shows satisfactory noise reduction but granular-like noise artifacts can be seen all over the image on flat areas. As expected, classic Beltrami (BL) efficiently reduces the noise while preserving edges, but showed artifacts particularly seen on small transitions and flat areas. However, this effect was clearly reduced on the vectorial isotropic Beltrami (BL-V). Qualitatively, the anisotropic method (BL-V-ANI) was seen to work as good as BL-V on flat areas (e.g. blood pool, myocardium) but particularly better on horizontal and vertical edges (e.g. papillary muscles and vessels, see arrows) with an overall map close to the reference T_1 map. The proposed denoising technique was able to minimize the noise and to reveal sharp details and structures.

Table 8.1 summarizes the PSNR (P), RMSE (R) and sharpness (S) results for the four denoising methods (NLM, BL, BL-V and BL-V-ANI) with different levels of noise. These results show that the proposed denoising outperformed the existing techniques in terms of PSNR, RMSE and sharpness. For high level of noise ($\sigma = 9\%$), the average PSNR of the proposed BL-V-ANI technique was 33.55 dB, showing improvements of about 9 dB, 5 dB and 3 dB over NLM, BL and BL-V respectively. Overall, the sharpness, as expressed by the sharpness index, showed good results with the proposed denoising and correlated well with the visual human perception (Figure 5.1). Compared to NLM, the proposed denoising technique consistently produced smaller RMSE for all noise levels.

The accuracy and precision of the four methods are compared in table 8.2. Myocardial T_1 values over the left ventricle obtained with BL-V-ANI (498 ± 5 ms), BL-V




Figure 5.1: Denoised T_1 maps for the simulation study by procedures NLM (b), BL (c), BL-V (d) and BL-V-ANI (e). Comparison with uncorrected (a) and gold standard (f) maps are shown. Close-up views and residual images are shown for each reconstruction. The conventional BL-V method considerably reduces the noise in the resulting T_1 map but fails to preserve edges on thin walls. The T_1 map reconstructed with the proposed BL-V-ANI denoising remains free from noise and provides better quality on thin structure, with a slight improvement at edges (arrows).

Table 5.2: Comparison of PSNR (P), RMSE (R) and Sharpness index (S) for the denoising of the numerical phantom (Simulation study) with different values of noise.

		Noise %				
		$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 7$	$\sigma = 9$
Noisy	P	33.87	24.24	19.45	15.76	11.18
	R	10.12	30.69	53.25	81.48	138.08
	S	0.66	0.16	0.06	0.05	0.14
NLM	P	45.63	35.96	31.19	27.76	24.61
	R	2.61	7.96	13.79	20.47	29.42
	S	1.02	0.91	0.82	0.62	0.48
BL	P	43.13	35.86	32.34	29.81	28.58
	R	3.49	8.05	12.08	16.15	18.61
	S	0.92	0.79	0.76	0.67	0.75
BL-V	P	45.85	38.16	34.50	31.95	30.30
	R	2.55	6.18	9.42	12.64	15.27
	S	0.93	0.82	0.79	0.72	0.78
BL-V-ANI	P	46.96	40.95	38.58	35.40	33.55
	R	2.24	4.48	5.89	8.49	10.50
	S	0.99	0.97	0.98	0.98	1.14

(499 ± 5 ms) and BL (499 ± 7 ms) were close to the reference T_1 map (500 ms) with significant improvement in the precision compared to the uncorrected map (507 ± 64 ms) and NLM (489 ± 14 ms). Mean T_1 values from BL-V-ANI over thin structures, e.g. right ventricle, were lower than the noisy T_1 map (mean 490 ms versus 507 ms, respectively), with a precision dropping from 64 ms to 7 ms. However, this minor trend toward shorter T_1 values, on thin structures, was less than 2% of the mean difference with respect to the gold standard. Overall, similar quantitative results were obtained with BL, BL-V and BL-V-ANI, but better precision of T_1 values were achieved on thin structures. The total average computation time over 100 trials for a given dataset of size 256×256 with eight TS images was 1.9 secs using a Matlab implementation of the proposed BL-V-ANI algorithm, which was significantly faster than NLM (6.2 secs) and slightly higher than the classic BL (1.6 secs) and BL-V (1.8 secs).

Table 5.3: Comparison of denoising time, accuracy (μ) and precision (SD) of the reconstructed numerical phantom map. Values are reported as the average over all noise levels.

	T ₁ values (ms)		
	Left Ventricle (T ₁ = 500 ms)	Right Ventricle (T ₁ = 500 ms)	Blood Pool (T ₁ = 230 ms)
Noisy	507 ± 64	507 ± 64	239 ± 62
NLM	489 ± 14	489 ± 17	229 ± 10
BL	499 ± 7	484 ± 16	230 ± 5
BL-V	499 ± 5	492 ± 13	231 ± 3
Noisy	498 ± 5	490 ± 7	230 ± 3

5.3.2 Ex Vivo Imaging

T₁ maps reconstructed using uncorrected and corrected SMART₁Map are shown in Figure 5.2a. Good agreement can be observed qualitatively between the reference Multi-NEX map and the corrected map with sharp results and faithful restoration of edges without introducing blurring and staircasing. Quantitatively, denoised maps provided better T₁ agreement with the reference Multi-NEX method with a slope of $s = 0.8$ (T₁ [denoised] = $s \times$ T₁ [gold standard]) as seen in Figure 5.2b-c, respectively equal to 0.6 for the uncorrected map.

Additionally, standard deviation over the left ventricle, observed on the Bland-Altman plots, confirmed the good agreement between reference and corrected map with a decreased of 37.5% (5.2d-e), showing improvement in T₁ values precision. For the uncorrected map, the Bland-Altman plot showed that the lower and upper 95% limits of agreement (-160 ms / 160 ms) contains 95.7% (506 / 529) of the difference (gold standard - uncorrected T₁ values). Correspondingly, the lower and upper 95% limits of agreement for the corrected map (-96 ms / 100 ms) contains 95.5% (505 / 529) of the difference (gold standard - corrected T₁ values).

5.3.3 In Vivo Imaging

Volunteer Studies

Exemplary T₁ maps acquired in a healthy volunteer with the standard SMART₁Map

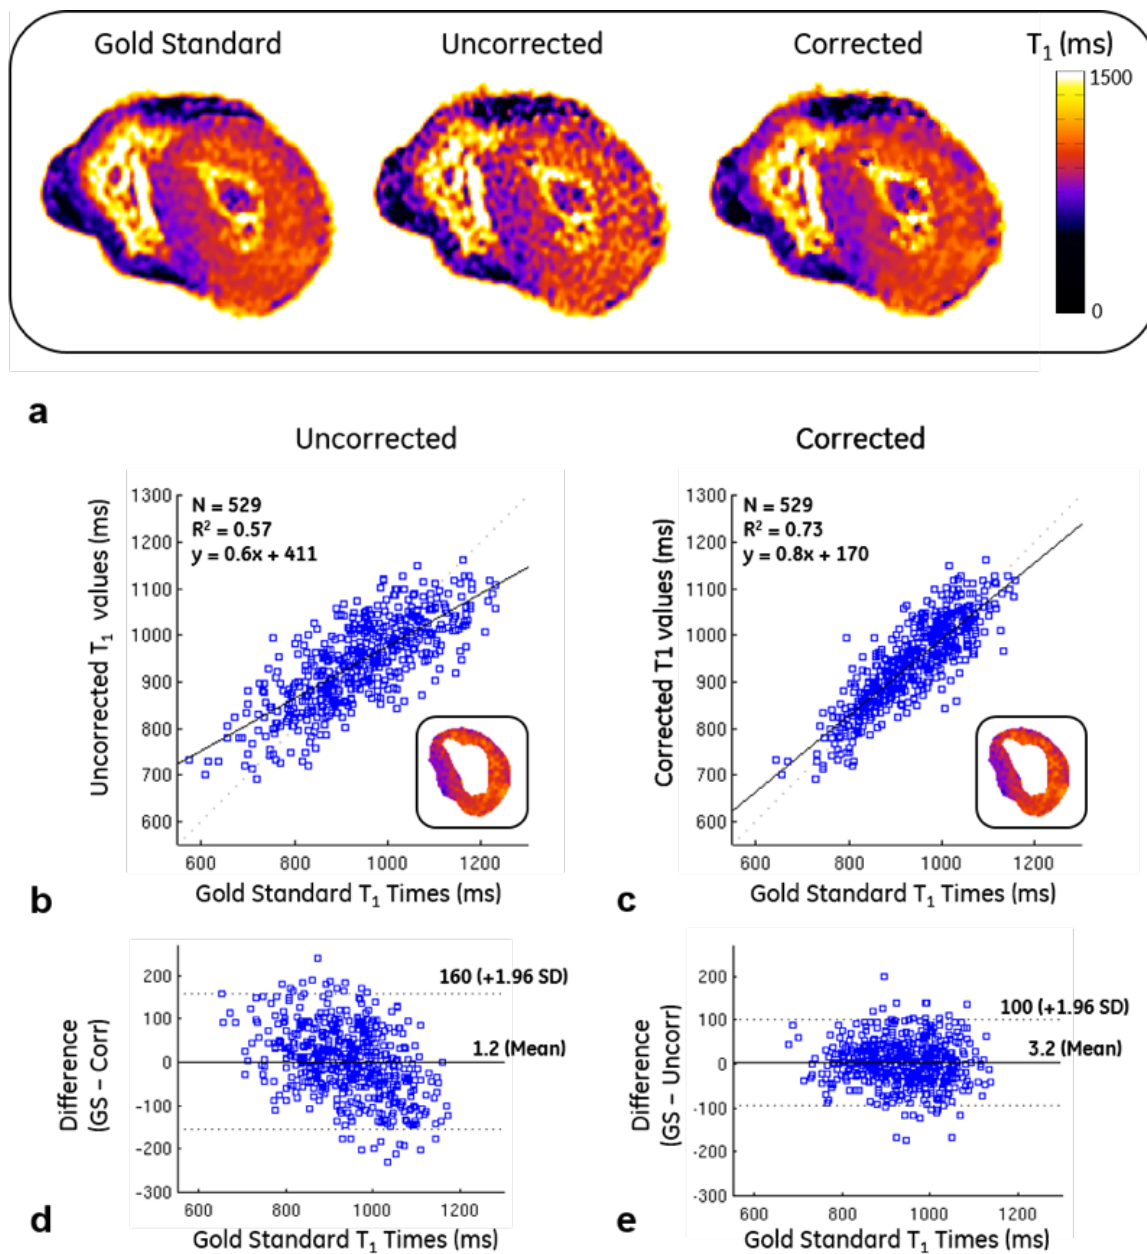


Figure 5.2: Ex vivo reconstructions and related Bland-Altman plots for the uncorrected and corrected maps. a: T_1 maps comparison on an ex vivo heart with SMART $_1$ Map, corrected SMART $_1$ Map and Multi-NEX reference. b-c: Linear regression analysis of the corrected T_1 map (y-axis) with respective gold standard derived from the Multi-NEX acquisition (x-axis). The dotted blue line represents the ideal correlation. d-e: Bland-Altman plots of the difference between corrected and theoretical T_1 values measured across the left ventricle with the two methods. The middle black line represents the mean of the difference in T_1 values. The top and bottom dotted lines represent the 95% confidence interval. The proposed method preserves accuracy of the original SMART $_1$ Map sequence and indicates better precision.




Figure 5.3: An example of the T_1 maps of a 51-year old healthy volunteer with a heart rate of 53 bpm, before (top) and after (bottom) denoising, using a three-parameter fitting. Better overall quality and more homogenous map (e.g. myocardial and blood) are seen with the proposed method. The anisotropic feature has the advantage to suppress noise while preserving sharp edges and thin structures, even for low contrast images (e.g. $TS = 100$ ms).

and our denoising method are shown in Figure 5.3. In addition to the T_1 maps, the figure displays the acquired T_1 -weighted images at three different TSs ($TS = 100$ ms, $TS = 460$ ms and $TS = T_\infty$ ms) before and after denoising. Noise observed among the T_1 -weighted images was considerably reduced by the proposed BL-V-ANI denoising, with well-preserved edges and good depiction of local structures (e.g. right ventricular wall, papillary muscles, ...), without introducing staircasing artifacts.

Correspondingly, the mean and SD of T_1 values measured in the myocardium of 16 healthy volunteers are shown in table 5.4. The average myocardial T_1 values without/with the proposed denoising method were 1433 ms / 1426 ms (pre-contrast) and 766 ms / 759 ms (post-contrast). There was no statistical difference ($p > 0.05$) between these mean values, suggesting no significant impact on accuracy. Conversely, the SD of T_1 values within the myocardium without/with the proposed denoising method were 223 ms / 151 ms (pre-contrast) and 176 ms / 135 ms (post-contrast). These differences were statistically significant ($p < 0.05$) suggesting an improved

Table 5.4: Mean and standard deviation (SD) of T_1 values within the myocardium ROI in the healthy volunteer population (N=16) using SMART $_1$ Map with and without the proposed denoising technique.

	Mean T_1 (ms)		SD T_1 (ms)	
	Uncorrected	Corrected	Uncorrected	Corrected
Pre-contrast	1433 \pm 67	1426 \pm 67	223 \pm 73	151 \pm 47
Post-contrast	766 \pm 95	759 \pm 96	176 \pm 58	135 \pm 40

Differences with statistical significance are identified by $*p < 0.05$.

precision.

Representative T_1 maps on one additional healthy volunteer before and after denoising are shown in Figure 5.4. Compared with the T_1 map at baseline (uncorrected), better image quality over the whole heart was observed with the three denoising techniques (BL, BL-V and BL-V-ANI). The proposed BL-V-ANI technique, however, exhibited a sharper T_1 map with improvement at edges (e.g. epi- and endocardium), and better visualization of cardiac features (e.g. right ventricular wall, vessels and papillary muscles).




Figure 5.4: Example of a noisy T_1 map and the denoised maps resulting from NLM, BL, BL-V and BL-V-ANI denoising techniques on a patient on a mid-ventricular short axis view. The data was acquired 5 minutes after Gadolinium injection. Applying the vectorized correction (BL-V) clearly improves the map quality. The anisotropic correction (BL-V-ANI) exhibits sharper details and substantial improvement at edges (e.g. epi- and endocardium) and thin structures (e.g. papillary muscles).

Patient Studies

The same qualitative observation can be done in the patient acquisition as depicted

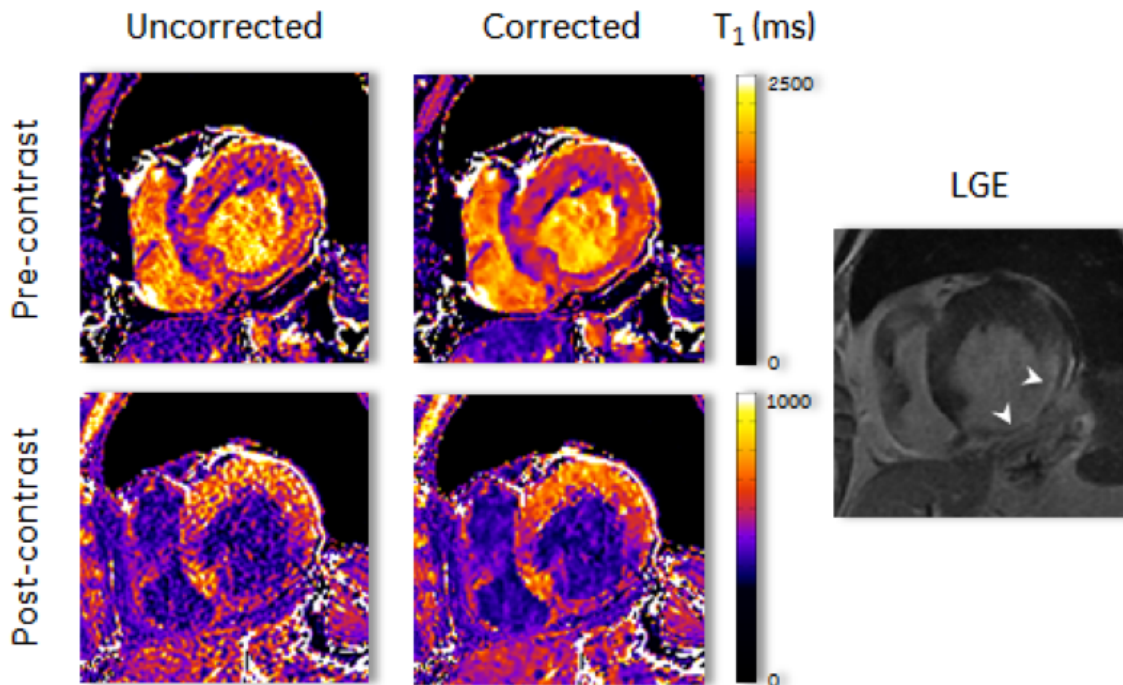


Figure 5.5: Short-axis views on a patient with history of myocardial infarction. Visual sub-endocardial enhancement in the left ventricular inferolateral wall can be depicted (bright area). Representative pre- and post-contrast cardiac T_1 maps acquired using SMART₁Map without correction (left) and with noise correction (right) are shown. A late gadolinium enhancement image is displayed on the right. Corrected T_1 maps exhibit a higher overall signal-to-noise ratio and better homogeneity. Both T_1 maps result in comparable T_1 measurements with reduced standard deviation in the myocardium and blood for the denoised map.

in 5.5. Pre-contrast and post-contrast T_1 maps with clear late gadolinium enhancement are shown. The quantitative T_1 results are summarized in table 5.5, together with the average values of the coefficient of determination R^2 of the three-parameter fitting, as returned by the Levenberg-Marquardt algorithm. As expected, maps without noise correction showed a lower overall SNR than the corrected maps derived from the anisotropic denoising method. Better homogeneity of T_1 maps can be observed on the denoised method in both pre- and post-contrast maps. Post-contrast maps exhibited differences in T_1 values between normal and infarcted myocardium, with clear depiction of the scar. Better goodness of fit was also found with noise correction.

Table 5.5: Impact of the proposed noise correction technique on T_1 values on a patient with myocardial infarction (see Figure 5.5).

	T_1 blood (ms)		T_1 scar (ms)		T_1 myocardium (ms)	
	Uncorrected	Corrected	Uncorrected	Corrected	Uncorrected	Corrected
Pre-contrast	$1990.8 \pm \mathbf{250}^a$ ($R^2 = 0.995$) ^b	$1987.4 \pm \mathbf{180}$ ($R^2 = 0.995$)	$1707.5 \pm \mathbf{260}$ ($R^2 = 0.995$)	$1696.4 \pm \mathbf{140}$ ($R^2 = 0.999$)	$1602.3 \pm \mathbf{300}$ ($R^2 = 0.991$)	$1581.1 \pm \mathbf{160}$ ($R^2 = 0.998$)
Post-contrast	$427.3 \pm \mathbf{81}$ ($R^2 = 0.984$) ^b	$422.0 \pm \mathbf{51}$ ($R^2 = 0.991$)	$536.7 \pm \mathbf{98}$ ($R^2 = 0.989$)	$530.9 \pm \mathbf{78}$ ($R^2 = 0.995$)	$778.8 \pm \mathbf{140}$ ($R^2 = 0.985$)	$759.1 \pm \mathbf{60}$ ($R^2 = 0.996$)

^a T_1 is expressed as mean \pm standard deviation in the ROI

^b R_2 is the mean value of the coefficient of determination in the ROI, as returned by the T_1 Levenberg-Marquardt fit

Table 5.5 shows mean segmental T_1 values for corrected and uncorrected methods performed on the same patient in pre- and post-contrast images. With both methods, lower post-contrast T_1 values were estimated in the scar (Post-contrast: $T_{1\text{corrected}} = 531 \pm 78$ ms vs. $T_{1\text{uncorrected}} = 537 \pm 98$ ms) compared to the healthy myocardium ($T_{1\text{corrected}} = 759 \pm 60$ ms vs. $T_{1\text{uncorrected}} = 779 \pm 140$ ms) but were higher than the T_1 values estimated in the left ventricle blood pool ($T_{1\text{corrected}} = 422 \pm 51$ ms vs. $T_{1\text{uncorrected}} = 427 \pm 81$ ms). Globally, the measured wide range of T_1 values across the scar, myocardium and blood pool was considerably reduced. Additionally, the coefficient of determination R^2 of the T_1 fitting algorithm was consistently improved by the denoising.

5.4 Discussion

The results presented in this study show that MR T_1 mapping can be improved, both qualitatively and quantitatively, using the proposed denoising algorithm. We demonstrate the benefit of using an anisotropic denoising technique in myocardial T_1 mapping. To the best of our knowledge, this is the first time a denoising model exploiting spatiotemporal correlations in the T_1 -weighted space with Beltrami regularization is proposed and its impact on accuracy and precision is investigated.

The method provides improved T_1 maps of the myocardium from T_1 -weighted images acquired with SR techniques, which results in improved reproducibility of myocardial T_1 estimates. The technique exploits the spatio-temporal correlations in

the T_1 -weighted images based on Beltrami regularization [63] and additionally uses a vectorial anisotropic approach to improve T_1 map sharpness. Our results in simulation and ex vivo experiments demonstrate that BL-V-ANI performs as good as BL-V on homogenous areas, i.e. small gradients, but further preserves edges and local structures, i.e. large gradients, of T_1 -weighted images, achieving high quality denoising without introducing staircasing. This translated into higher SNR, enhanced sharpness and improved precision without impacting accuracy. This simple technique is fully automatic and comes with negligible computational costs, which makes the method highly suitable for clinical practice. T_1 values found in this study at 3T (pre- and post-contrast) are in good agreement with previous studies [59, 61, 98]. As expected, higher noise level was observed using standard SR technique when combined with a three-parameter fit without noise correction. Large SD observed in ex vivo data and volunteers/patients confirm the benefit of using a denoising method, even for data acquired with higher SNR and potentially increased precision. In both simulations and in vivo imaging, mean T_1 values were preserved after denoising, with tendency to be slightly lower after denoising, which was not statistically significant. On the other hand, precision in the myocardium, as measured by the SD, was shown to be significantly reduced after correction.

In ex vivo studies, we compared the method with a Multi-NEX acquisition. Additional evaluation can be done by considering, for example, inversion-recovery spin echo acquisitions as gold-standard. This choice was made to isolate the impact of noise from all other potential sources of variability. The results show that noise can impact the accuracy of certain individual pixels in the myocardium. This did not translate into a loss of accuracy when considering average T_1 values over the whole myocardium but rather into a loss of precision. Indeed, in the Bland-Altman plots in Figure 5.2, no bias was observed with/without denoising; however, denoising did reduce the variability drastically.

Since the technique does not use any signal variation prior, and is therefore not biased by any application-specific assumption, the proposed strategy can potentially benefit any parametric mapping applications such as T_2 , T_2^* or diffusion tensor imaging (DTI) where precision is influenced by noise [77, 83]. Nevertheless, and although

the method is expected to provide similar quality performance, the impact of denoising on the reproducibility and spatial variability of other mapping sequences may be different and is beyond the scope of this study.

Imperfect breath-holding, observed during the experiments on few patients, has shown to degrade the final map. The proposed imaging protocol included a non-rigid image registration, before denoising and fitting, to align pixels across the T_1 -weighted images. While the technique substantially reduces artifacts and blurring in the reconstructed T_1 map, it would be potentially of interest to examine the impact of the motion correction step on the noise distribution.

Acquisition noise is only one factor affecting precision and accuracy in QMRI. In particular, bias in T_1 measurements can also partly be explained by the B_1 field inhomogeneities, mainly due to B_1 penetration effects or transmitting coil geometry, particularly at 3T [89, 9]. Integrating actual B_1 maps into the fitting procedure, instead on adding a third parameter accounting for the uncertainty on the applied flip angle, might further reduce the variability in the T_1 values and improve the precision of myocardial T_1 mapping techniques by SR. Patient motion is another source of variability. Precision could be further improved using, for instance, free breathing acquisition and motion correction [77, 83, 98]. This would allow multi shots acquisitions e.g. a shorter acquisition window and consequently less cardiac blurring, with an increased matrix size. In addition, this would also allow acquiring more samples of the T_1 recovery curve and not being limited to eight TSs. This method would however require more time to reconstruct the T_1 map.

In the present application, the systematic noise bias introduced by the Rician noise distribution was disregarded in the denoising stage because it was intrinsically taken into account by the three-parameter fitting stage. In other circumstances, it may be desirable to consider the actual noise distribution directly in the denoising stage (e.g. if a two-parameter T_1 fit is desired). This can be done by applying the fit on the complex data with phase restoration, which can be particularly efficient to preserve noise statistic, especially for low SNR data [10].

Finally, the study was only performed in volunteer and patient subjects with a small sample size. Further studies with larger samples are warranted to confirm the

benefit of the denoising approach to improve the reproducibility and spatial variability of myocardial T_1 mapping in patients. Further investigations in cross-center and cross-vendor studies with additional comparisons with inversion-recovery type methods (e.g. MOLLI) are desired and currently investigated.

In conclusion, we have developed a denoising technique for magnetic resonance T_1 mapping using an anisotropic denoising model that preserves accuracy and improves precision in myocardial T_1 mapping. Ultimately, this technique could enable accurate and precise reconstruction of myocardial T_1 maps, with the potential to offer better visual image analysis and improved performance of post-processing procedures such as registration and segmentation.

5.5 Sub-Study A: T_1 mapping denoising: comparison with MOLLI²

Background

Quantitative myocardial T_1 mapping is a valuable tool that intrinsically studies tissue properties in the human body. Among all cardiac mapping sequences, inversion-recovery (IR) techniques for T_1 assessment, such as MOLLI (Messroghli et al., MRM 2004), have received most clinical interest due to their high precision and satisfactory SNR, but suffer from the underestimation of myocardial T_1 values. On the other hand, saturation-recovery (SR) based myocardial T_1 mapping techniques, such as SMART₁Map (Slavin et al., JCMR 2013), have shown better accuracy but are more sensitive to acquisition noise and are therefore limited by low precision of T_1 values. We propose to further improve the precision of SR-based techniques by employing a novel denoising method which exploits the spatio-temporal correlations in the T_1 -weighted images.

Methods

16 healthy volunteers (12 males, 39.8 years) were included in the study. Both SMART₁Map and MOLLI sequences were scanned on a 3T MRI scanner (HDxt signa, GE Healthcare, Milwaukee, WI). Imaging parameters for SMART₁Map were: FOV = 220 x 220 mm², slice thickness = 8 mm, TR/TE = 3.74 / 1.63 ms, FA = 45 degrees. T_1 -weighted images were acquired in short-axis views before and 15 min after gadolinium injection. Our proposed denoising technique uses a coupling between the T_1 -weighted images by employing a Beltrami constraint along the T_1 -weighted images. T_1 maps were then reconstructed using a 3-parameter-fitting. Regions of interest in the left ventricular septum were drawn by an expert observer and statistical analysis was performed to compare accuracy and precision of T_1 values.

²Published in: Bustin A, Menini A, Liu S, Burschka D, Codreanu A, Al-Mallah H M, Janich A M, Brau CS A, Odille F. Saturation-based myocardial T_1 mapping with denoising: initial comparative study with MOLLI. *Journal of Cardiovascular Magnetic Resonance*, 2017

Results

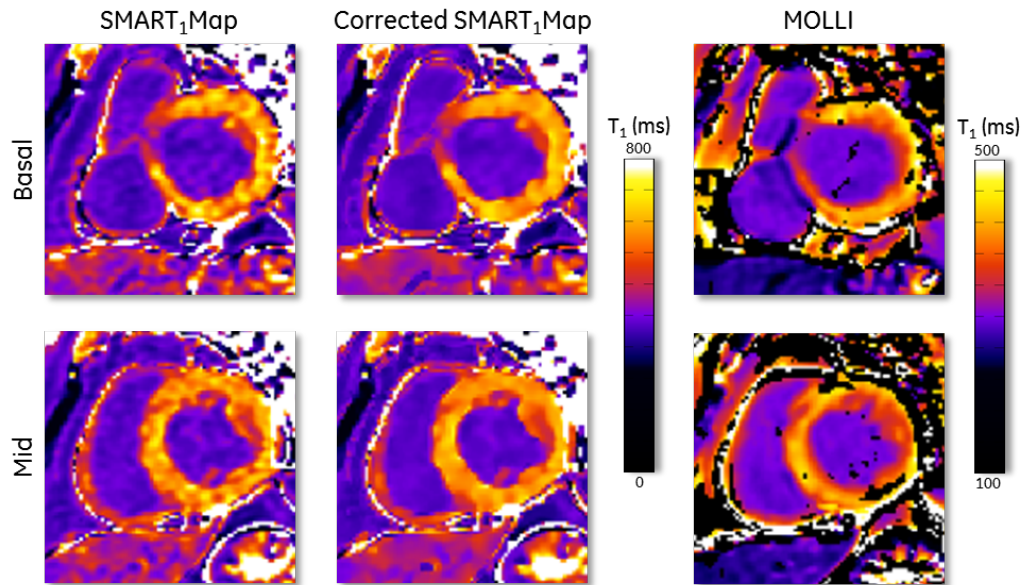


Figure 5.6: Visual comparison of the standard SMART₁Map, corrected SMART₁Map and MOLLI in a 56 year-old patient undergoing cardiovascular magnetic resonance imaging. Post-contrast basal (top) and mid-short axis slices (bottom) are shown.

The average time to complete denoising was ≈ 1.4 sec for one map. Precision of T_1 values was significantly improved using the proposed technique (Figure 5.7), with a decrease in standard deviation (pre-contrast: 33%, post-contrast: 27%, $p < 0.05$). Precision tends to be closer to MOLLI after correction in pre-contrast (SMART_{uncorr}: 182 ms / SMART_{corr}: 122 ms / MOLLI: 55 ms, $p < 0.05$) and post-contrast (153 / 111 / 37 ms, $p < 0.05$). Accuracy was preserved as no difference in mean T_1 values was observed in the myocardium (pre-contrast: 1411 / 1407 ms, post-contrast: 767 / 760 ms, $p > 0.05$). Noise observed in the original SMART₁Map was considerably reduced by the denoising, with sharp edges, better visualization of cardiac features and homogeneous maps, as observed in MOLLI maps (Figures 5.6 and 5.8).

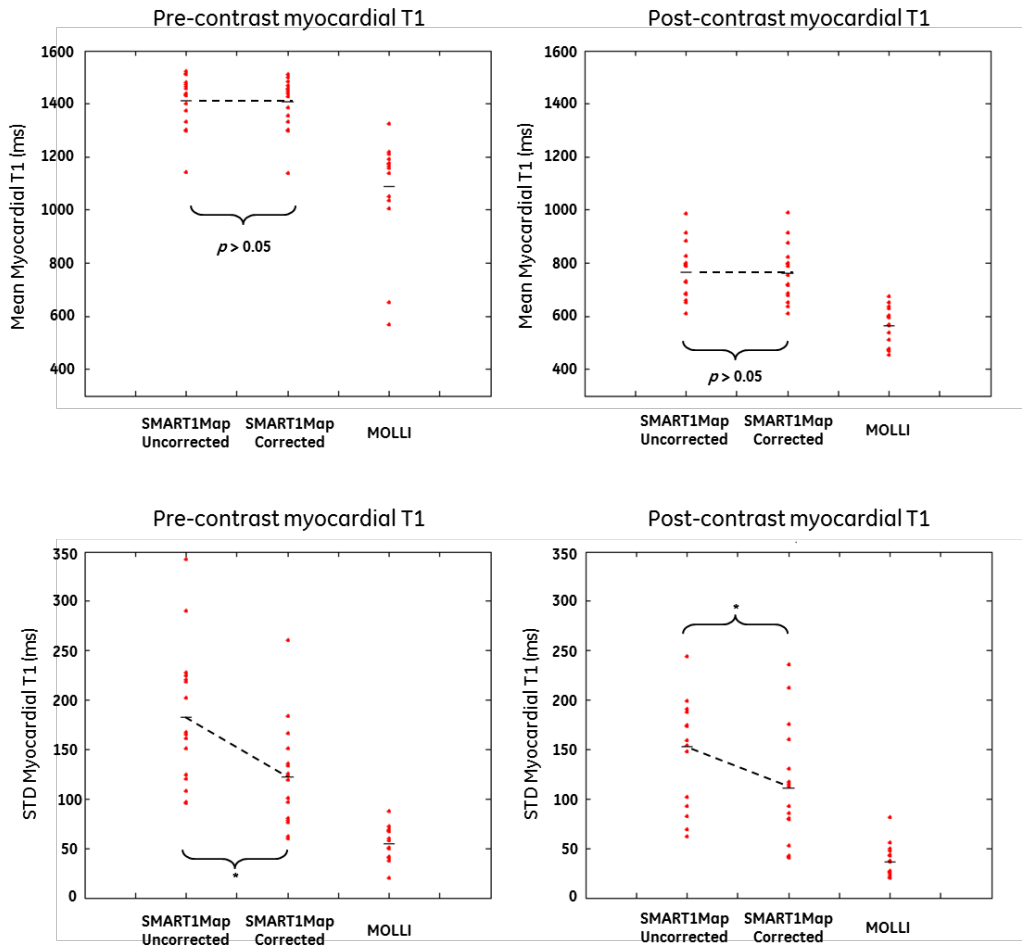


Figure 5.7: Qualitative analysis in volunteers, Plots comparing (left) pre-contrast and (right) post-contrast cardiac T_1 measurements derived from SR-based SMART $_1$ Map and IR-based MOLLI cardiac T_1 mapping ($*p < 0.05$).




Figure 5.8: Pre-contrast SMART₁Map with and without denoised and MOLLI in a 29-year-old male on a mid-ventricular short axis view. Applying the anisotropic correction clearly improves the map quality, exhibiting sharper details. Both corrected and uncorrected SMART₁Map T_1 values result in comparable T_1 measurements with reduced standard deviation in the myocardium.

Conclusion

We demonstrated the feasibility of denoised SR-based myocardial T_1 mapping to perform accurate and precise T_1 analysis even in the setting of significant noise. Therefore, this approach could be beneficial to accurately detect and quantify myocardial fibrosis at no additional cost.

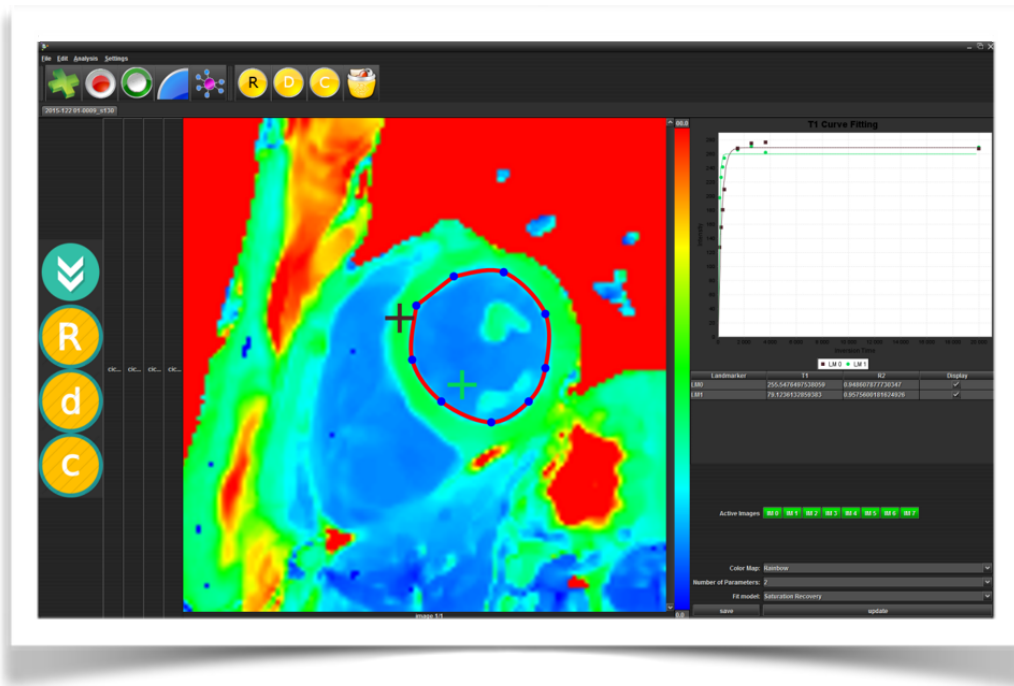


Figure 5.9: Screenshot of the developed parameter mapping software for T_1 and ECV analysis. The software generates T_1 map in both pre- and post-contrast from DICOM (digital imaging and communications in medicine) image data. The user can view and register images within one T_1 series, using a developed non-rigid registration algorithm. The user can then perform the proposed automatic denoising to improve precision and quality of the final T_1 map. The software can be applied on any saturation and inversion-recovery-based sequences with both two and three-parameter fitting. Reconstructed T_1 maps are displayed on the graphical user interface and can be analyzed using the assessment toolkit, which includes the standard six AHA segments extraction, ROIs selection, endocardial and epicardial contours segmentation, and mean and standard deviation analyses. On this screenshot, a T_1 map on a patient is shown on the map viewer and the related curve fitting analysis section is displayed on the right side. (Courtesy of Shufang Liu)

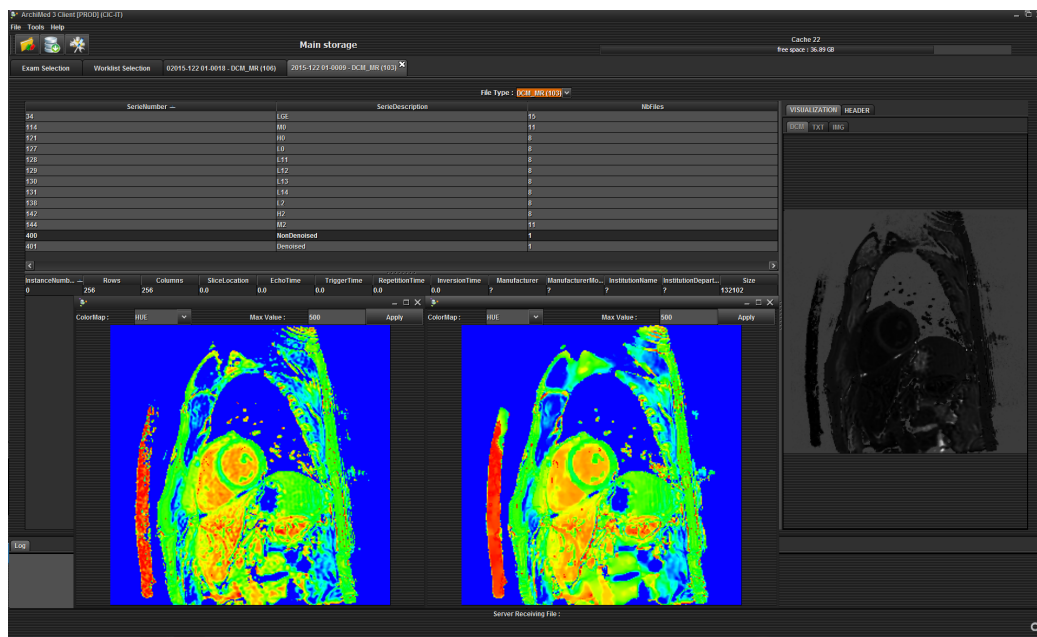


Figure 5.10: Screenshot of the integrated automatic denoising tool on the hospital PACS workstation. Myocardial T_1 map can be reconstructed and denoised directly in the hospital PACS system with minimal radiologist interaction. (Courtesy of Emilien Micard)

Chapter 6

Motion-Correction in Free-Breathing Single-Shot Cardiac MRI¹

¹Published in: Bustin A, Menini A, Janich M, Burschka D, Felblinger J, Brau A, Odille F. Motion Estimated-Compensated Reconstruction with Preserved-Features in Free-Breathing Cardiac MRI. MICCAI 2016, and in Bustin A, et al. Joint Denoising and Motion Correction: Initial Application in Single-Shot Cardiac MRI. *Journal of Cardiovascular Magnetic Resonance*, 2015, 17(Suppl 1):Q29 doi: 10.1186/1532-429X-17-S1-Q29. and presented in part at Proceedings of the 23th Annual Meeting of ISMRM. Toronto; 2015. 0810

6.1 Introduction

Cardiac magnetic resonance imaging is a valuable tool for myocardial structure, function, and tissue assessment, providing essential information for clinical diagnosis and treatment decisions in cardiovascular disease. Using standard segmented sequences in which data acquisition is segmented over multiple heart beats, good image quality can be obtained in patients with regular cardiac rhythm and good breath-holding ability; however, image quality can be degraded by motion artifacts when scanning patients with arrhythmia or poor breath-hold compliance.

In comparison to segmented acquisitions, single-shot techniques can be applied for rapid image acquisition of a whole slice within a single shot, greatly reducing the scan time. Due to the short acquisition duration of single-shot techniques (typically ≤ 200 ms), artifacts from intra-shot motion are negligible, therefore such methods tend to be robust against cardiac and breathing motion. However, this motion robustness comes at the expense of lower spatial resolution and signal to noise ratio (SNR). An example of the benefit of single-shot over segmented late Gadolinium enhanced (LGE) imaging in a patient who could not breath-hold is shown in Figure 6.1.

Recent techniques proposed to enhance the SNR of single-shot methods by motion correcting and then averaging multiple single-shot images acquired in free-breathing [60]. While this technique shows good results with low acceleration factors, it may not provide optimal image quality for higher undersampling, introducing blurring and undersampling artifacts, mainly due to the higher weight given to the regularization. Moreover, frames with large registration errors are discarded, leading to a loss of information and scan efficiency.

Motion estimation and correction has been studied intensively in recent years. Batchelor et al. [11] proposed a first generalized reconstruction framework for motion compensation. The method allows arbitrary motion to be compensated by solving a general matrix inversion problem. This technique, however, requires an adequate knowledge of the displacement fields. The recent GRICS method [84] extended this work by jointly estimating the motion and the recovered image through the usage of a motion model provided by external sensors (e.g. ECG, respiratory belt). Most




Figure 6.1: Comparison between 2D segmented LGE (top) and 2D single-shot LGE (bottom) technique in short axis and long axis vertical on a 77 years old male with chest pain. Patient experienced difficulty holding his breath during the acquisition. The LGE pattern is poorly defined on the routine 2D LGE images because of severe artifact due to poor breath holding. However, free-breathing single-shot LGE images clearly showed sub-endocardial enhancement in the septum and anterior wall from mid-cavity to the apex (arrow). (Courtesy of Morrision Hospital, Swansea, UK)

motion artifacts are removed by solving the motion compensated model iteratively. However, the algorithm is based on a linearity assumption of the motion model which challenges the reliability of the result if the displacement of the heart is not perfectly correlated with the available physiological signals (this can be the case when large displacements occur). Furthermore, this technique requires heavy computational effort since the motion is extracted in the k-space domain and the process is part of an iterative optimization.

In this work, we sought to develop an efficient motion correction implementation

suitable for reconstructing a high-resolution, high-SNR image from multiple accelerated single-shot images. The proposed method combines the benefits of using a hybrid self-navigated sampling scheme (see Figure 6.1) with a joint reconstruction framework. We use single-shot imaging because it allows motion fields to be easily extracted and used for motion compensation. In the image reconstruction step, a highly efficient feature-preserving regularization scheme (Beltrami) is proposed for recovering sharp details. We show that the proposed method is robust to high acceleration factors and yields results with efficient noise reduction and better overall image quality at a low computational cost.

6.2 Theory

We begin by reviewing the principles of general motion correction techniques to set the framework for the technical description of the proposed method.

6.2.1 General Motion Compensation Framework

Motion compensation techniques aim to solve the following inverse problem: find an underlying image ρ free of motion artifacts, given derived measurements s through the system E , affected by noise μ :

$$S = E\rho + \mu \tag{6.1}$$

Where E is the encoding matrix, generally composed of a sampling operator ξ , a Fourier transform F , coil sensitivity maps σ (in case of parallel imaging reconstruction), and a motion warping operator W describing a nonrigid deformation for each shot. Here $\rho \in \mathbb{C}^{n_x \times n_y}$ and $s \in \mathbb{C}^{n_x \times n_y / acc \times n_r \times n_c}$ are complex data with n_c the number of receiver coils, n_r the number of repetitions and acc the acceleration factor. In this work the acquired data s represents the k-space data from multiple single-shot images and is generally corrupted by noise. A typical approach to solve this problem is to minimize the squared difference as assessed by the Euclidian norm (also called least-norm-squared) as seen in Equation (6.2), which has the advantage to be

differentiable.

$$\|E\rho - S\|^2 \quad (6.2)$$

However, this problem is generally ill-posed (e.g. due to undersampling and noise) leading to non-uniqueness of the solution, if it exists. Thus, regularity constraints on the unknown solution ρ have to be considered. Furthermore, motion should also be considered as unknown. The general joint optimization framework is then defined as

$$\rho = \underset{(\rho, \vec{u})}{\operatorname{argmin}} \left\{ \|s - E(\vec{u})\rho\|_2^2 + \lambda\phi(\rho) \right\} \text{ where } E(\vec{u}) = \xi F \sigma W(\vec{u}) \quad (6.3)$$

Here E is the encoding matrix, \vec{u} represents the displacement fields, ϕ is the chosen regularization function and $\lambda > 0$ is the corresponding regularization parameter. The optimization problem in Equation (6.3) is solved in four steps: (i) we first use the k-space center of the single-shot images to extract a self-navigation signal and to cluster the raw data into a reduced number of respiratory bins; (ii) we reconstruct the images from each bin independently using a Beltrami-regularized SENSE (B-SENSE) reconstruction; (iii) then an estimate of the motion is obtained using a non-rigid registration (minimization of Equation (6.3) with regards to \vec{u}) and (iv) a high resolution/SNR image is generated using the proposed motion-compensated reconstruction process (minimization with regards to ρ). A general description of the method is shown in Figure 6.2.

6.2.2 Respiratory Binning for Raw Data Clustering

The concept of respiratory binning is widely used in cardiac imaging [115] and has opened ways to efficient whole-heart coronary imaging (6,7). Our objective here is to extract a respiratory signal in order to combine the acquired k-spaces according to their position in the respiratory cycle. Here we propose a fully automatic self-navigated method: at each shot, the center region of the k-space is fully sampled and the periphery is undersampled. A processing of the self-navigation data (i.e. k-space center area), based on a principal component analysis and temporal filtering, allows

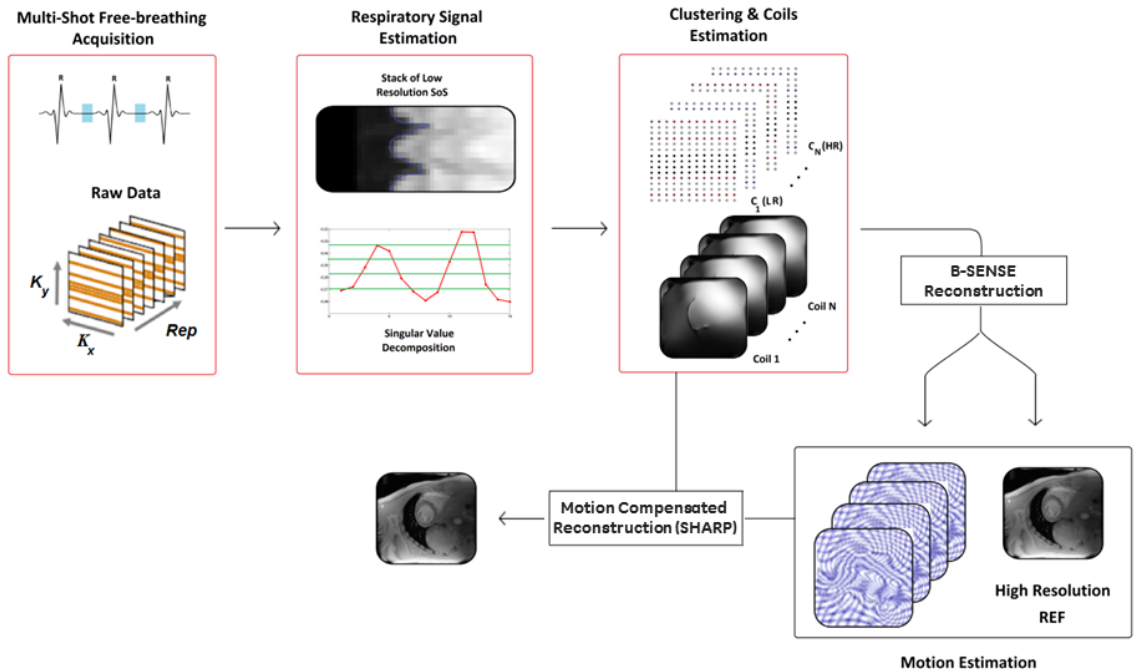


Figure 6.2: Schematic illustration of the proposed motion-compensated reconstruction including the non-rigid motion extraction. The method acquired the single-shots with complementary interleaved trajectories, leading to uniform samplings in the phase encoding direction which allows an optimal combination of the k-spaces. Respiratory signal is automatically extracted from the low-resolution images and allows k-spaces to be combined according to their positions in the breathing signal. The motion model, initialized by registering the images from each respiratory bin, is incorporated into the reconstruction process.

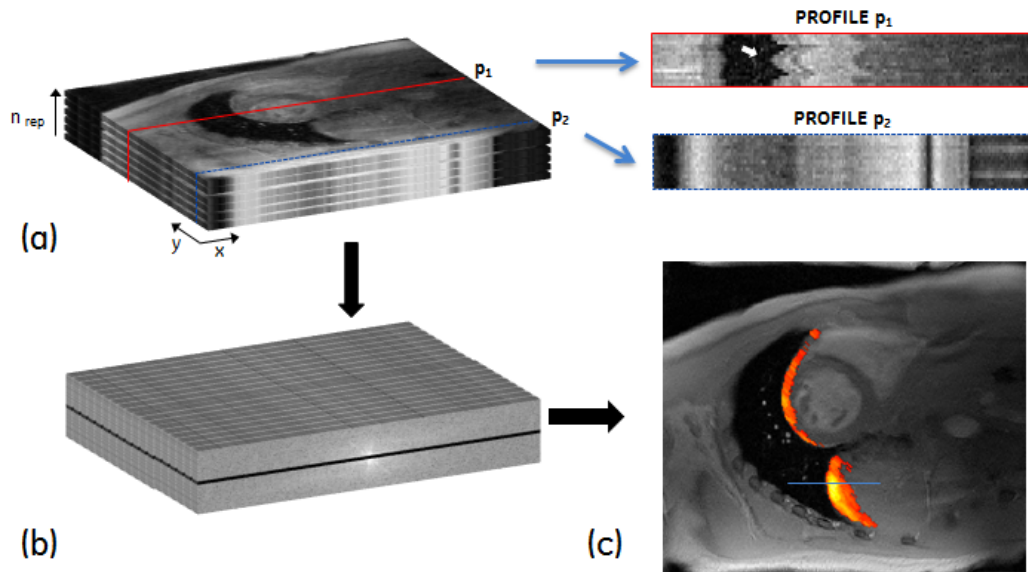


Figure 6.3: Schematic illustration of the binning method. Low-resolution images are stacked along the time dimension (a). One cross-section through the 3D spatial-temporal volume shows a 2D image in the $x-t$ plane. If the pixels are static over time (profile p_2), the $x-t$ image would exhibit horizontal and vertical structures, while a profile going through the myocardium (profile p_1) with moving objects along the time dimension, would show more inhomogeneous structures, thus generating a pattern with breathing motion information. Low-resolution images were obtained from the k -space center of each shot and were stacked together along the time dimension. The moving parts of the image were automatically extracted by band pass filtering along the time-dimension with the frequency interval $[0.05, 0.5]$ Hz corresponding to the frequency range of respiration (b). The motion mask (d) was obtained by summing over time and thresholding the back-transformed volume.

a respiratory signal to be extracted. The principle is illustrated in Figure 6.3 and is explained in details in Appendix A.

The next step is to separate the respiratory signal into several phases (also called "bins") and to assign each acquired shot to a specific respiratory phase (8). This binning strategy would split the data into fewer motion states with negligible respiration motion and lower undersampling into each of them. In this work we used an irregular sampling pattern (described in the Methods Section 6.7) which ensures complementary acquisitions across shots

6.2.3 Beltrami-Regularized SENSE

Images from each respiratory bin $(\rho_i)_{i=1,\dots,n_r}$ are first individually reconstructed by solving

$$\rho_i = \underset{\rho}{\operatorname{argmin}} \left\{ \|s_i - E_i(\vec{u}) \rho\|_2^2 + \lambda \sqrt{1 + \beta^2 \|\nabla \rho\|^2} \right\} \quad (6.4)$$

where $\lambda > 0$ is a given balancing parameter tuning the weight between the data consistency and the regularization penalty. The first term in Equation (6.4) is a data fidelity term that aims to minimize the difference between the reconstructed image and the acquired data. The Beltrami regularization $\sqrt{1 + \beta^2 \|\nabla \rho\|^2}$ has been introduced in the field of string theory for physics and has shown high potential in several imaging problems, including image denoising and enhancement [91]. The potential of this geometric framework lies in the general definition of the space-feature manifold and the choice of its metric. In particular, the metric can be chosen such that the Beltrami energy corresponds to an arbitrary interpolation between quadratic or total variation gradient penalties. As a consequence, this regularization has very interesting fundamental properties, such as geometric regularization interpolating between Gaussian diffusion ($\beta \gg 1$) and TV regularization ($\beta \rightarrow 0$). One can indeed recognize some equivalence to the well-known total variation (TV) regularization introduced in reference [100]. TV has the advantage to suppress noise while preserving sharp edges, but suffers from the presence of flat regions, also called "staircasing" effect. In Chapter 4, we showed that Beltrami regularization is able to maintain the advantage of TV (edges preserving, noise reduction) as well as reducing the effect of staircasing observed in the reconstructed image.

B-SENSE is very similar to compressed sensing SENSE (CS-SENSE) methods presented by other authors [70], where here Beltrami is making the image sparse in the gradient domain. Even though this suggests that B-SENSE has a close relationship with the compressed sensing (CS) theory, it is however not CS as defined by Candes [26], especially due to the pseudo-random undersampling pattern used here (i.e. a uniform random pattern is used in [26]).

While Equation (6.4) can be solved using a conjugate gradient (CG) algorithm

or similar techniques (LSQR, GMRES, ...), we propose to use an advanced optimization scheme by using a primal-dual projected gradient approach [29]. The image ρ is reconstructed by solving the optimization problem retrospectively in the primal and dual space, thus taking advantage of the joint information. The minimum is independent of the initialization which can be either equal to zero or equal to a previous solution. This optimization is practical and simple to implement and has the potential to converge faster than classical primal gradient-descent (see 4). The only computationally expensive part of the algorithm is the calculation of the operator $E^H E$ required for the primal variable update. The Beltrami primal-dual framework is summarized in the three-step algorithm 5.

Algorithm 5 Beltrami Primal-Dual reconstruction for efficient motion correction reconstruction in cardiac single-shot imaging.

Input: raw data s , operator E , λ , τ_1 , τ_2 , β .

Initialize:

$$\begin{aligned}\rho^0 &= 0 \\ d^0 &= \nabla \rho^0\end{aligned}$$

repeat

$$\left\{ \begin{array}{l} \hat{d}^{k+1} \leftarrow (1 - \tau_1) d^k + \beta \tau_1 \nabla \rho^k \sqrt{\beta^2 - \|d^k\|^2} \\ d^{k+1} \leftarrow \frac{\beta \hat{d}^{k+1}}{\max(|\hat{d}^{k+1}|, \beta)} \\ \rho^{k+1} \leftarrow \rho^k + \tau_2 (E^H (s) - E^H E (\rho^k) + \lambda^{-1} \text{div} d^{k+1}) \end{array} \right.$$

until stopping criterion is reached

Output: motion-free reconstructed image ρ

6.2.4 Motion Estimation

Respiratory motion estimation is accomplished using independent non-rigid registration of the images reconstructed from each respiratory bin with B-SENSE. We opt for a classical multi-scale regularized nonrigid registration method (with sum-of-squared




Figure 6.4: Left: Comparison of the RMSE between a no motion-corrected dataset and our non-rigid regularized registration model. The x-axis represents the set of single-shot images (here we considered 15 single-shots). Right: Visual comparison of the accumulated residual error $\sum_{shot_i} \|W(\vec{u}_i \rho_i - \rho_{ref})\|^2$ for the two cases. One can clearly see the respiratory motion on the dashed line (no motion correction). The respiratory-related myocardium motion is mostly seen on the anterior and inferior wall and can be corrected with a non-rigid registration model.

differences matching criterion as described e.g. in [82]). The reference image used for the nonrigid registration is chosen as the bin with the highest amount of data. A typical respiratory motion error is shown in Figure 6.4, where the accumulated residual error is displayed when a nonrigid registration model is applied (right) and is compared to no correction (middle). Although highly patient specific, the respiratory-related myocardium motion is mostly seen on the anterior and inferior wall and has a maximum displacement of 10 mm in this example [102].

6.2.5 Motion Compensated Reconstruction with Preserved-Features

This section presents the final step for solving the motion compensated problem in Equation (6.3). The aim of the method is to reconstruct the high resolution, high SNR image from the acquired raw data $s = (s_i)_{i=1, \dots, n_r}$. For the motion compensated reconstruction, we solve the following optimization problem, with the acquisition model now including the estimated motion fields:

$$\rho = \underset{\rho}{\operatorname{argmin}} \left\{ \|s - E(\vec{u})\rho\|_2^2 + \lambda \sqrt{1 + \beta^2 \|\nabla \rho\|^2} \right\} \quad (6.5)$$

As in the previous Section 6.2.3, we use a primal-dual projected gradient approach, employing the Beltrami energy as regularity prior. Note that regularization is always preferred in motion compensated reconstruction due to the ill-conditioning induced by the motion operators, as shown in [8].

6.3 Methods

The proposed motion-compensated reconstruction algorithm was applied and validated with different experiments, both quantitatively and qualitatively, using Matlab (The MathWorks, Natick, MA) on a PC with Intel Xeon 3.3 GHz CPU and 64GB ram. The experiments were performed on 3T MR750w and 1.5T MR450w systems (GE Healthcare, WI, USA). All reconstructions were implemented using the motion compensated reconstruction as described in the previous Section 6.2.5.

6.3.1 Offline Simulation on Synthetic Data

In order to perform a realistic simulation we first created a synthetic dataset based on actual LGE patient images. In one patient with suspected cardiovascular disease, 4 repetitions of a cardiac-gated single-shot LGE scan were acquired in free-breathing 7-10 minutes after Gadolinium injection. Informed consent was obtained from the patient. Cardiac images were obtained with spoiled fast gradient echo and the following parameters: matrix size 192 x 256, in-plane spatial resolution 1.52 x 1.52 mm² in short axis with slice thickness = 8 mm, readout flip angle = 20°, echo time (TE) = 2.02 ms, mid-diastolic trigger delay, pulse repetition time (TR) = 4.43 ms and SENSE factor = 2 with partial Fourier.

Synthetic k -space data were created by application of synthetic coil sensitivity maps (with Gaussian profiles) to the LGE images, Fourier transformation and undersampling in the phase encoding direction. A full sampling of the central k -space area (17 lines) was used and the peripheral area was undersampled with a Golden

Table 6.1: Parameters used for the different experiments. The acquisition matrix size was 192×256 .

	#repetition n_r	#calib lines	#periphery lines	acc peri	acc shot pre-bin	acc shot post-bin	NEX sequence
Simulation 1	4	32	48	3.3	2.4	-	1.67
Simulation 2	4	32	32	5	3	-	1.33
Phantom	6	32	48	3.3	2.4	-	2.5
In vivo	15	17	43	4.1	3.2	1.1 (5 bins)	4.7

Step Cartesian trajectory [38] with an acceleration factor $R = 3.3$. Spacing between samples was proportional to the Golden fraction 0.618. This trajectory enables an irregular but almost uniform distribution of the measured data for any arbitrary number of repetitions, leading to incoherent aliasing. The motion-free image was reconstructed using our motion-compensated reconstruction and compared to a motion correction method similar to that proposed by Kellman et al. [60], where the motion-free image is recovered by averaging the registered images obtained after the B-SENSE reconstruction step. We call this prior method reconstruction-registration-average (RRA).

6.3.2 Phantom Imaging

Single-shot pulse sequences were used to acquire phantom images with a 26-channel cardiac coil. The sequence was modified to take into account two sampling schemes we developed: the same Golden ratio sampling as in our offline simulation experiment and a classic interleaved trajectory. The protocol was first applied to acquire phantom images with a resolution of $1.48 \times 1.48 \text{ mm}^2$. A translational motion was imposed to the table to mimic respiratory motion. The maximum displacement of the phantom in this acquisition was about 12 mm. The SNR of our accelerated acquisition was compared with a full k -space image acquired without motion.

6.3.3 In-Vivo Validation Experiment with Self-Navigation

In vivo cardiac datasets from two healthy adult volunteers were acquired on a 1.5T scanner using a 32-channel cardiac coil. Fifteen multi-shots slices of a free breathing,

Table 6.2: Recommended single-shot LGE protocol on a GE system.

Patient Position	
Patient position	supine
Patient entry	feet first
Coil	cardiac coil
Series description	SSHLGE
Scanning Range	
FOV	35
Phase FOV	0.75
Slice thickness	8 mm
Number slices	enough to cover LV
Skip	2 mm
Imaging Parameters	
Plane	Oblique
Mode	2D
Family	Gradient echo
Pulse sequence	FGRE, FSPGR, or FIESTA
Imaging options	Cardiac gating, IR prep, ASSET
Details	
TE	MinFull
Flip angle	45 (FIESTA) or 25 (FSPGR, FGRE)
Prep time	appropriate time from CINE IR
Acceleration factor	2
Frequency	160 (FIESTA) or 140 (FSPGR, FGRE)
NEX	1
Phase	160 (FIESTA) or 140 (FSPGR, FGRE)
Bandwidth	68 kHz (1.5T), 100 kHz (3.0T) 83 kHz (1.5T FIESTA), 100 kHz (3T FIESTA)
Shim	auto
Cardiac Gating	
# RR interval	1 or 2 (recommended)
Trigger window	auto
Views per segment	fixed
Trigger delay	appropriate delay from CINE scan

cardiac-gated, inversion-recovery-prepared spoiled fast gradient echo sequence were collected in short axis with the following parameters: TE = 2.10 ms, TR = 4.52 ms, 8 mm slice thickness, FOV = 253 x 338 mm², matrix size 192 x 256 and 1.32 x 1.32 mm² in-plane resolution, diastolic trigger delay. A fully sampled reference was acquired additionally in breath-hold for visual comparison. Each shot consists of 60 k-space lines: the central k-space was fully sampled with 17 lines and the periphery (43 lines) was undersampled, leading to a global acceleration factor of 3.2. The raw data were acquired using the same acquisition scheme as defined in the previous Section 6.3.2. An estimate of the respiratory signal was extracted using the proposed self-navigated technique, and was subsequently used to separate the acquired data into five respiratory bins. An overview of the parameters used in this study is given in table 1. The total time needed to acquire the 15 single-shot images with an acceleration of 3.2 was approximately 1 min and 8 seconds in free breathing using an acquisition window of 271 ms per heartbeat (27% of the R-R interval) at a heart rate of approximately 60 beats per minute. The total measurement time for a segmented acquisition in breath-hold with same resolution was approximately 20 seconds. A 17×17 region in the center of the k-space was used for coil sensitivity estimation. Here, coil sensitivity maps were obtained by dividing the complex image from each coil by the sum-of-squares (SoS) of the images for all coils. In order to reduce the Gibbs ringing effect commonly seen and the potential ghosting artifacts coming from the shot-averaging, a Fermi window was applied on the magnitude of the k-space, leading to smooth and artifact-free coil sensitivity maps.

Reconstruction results were again compared to the RRA method, where the motion-corrected image was obtained by averaging the 8 single-shot images reconstructed with lowest registration error.

6.4 Results

6.4.1 Computing Resources

The time needed to run the motion-compensated reconstruction for 15 shots of matrix size equal to 192×256 with 32-channel cardiac coils was about 4 min 52 seconds and is reduced to 1 min 35 seconds when the data are clustered into 5 bins. This time included the computation of the motion parameters as well as the Beltrami reconstruction. The mean computational time to get the respiratory signal was about 1 sec when applying a MATLAB implementation of the self-navigated algorithm on 20 single-shots.

6.4.2 Offline Simulation and Phantom Imaging

Example reconstruction results on the simulated data generated from a patient with non-ischemic cardiomyopathy are shown in Figure 6.5. Late Gadolinium imaging enables here the depiction of hyper-enhancement in the ventricular septum. One can see a spatially blurred result with a standard reconstruction-registration-average (RRA) method. The benefit of incorporating a motion correction step directly into the reconstruction process is seen on (d). The proposed joint reconstruction and motion correction model exhibits significant quality improvement over each method with an acceleration factor $r = 3.3$ while reconstructing sharper edges (arrows) and small structures. Moreover, when r increases, the standard RRA method shows a more blurry result while the proposed technique preserves image quality. For both acceleration rates, we can see that RRA effectively reduces the motion artifacts at the expense of a spatial blurring. This method is well suited for low undersampling ($r = 2$) and performs well for SNR improvement by averaging the images (SNR grows with the square root of the number of images to average). On the other hand, the proposed motion-compensated reconstruction not only reduces motion artifacts, but more importantly enhances the sharpness and SNR of the reconstructed image without introducing blurring.

We varied the undersampling ratio and compared the root mean squared error




Figure 6.5: Cardiac short axis reconstruction of a synthetic dataset generated from 4 single-shot LGE acquisitions in free-breathing on a 37-year old patient with acceleration factors $r = 2.5$ and $r = 3.3$. a) One reconstruction using a classic SENSE (192 x 256), b) Sum-of-Squares (all repetitions), c) Reconstruction-Registration-Average, d) Iterative reconstruction with Beltrami regularization and motion correction.

(RMSE) of the proposed technique, RRA, and uncorrected sum-of-squares (SoS) reconstructions. When the undersampling is small enough, our method ($\text{RMSE}_{\text{acc3}} = 0.053$) and RRA ($\text{RMSE}_{\text{acc3}} = 0.064$) have similar performance with small errors. However for higher acceleration factors the performance of the proposed reconstruction ($\text{RMSE}_{\text{acc6}} = 0.061$) is far better compared to RRA ($\text{RMSE}_{\text{acc6}} = 0.110$) both in terms of reconstruction accuracy and image quality (Figure 6.5).

As illustrated in Figure 6.5, one can adjust the regularization parameter according to the reduction of noise needed. A high regularization parameter will give more importance to the regularization term and will lead to a less noisy result while preserving fine details. Experimental results indicate that $\lambda \in [0.1, 1]$ yields good results. Comparisons with a classic Tikhonov reconstruction are shown Figure 6.7. The results present the reconstructed phantom motion experiments where, here, the motion




Figure 6.6: Quality comparison of the Beltrami-regularized motion compensated reconstruction for different values of λ when the acceleration factor is equal to 3.3. The effect of denoising is clearly seen with a larger λ . The model reconstructs a motion-free image with sharp edges while reducing the effect of staircasing usually seen with a total variation norm.

has been applied with the table. The sum-of-squares reconstruction (Figure 6.7, left) clearly exhibits the effect of motion. As in the simulation section, the RRA method exhibits blurry result (due to the undersampling), although giving a motion-corrected denoised image. A visual improvement can be noticed when applying a motion compensated reconstruction with Tikhonov regularization. The latter method performs well but is, however, unable to recover sharper edges and some residual aliasing artifacts can still be seen on the recovered image. The use of a fast primal-dual algorithm combined with Beltrami regularization makes the proposed reconstruction robust with better performance in terms of image quality, reduced artifacts and sharpness (Figure 6.7, right).

6.4.3 In-Vivo Validation Experiment with Self-Navigation

Scans and reconstructions were completed successfully. Short-axis images of the myocardium of a healthy subject without Gadolinium injection and without inversion recovery preparation are shown in Figure 6.8. The figure gives a visual comparison of a classic sum-of-squares reconstruction with our model and a breath-hold acquisition. Both cardiac structures (myocardium wall, papillary muscles) and non-cardiac




Figure 6.7: Motion-compensated reconstructions on a phantom using two different regularization methods with acceleration factor 3.6 from a 26 channel phantom scan, 10 single-shot repetitions have been acquired. Continuous motion was simulated with table movement. From left to right: Sum-of-Squares (SoS), Reconstruction-Registration-Average (RRA), Tikhonov, Beltrami.

structures (blood vessels) are very well recovered with the proposed reconstruction. The method yields significant sharpening of the myocardium wall and papillary muscles. However, due to the relatively high-undersampling, the RRA method is unable to recover a good quality image, exhibiting blurry structures and losing some of the details in the image such that blood vessels (yellow arrows).

This particularity is also seen in Figure 6.8d where a specific intensity profile is plotted. The white line shows the position of this profile. The sharpness of the edges on the motion-corrected reconstruction is confirmed as well as the fidelity to the breath-hold acquisition. Sharpness is evaluated with local gradient entropy as proposed by Vuissoz et al. [118] and is shown in Figure 6.8e for two volunteers. Again, the combination of nonrigid registration model with the motion compensated Beltrami reconstruction yields a motion-free image with high SNR. The method is able to recover good quality images with an acceleration factor of up to eight, which clearly promotes the acquisition of high spatial resolution images.

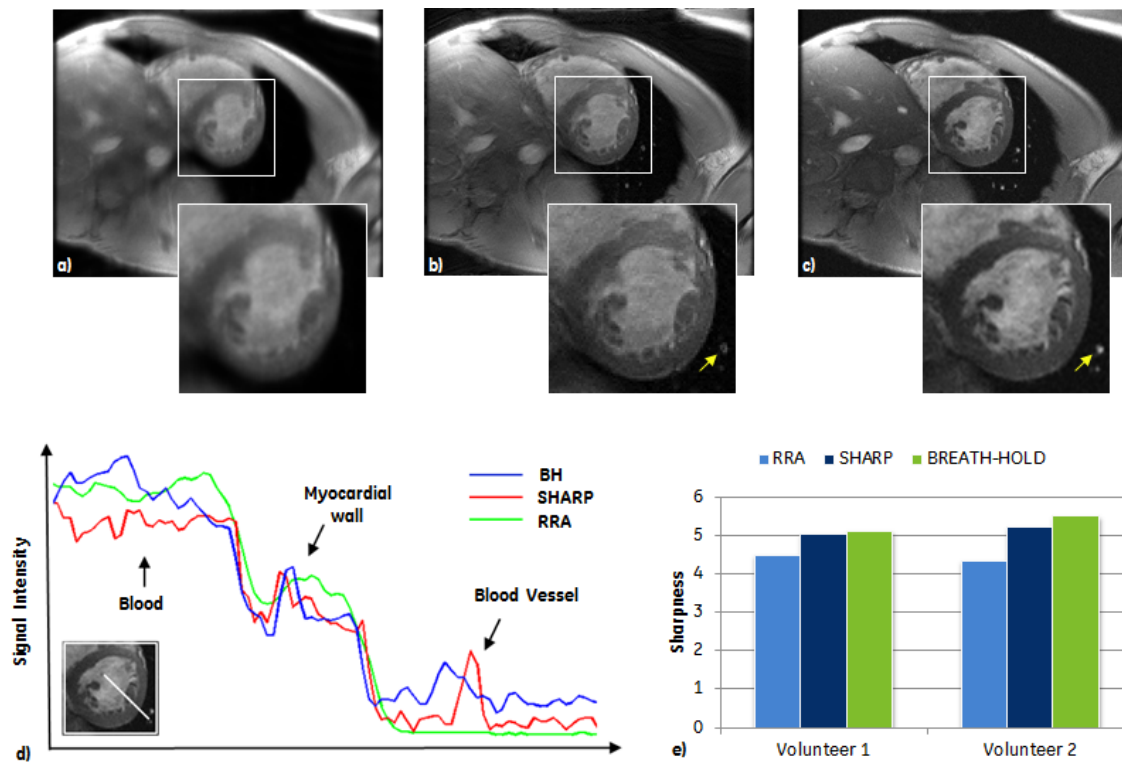


Figure 6.8: Top: Cardiac single-shot spoiled fast gradient echo images from a conventional breath-held acquisition (c) and from a free-breathing acquisition with a naive reconstruction (a) and with the proposed motion-compensated reconstruction method (b). The acquisition has been accelerated by a factor of 3.2. Bottom: A cardiac profile of the reconstructed short axis image is shown (d). The white line displays the position of the considered profile across the myocardium. The proposed reconstruction is able to reconstruct sharpen features (myocardium wall, blood vessel) with high-fidelity, while the RRA method results is a more blurry (smooth profile on the edges) image with loss of information (blood vessel). Myocardium sharpness is evaluated for two volunteers (e).

6.5 Discussion

We introduced a new free-breathing single-shot LGE pipeline including an optimized sampling and the associated joint reconstruction and motion correction algorithm designed for fast and robust cardiac imaging. By incorporating the estimated motion into the reconstruction process, we increased the robustness of the model and exhibited good quality images. In this work, our self-navigation scheme was used in free-breathing single-shot imaging but it might also be used to improve the reconstruction of breath-held images. Both qualitative and quantitative analysis of the model applied on in-vivo cardiac imaging positively have shown improvement in speed and image quality over widely used motion correction method.

In this study, we used a fast and automatic self-navigated binning strategy that aims to cluster the acquired raw data into similar motion states. While preliminary results have shown improved image quality and better motion estimation, additional optimization of number of bins and number of repetitions is still required to maintain an optimal tradeoff between reconstruction quality, reconstruction time and accuracy of motion estimates. The motion corrected images show better visual quality than classic reconstructions but appear less sharp than corresponding breath-held acquisitions, especially for high accelerations. Possible explanations are the inaccuracies in motion estimates or other effects related to MR physics, such as spin history or changes in B_0 and B_1 inhomogeneities induced by breathing.

The proposed nonrigid registration model was applied independently for each frame. However, additional information can be provided by considering the whole set of frames. In this way, a multi-frame nonrigid registration combined with a rank constraint and a total variation norm, should improve motion estimation accuracy on regions suffering from large undersampling artifacts. More advanced registration methods and non-Cartesian sampling deserve further investigation as well.

Meanwhile, it would be of great interest to adapt the local Beltrami algorithm to a patch-based non-local regularization which should perform better to recover local image information. Each region of the heart is thus reconstructed and weighted using all similar regions in the image. Besides, dynamic weights can be optimized

for our Beltrami regularity by solving a specific entropy minimization problem. This nonlocal-regularized reconstruction model associated with parallel imaging and motion correction will be investigated in future work.

One interesting application of the proposed motion correction model is for high-resolution 3D isotropic LGE single-shot imaging of the heart such as the one proposed recently in [43] for myocardial scar assessment. This will allow the reconstruction of 3D isotropic motion corrected volumes by keeping the advantages of a 2D acquisition (high tissue and vessel contrast, short acquisition time), e.g. using super-resolution techniques [81]. Other applications, such as abdominal imaging and coronary vessel imaging, are being investigated.

A limitation to the method is that potential through-plane motion cannot be corrected, although it remains small compared to the slice thickness. To overcome this problem, one could consider weighting the images according to the motion amplitude compared to the target image or acquiring 3D slab instead of 2D slice data and applying motion compensation. The preliminary results presented in this work should be confirmed with further patient studies.

The feasibility of the proposed reconstruction has been evaluated in simulation, phantom and volunteer experiments. The method has been shown to allow non-rigid motion correction while efficiently recovering features, thanks to the Beltrami regularization scheme. The conventional segmented LGE acquisition is limited by the maximum breath-hold time, which limits the signal-to-noise ratio and/or spatial resolution. This limitation is overcome by the presented free-breathing approach. Ultimately, this method could enable accurate motion corrected reconstruction of single-shot images with higher spatial resolution and a higher signal-to-noise ratio compared to conventional segmented methods, with the potential to offer high-quality LGE imaging in challenging patients. The extension of the proposed method to other types of single-shot cardiac pulse sequences can be envisioned as well,

6.6 Appendix A: Processing of the self-navigation data²

Low-resolution images are obtained from the k-space center of each shot by the root sum of squares of the individual coil images, and are then stacked together along the time dimension (see example in Figure 6.3). In order to extract the moving part of the image, the high frequency component of every profile along the y-dimension is removed by band-pass filtering. Here we chose a frequency interval of [0.05, 0.5] Hz corresponding to human adult respiration (22). Finally, moving objects are separated from the background by averaging along the time dimension all back-transformed cross-sections and by applying morphological operations (opening/closing and thresholding). Large deformations are mainly visible on the diaphragm (where navigators are usually positioned). Obtained mask is overlaid onto the image as seen on (Figure 6.3c) and a respiratory signal can then be extracted.

This step is completed by using a singular value decomposition (SVD) on the 2D stack $T = UDV^*$ where U and V represents the left and right singular vectors and D is the diagonal matrix containing the singular values of T (here, only the moving pixels from the previous step are considered). The SVD has the advantage to of being independent of the order of the columns and can thus provide information on the image with randomly concatenated columns. The best fit line representing the temporal stack turns out to be the first column of the left-singular vectors (first singular vector of V) which corresponds to the highest singular value. We can thus have a good approximation of the true respiratory signal directly from this vector. The algorithm is computationally inexpensive since it only requires filtering step and SVD of a stack of 2D small patches.

The next step is to separate the respiratory signal into several phases (also called "bins") and to assign each acquired shot to a specific respiratory phase (8). This binning strategy would split the data into fewer motion states with negligible respiration

²Published in: Bustin A, et al.. A fully automated binning method for improved SHARP reconstruction of free-breathing cardiac images. *Journal of Cardiovascular Magnetic Resonance* 2016, 18(Suppl 1):P267 DOI: 10.1186/1532-429X-18-S1-P267.

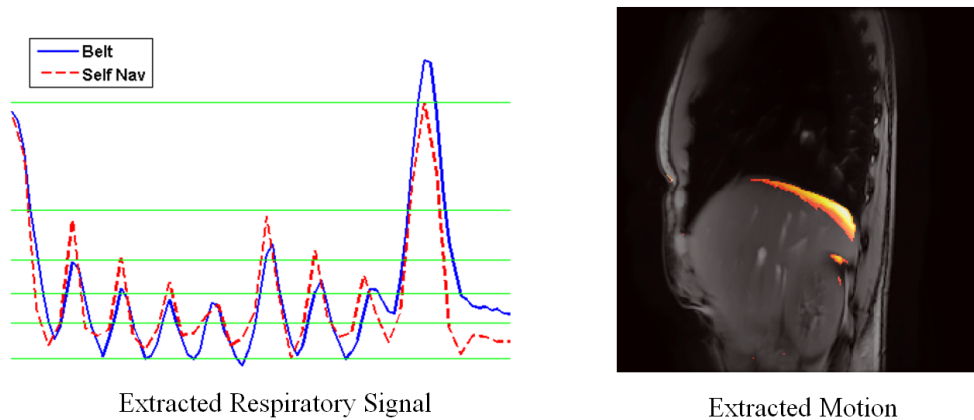


Figure 6.9: Data clustering step of the proposed automatic method for liver motion tracking. Single-shot 2D spoiled fast gradient echo sequence has been used to acquire 150 consecutive slices in free-breathing at a temporal rate of 400 ms corresponding to a total duration of 1 min. One pneumatic belt placed on the subjects thorax is used for comparison. (left) A plot of the extracted respiratory signal versus number of slice. The green lines represent the binning of the shots into five different respiratory states. The extracted respiratory signal (red) shows good agreement with the respiratory belt (blue), with a coefficient of determination $R^2 = 0.76$. (right) The image illustrates the moving parts of the liver (color) used to extract the breathing signal via SVD.

motion and lower undersampling into each of them. Here we see the advantage of using an irregular sampling such as the interleaved trajectory which ensures complementary acquisitions across shots and enables a full coverage of the k -space.

An example of automatic respiratory signal is shown Figure 6.9 on a liver scan from a healthy volunteer.

6.7 Sub-Study A: Free-Breathing 2D Radial Cardiac MRI³

Synopsis

Cardiac MRI is affected by both cardiac and respiratory motion. While ECG-gated imaging in breath hold is the clinical method of choice, free-breathing methods are needed in patients with limited breath hold capability. This work describes a method to obtain free-breathing cine datasets with high SNR and high spatial resolution (1.4 mm in-plane) from a completely self-gated Golden Angle radial scan within an 11 seconds scan time. The motion compensated reconstruction technique takes advantage of calibrated displacement fields extracted from the radial data to recover motion artifact-free cardiac phases. Beyond cine imaging, contrast-enhanced cardiac imaging can also be expected to benefit from this motion compensated reconstruction strategy.

Motivation

Cardiac magnetic resonance imaging (CMR) is affected by both cardiac and respiratory motion which are commonly compensated with ECG-gated imaging during a breath hold. However, free-breathing methods are preferred or even required in patients with limited breath hold capability. Free-breathing approaches utilize either rapid single-shot scans to reduce respiratory motion artifacts at the expense of spatial resolution or segmented respiratory-gated scans at the expense of scan time efficiency. Golden Angle radial [124] sampling has favorable properties for free-breathing CMR, including motion robustness and flexible reconstruction options. These properties

³Published in: Fischer A, Menini A, Bustin A, et al. Motion Compensated Reconstruction from Free-Breathing 2D Radial Cardiac MRI Data. *Journal of Cardiovascular Magnetic Resonance*, 2016, 18(Suppl 1):O110, and presented in Fischer A, Menini A, Bustin A, et al.. Fully Self-Gated Motion Compensated Cine Reconstruction from Free-Breathing 2D Radial Cardiac MRI Data, *Proceedings of the 24th Annual Meeting of ISMRM*. Singapore; 2016. p. 0789

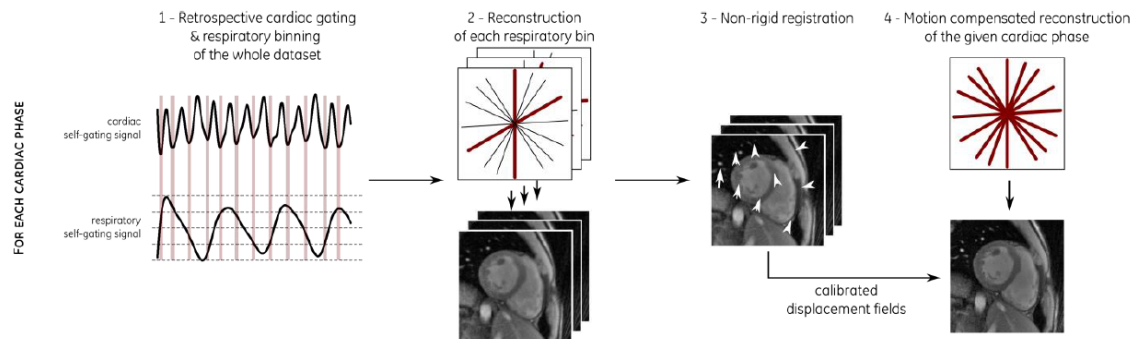


Figure 6.10: Scheme of the MCR. For each cardiac phase, steps 1-4 are performed. The vertical red lines in step 1 indicate the desired cardiac phase in the cardiac self-gating signal. The dashed black lines indicate the targeted bins in the respiratory self-gating signal. (Courtesy Andre Fischer)

also make Golden Angle radial sampling a good sampling strategy for motion compensated reconstructions [45, 33]. In this work, we propose the combination of a self-gated 2D radial Golden Angle data acquisition scheme with motion compensated reconstruction strategy in order to obtain high-resolution motion compensated CMR data in every desired cardiac phase.

Methods

A 2D golden angle radial spoiled gradient echo sequence was employed to collect short-axis cardiac data from four volunteers using the following parameters: 3.0T scanner (MR750w, GE Healthcare, Waukesha/WI, USA), $\alpha = 15$ deg, $BW = \pm 125$ kHz, 256 readout points, $TR = 4.3$ ms, $TE = 1.5$ ms, slice thickness = 8 mm, FOV = 360 mm circular, resolution 1.4×1.4 mm², 2584 projections resulting in 11 seconds total scan time. Self-gating signals for cardiac and respiratory motion were obtained from the k-space DC signal [46] and compared by regression to simultaneously recorded peripheral gating (PG) and respiratory belt signals. Using the derived self-gating signals, images in various cardiac and respiratory phases were reconstructed using a binning scheme as previously described, e.g., in [46] (Figure 6.9, step 1). For each cardiac phase, four images were reconstructed as follows: 1) Combining all




Figure 6.11: The derived self-gating signals (blue dashed lines) show a very high correlation ($R^2 = 95.9\%$) to simultaneously recorded PG and respiratory belt signals (red solid lines).

free-breathing data per cardiac phase without motion management (NoMCR); 2) Retrospective self-gating selecting data (approx. 20-25%) closest to end-expiration of the desired cardiac phase (RG); 3) Motion Compensated Reconstruction (MCR); and 4) a filtered backprojection (FBP) of the RG data, which was viewed as ground truth to assess whether MCR accurately reconstructs the desired cardiac phase. NoMCR and RG recovered missing data using a non-Cartesian iterative SENSE [93]. The MCR was accomplished in four steps using the cardiac and respiratory self-gating signals: 1) clustering of all data of the desired cardiac phase into eight respiratory bins; 2) filtered backprojection of these eight respiratory states; 3) extraction of the calibration displacement fields between the bins by applying a non-rigid registration; 4) motion compensated SENSE-like reconstruction using the previously extracted calibration motion fields [11] (Figure 6.9).

Results

A very high correlation ($R^2 = 95.9\%$) between the self-gating signals, the recorded PG curve, and respiratory belt was observed (Figure 6.10). Data from 11 seconds scan time were found sufficient to result in high-quality motion compensated cine images. Figures 6.11 and 6.12 show representative images and temporal profiles of the four different reconstruction schemes. Combining all free-breathing data into one dataset (NoMCR) leads to high SNR images with significant respiratory motion induced blurring. RG improves apparent sharpness (e.g. diastole), but can suffer artifacts from severe undersampling (e.g. systole). MCR takes advantage of all projections in all respiratory states by including the calibration displacement fields into the reconstruction. Thus it recovers a sharp image with high SNR where small anatomical features are preserved. Furthermore, the cardiac phases in the MCR are in good agreement with the FBP ground truth images (Figure 6.13c), demonstrating the ability of the MCR to accurately recover the desired cardiac phases.

Discussion

Cardiac cine images from 11 seconds fully self-gated, free-breathing Golden Angle radial datasets were successfully reconstructed using MCR. This scheme proved to be superior to non-Cartesian iterative SENSE reconstructions without motion correction (NoMCR, RG). The described technique offers several potential advantages: improved patient comfort and clinical workflow since neither breath holds nor ECG leads are necessary; high image quality with both high SNR and high spatial resolution; time-efficient acquisition since no data are discarded. Instead, every acquired sample is used in the motion compensated image.

The presented technique could also be extended to simplify and improve contrast-enhanced cardiac exams such as perfusion or late enhancement.




Figure 6.12: Results obtained from volunteer 1. (a) MCR exhibits improved sharpness (white arrows, (c)) and recovers more details than NoMCR and RG (e.g., near anterior wall in diastole). (b,c): Temporal profiles (see dashed lines in (a)).




Figure 6.13: Results obtained from volunteer 2. (a) MCR exhibits improved sharpness (white arrows, (b)) and recovers more details than NoMCR and RG (e.g., endocardium of lateral wall in diastole). (b,c): Temporal profiles (see dashed lines in (a)). Note the pulmonary vessel moving in and out of the profile position in MCR (white arrows, (c)) as confirmed by FBP.

Chapter 7

Super-Resolution in Cardiac Cine MRI¹

¹Published in: Odille F, Bustin A, Chen B, Vuissoz PA, Felblinger J. Motion-Corrected, Super-Resolution Reconstruction for High-Resolution 3D Cardiac Cine MRI. MICCAI 2015, Part III, LNCS 9351, pp. 435-442, 2015. DOI: 10.1007/978-3-319-24574-4 52

7.1 Introduction

High-resolution, 3D isotropic cine imaging of the heart is challenging because it requires lengthy acquisitions, even with efficient imaging sequences, and therefore advanced patient motion management. Such an imaging technique would be useful for understanding complex anatomy and function in congenital heart diseases or for imaging small cardiac structures such as the atrium or the valves. It might also help reducing the variability of ventricular volumetric assessment in cardiovascular diseases (stroke volume, ejection fraction) which is generally high with the conventional 2D cine image stacks due to the difficulty of segmenting the myocardium near the base of the ventricles.

The b-SSFP (balanced steady-state free precession) sequence is one of the most widely used imaging techniques for cardiac cine imaging. This is because it provides the highest signal-to-noise ratio per unit time among all known sequence, with good T2/T1 contrast [101]. Moreover when the slice thickness is not too small (5 to 10 mm) the 2D b-SSFP does not suffer from severe motion-induced signal dropouts thanks to its fully balanced gradients. Unlike its 3D version, the 2D b-SSFP provides excellent contrast between tissues and blood/vessels due to the inflow effect (fresh spins moving in/out of the slice) [101]. A multi-slice 2D b-SSFP acquisition is therefore a good candidate but it suffers from poor resolution in the slice direction. One possible way to overcome this limitation is to combine multiple 2D stacks acquired in different orientations using a super-resolution (SR) reconstruction [50, 48, 90]. Applying SR techniques to chest or abdominal imaging is difficult because the image stacks to be combined need to be perfectly aligned. Rigid motion correction techniques have been combined with SR algorithms [99] but nonrigid motion is even more challenging and has not been addressed yet.

In this study, we propose a method for high-resolution, 3D isotropic cine imaging of the heart. The acquisition strategy consists of multiple 2D image stacks with different orientations obtained by a 2D multi-slice b-SSFP sequence during free breathing. The image reconstruction consists of two steps: (i) a motion-compensated 3D cine

	2D multi-slice SSFP	3D SSFP
Blood/tissue contrast	☑ Good (in-flow effect)*	☒ Poor (no in-flow effect)*
Spatial resolution	☒ 5-10 mm slice thickness	☑ Isotropic possible

Figure 7.1: 2D multi-slice SSFP vs. 3D SSFP

reconstruction of each stack; (ii) a super-resolution reconstruction combining the multiple cine volume stacks into an isotropic cine volume. The objective of this study is to show the feasibility of the method in actual volunteer and patient experiments.

7.2 Image Reconstruction

7.2.1 Motion-Compensated 3D Cine Reconstructions

In the first step each stack of 2D slices (from one particular volume orientation) is reconstructed with a nonrigid motion compensation approach named cine-GRICS (generalized reconstruction by inversion of coupled systems) [118, 84]. Here the method reconstructs $N_{ph} = 32$ volumes corresponding to 32 phases of the cardiac cycle. In order to reconstruct a given cardiac phase, k-space data within a cardiac acceptance window (here $W_{card} = 120$ ms) are selected. As described hereafter the motion-compensated reconstruction for this cardiac phase is formulated as the joint reconstruction of a 3D image ρ_i ($i = 1, 2, 3$ for SA, HLA, VLA) and a motion model α accounting for respiratory motion and for cardiac motion within the acceptance window. The motion model approximates 3D time-varying nonrigid displacement fields $u(x, y, z, t)$ in the imaged volume such that $u(x, y, z, t) \approx \sum_{k=1}^{N_s} \alpha_k(x, y, z) s_k(t)$ as described in Ref. [84], with some motion signals $s_k(t)$ acquired simultaneously with the k-space data m . Here we use $N_s = 5$ motion signals: $[s_1(t) s_2(t)]$ are the signals from two pneumatic belts (chest and abdomen); $[s_3(t) s_4(t)]$ are their first-order time derivatives; and $[s_5(t)]$ is the cardiac phase signal derived from the ECG recordings ($[s_5(t)]$ has values between 0 and the mean cardiac cycle length in ms).

For each cardiac phase the joint reconstruction is achieved by solving:

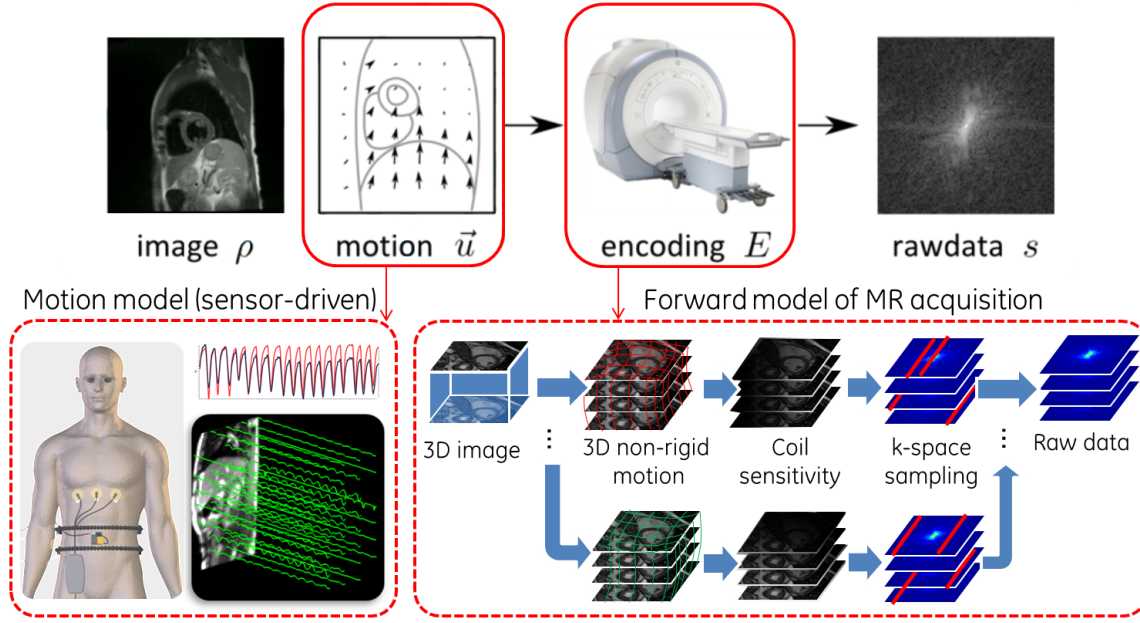


Figure 7.2: Motion-correction with GRICS reconstruction. (Courtesy of Freddy Odille)

$$\min_{(\rho_i, \alpha)} \left\{ \|E(\alpha, s) \rho_i - m\|_2^2 + \lambda \|\nabla \alpha\|^2 \right\} \quad (7.1)$$

Here $E(\alpha, s)$ denotes the forward acquisition model, i.e. the motion-corrupted 2D multi-slice acquisition. It includes a 3D nonrigid transformation operator for each motion state, followed by a slice selection, receiver coil sensitivity weightings, 2D Fourier transforms and k-space sampling operators similar to the description in Ref. [84]. The respiratory signals $s_1(t) \dots s_4(t)$ are centered on their most likely value (i.e. end-expiratory plateaus) so that all image stacks are reconstructed in a consistent motion state. Problem (7.1) is solved by a multi-resolution Gauss-Newton scheme, alternating between optimization with respect to ρ and α as described in [84].

7.2.2 Super-Resolution Reconstruction

In a second step the N motion-compensated 3D cine volumes (obtained from different orientations, $N = 3$ in this study) are combined using SR reconstruction. For a given

cardiac phase if $\rho_1 \dots \rho_N$ denote the N motion-compensated 3D images, the SR image is found by solving:

$$\min_x \left\{ \sum_{i=1}^N \|D_i B_i T_i - \rho_i\|^2 + R(x) \right\} \quad (7.2)$$

Here T_i is a rigid image transformation that takes the SR image from the desired reconstructed orientation to the orientation of the i^{th} acquisition (i.e. it is an interpolation operator that can describe an arbitrary orientation); $D_i B_i$ is a slice selection operator including a blurring operator B_i (i.e. a sum in the slice direction in the range of the slice thickness) and a downsampling operator D_i (in the slice direction). Two regularizers were tested in this study: Tikhonov regularization, i.e. $R(x) = \|x\|^2$ and Beltrami regularization (see Chapter 4, i.e. $R(x) = (1 + \beta^2 \|\nabla x\|^2)^{1/2}$). The Beltrami regularizer is a modified version of the well-known total variation and has better feature-preserving properties. In the case of Tikhonov regularization the SR reconstruction is a linear least-squares problem and is solved by a conjugate gradient solver. For the Beltrami regularization a primal-dual projected gradient scheme is used as described in the previous Chapter 4.

7.3 MRI Experiments

Eight healthy volunteers and ten patients with major breathing difficulties (Duchenne muscular dystrophy) underwent a cardiac functional MRI. Experiments were conducted on either a 1.5 T Signa HDxt (10 patients and 5 volunteers) or a 3 T Signa HDxt scanner (3 volunteers) (General Electric Healthcare, USA). For both the volunteers and the patients written informed consent was obtained and both studies were approved by an ethical committee. A 2D multi-slice b-SSFP sequence was used to cover the left ventricle in three different orientations: short-axis (SA), horizontal long axis (HLA) and vertical long axis (VLA). The orientations were those planned by the MR technologist (they were chosen to be independent but not strictly orthogonal). All acquisitions were performed during free breathing and were ungated (retrospective gating was included in the cine-GRICS reconstruction step). Typical parameters

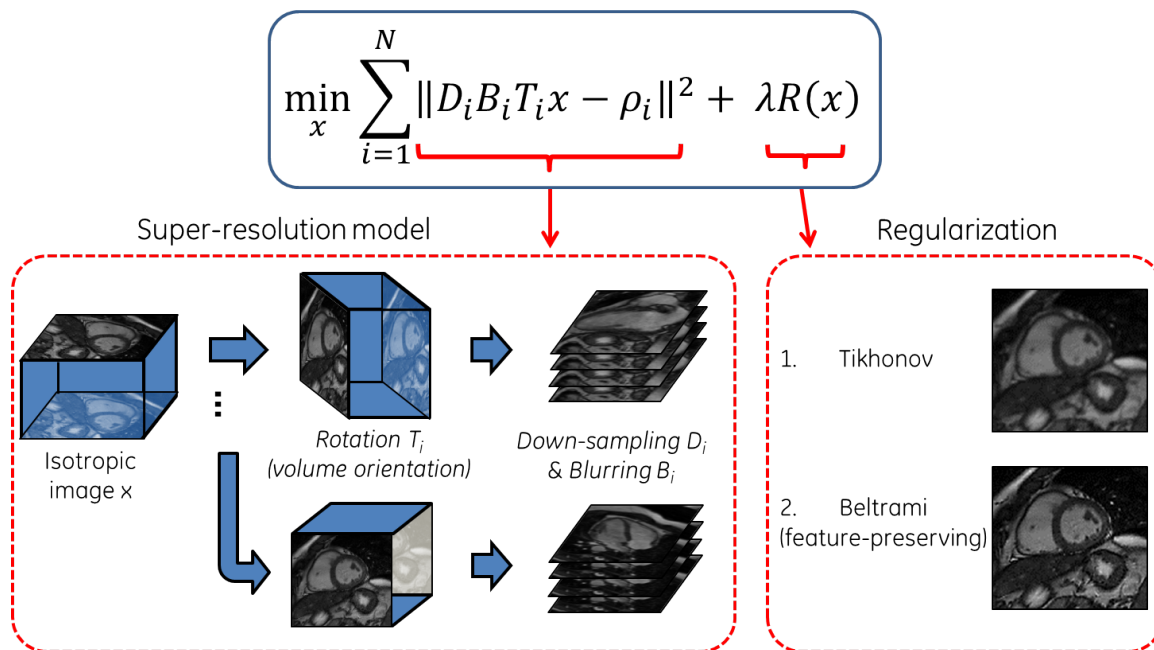


Figure 7.3: Super-resolution scheme in MRI.

were: 224x224 acquisition matrix, 20 frames (i.e. 20 k-spaces acquired per slice), native inplane resolution 1.6 x 1.6 mm², 8 mm slice thickness, TE/TR=1.8/4.1ms. The cine-GRICS motion correction step provided three cine volumes (with 32 cardiac phases) of resolution 1.6 x 1.6 x 8 mm³ (by solving Equation (7.1)) which were combined by the SR reconstruction to form an isotropic volume of resolution 1.6 x 1.6 x 1.6 mm³ (by solving Equation (7.2)). The SR reconstruction was performed on the smallest cube encompassing the intersection volume of the three image stacks.

7.4 Validation

In order to validate the method we compared the SR reconstructed images (with Tikhonov and Beltrami regularization) against the native images. We first analyzed the images visually with cine-loop movies. We also drew intensity profiles across the ventricle to highlight differences in effective resolution in various oblique directions. We then proposed a quantitative assessment. Isotropic images are expected to provide good depiction of features/edges in all three dimensions unlike the images with native

resolution which are blurred in the slice direction. It is proposed to assess the presence of features in specific directions using the structure tensor [17]. The structure tensor S of a 3D image I characterizes, for each voxel p in the image, the local orientation of the anatomical structure based on the gradients in a small neighborhood of p :

$$\begin{bmatrix} S_{xx}[p] & S_{xy}[p] & S_{xz}[p] \\ S_{xy}[p] & S_{yy}[p] & S_{yz}[p] \\ S_{xz}[p] & S_{yz}[p] & S_{zz}[p] \end{bmatrix} \quad (7.3)$$

$$\text{with } S_{ij}[p] = (W * (\partial I / \partial i \times \partial I / \partial j)) [p] \text{ for } (i, j) \in \{x, y, z\}^2$$

Here W is a Gaussian convolution kernel defining the size of the neighborhood. Before computing the structure tensor, the SA images (native resolution of $1.6 \times 1.6 \times 8 \text{ mm}^3$) were interpolated to the same isotropic resolution as the SR images (i.e. $1.6 \times 1.6 \times 1.6 \text{ mm}^3$) using a windowed *sinc* interpolation (Lanczos window). We compared the following metrics based on the structure tensor in the different reconstructions: $\langle S_{xx} \rangle$, $\langle S_{yy} \rangle$ and $\langle S_{zz} \rangle$ (i.e. the average of S_{xx} , S_{yy} and S_{zz} respectively over the reconstructed volume). S_{xx} and S_{yy} can be viewed as the amount of details (edges, features) present in the SA plane while S_{zz} relates to the direction orthogonal to SA. It is expected that the SR images can preserve directional information from the native in-plane resolution and that they can retrieve additional information in the orthogonal direction (z) from the VLA and HLA images (increased S_{zz}). Statistical differences between the native and SR images were assessed by Wilcoxon signed rank tests. A significance level of 5% was used.

7.5 Results

The time needed to acquire the three cine stacks was approximately 10 min (4 min for SA, 3 min for HLA and 3 min for VLA). The cine-GRICS reconstruction time was 180 min (SA+HLA+VLA) using a parallel C++ code running on a cluster of 16 workstations; the SR step for a full cine dataset took 14/16 min (Tikhonov/Beltrami) using Matlab (single-thread, sequential reconstruction of the 32 cardiac phases). Example




Figure 7.4: Example 3D cine datasets from a patient: native short-axis images (top row), super-resolution with Tikhonov (middle row) and with Beltrami regularization (bottom row). The arrows point out noticeable anatomical details which were not visible in the native images but were recovered by both SR techniques, especially near the base and apex of the left ventricle.

reconstructed images are shown in Figure 7.4. Images are shown both in enddiastole and in end-systole position. The visual comparison between the native SA images and the SR images shows preserved image quality in the short-axis plane with a slight noise amplification for the Tikhonov regularization. Important differences are observed in the through-plane direction, i.e. in the VLA and HLA orientations. Several anatomical structures and vessels were not visible in the native SA images but were recovered by both SR reconstruction techniques. Differences are mostly noticed near the base and near the apex of the left ventricle where the structures are orthogonal to the slice direction (see arrows in Figure 7.4).

Example intensity profiles drawn across the left ventricle are shown in Figure 7.5.

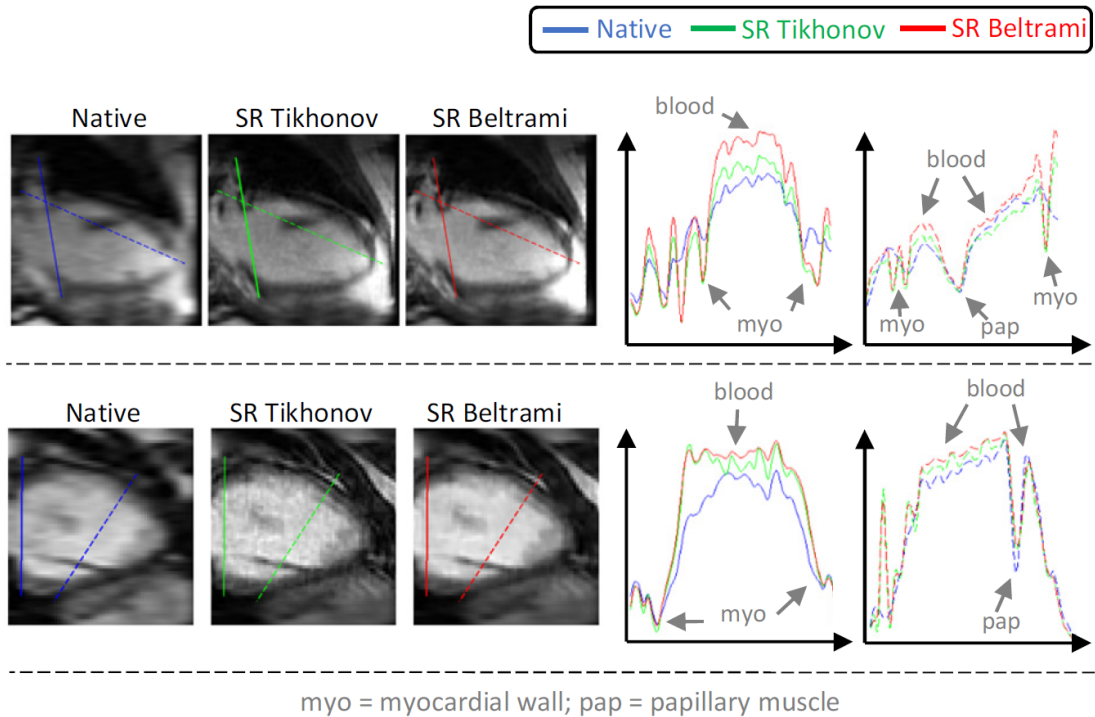


Figure 7.5: Example intensity profiles from two volunteers (top and bottom row). Two line profiles were drawn on the images (left side: native and SR images) with various oblique orientations and the corresponding intensity profiles are shown on the right. Finer details better observed in the SR images and Beltrami regularized images appear less noisy than Tikhonov ones.

Again the resolution improvement compared to the native images can be observed, especially near the mitral plane and at the apex. Structures like the myocardial wall and papillary muscles appear much sharper in SR images. Beltrami regularized images appear less noisy than Tikhonov images.

Measures of directional information, as measured by the structure tensor metrics in the 18 subjects (volunteers and patients), are summarized in Figure 7.6. In-plane information (SA plane) was preserved between native SA images and SR images, as differences in $\langle S_{xx} \rangle$ ($p > 0.1$) and $\langle S_{yy} \rangle$ ($p > 0.6$) were not statistically significant. Information in the orthogonal direction was improved with SR reconstruction compared to the native images since $\langle S_{zz} \rangle$ was significantly different ($p = 0.0002$). Beltrami denoising did not seem to affect this resolution enhancement as




Figure 7.6: Directional information (derived from the structure tensor) contained in the native short-axis images and in the SR images (with Tikhonov and Beltrami regularization) in the short-axis plane (S_{xx} and S_{yy}) and in the orthogonal direction (S_{zz}). Differences between native and SR images were not significant for x and y directions and were significant for z direction.

differences in all three metrics between Tikhonov and Beltrami regularized images were not significant ($p > 0.5$).

7.6 Discussion and Conclusion

The proposed reconstruction combines nonrigid motion-correction and super-resolution. Our hypothesis was that SR reconstruction can be guaranteed if two conditions are met: (i) images with good quality and good motion-consistency can be provided as inputs of the SR algorithm, i.e. the input volumes need to have high in-plane resolution, minimal blurring/ghosting artifacts, and the misalignments from one volume to another should be minimal; (ii) an efficient SR solver can be implemented, i.e. efficient regularization schemes can be implemented in order to solve this under-determined inverse problem. For these reasons the joint reconstruction of image and

motion (GRICS approach) and the feature-preserving Beltrami regularization were key technical choices. Our results in healthy volunteers and patients with breathing difficulties show the feasibility of the technique.

Alternative methods for SR cine reconstruction include dictionary learning approaches [25, 16] which do not explicitly require motion correction. Our cine-GRICS approach has the advantage of using a motion model that inherently compensates intra-image motion (ghosting/blurring), intra-stack motion (slice-to-slice consistency) and inter-stack motion (stack-to-stack consistency). Compared to other total variation schemes [113] our Beltrami scheme yields similar results with faster convergence 4.

The proposed approach still has limitations that should be mentioned. Despite its high efficiency the b-SSFP sequence is sensitive to B_0 field inhomogeneity that may result in dark banding artifacts. Although the requirement in terms of B_0 homogeneity is less for the 2D b-SSFP than it is for its 3D version, it does rely on efficient shimming in the volume to be scanned. It should also be noted that both the in-flow effect and the location of the dark-band artifacts depend on the slice orientation and may result in outliers in the native images. This might be overcome with robust SR algorithms [67]. Another limitation is that the motion correction and the SR steps were applied sequentially. The final reconstruction might still be improved by merging the two steps into a single optimization problem. Such an approach would consist of searching for the isotropic SR image directly from the motion-corrupted k-space data.

In conclusion we have proposed a strategy for 3D isotropic cine MRI. The data acquisition consists of multiple stacks of 2D b-SSFP datasets with different orientations. Efficient SR reconstruction of an isotropic 3D cine dataset was rendered possible by (i) advanced nonrigid motion correction (using 3D cine GRICS) providing artifactfree and motion consistent images for each stack; (ii) regularized inversion of the SR model (with either Tikhonov or feature-preserving Beltrami regularization). The feasibility has been demonstrated in healthy volunteers and patients with breathing difficulties. The approach might be adapted to other imaging sequences or applications such as 3D late gadolinium enhancement.

7.7 Super-Resolution MRI and 3D Printing²

Synopsis

Three-dimensional (3D) printing technology is an emerging tool with significant clinical and prognostic implications for the management of interventional and surgical planning in congenital heart disease. Computed tomography has been established as an essential tool to provide adequate anatomical 3D models but is a very invasive procedure. 3D cardiac MR is another attractive imaging modality in terms of patient safety and excellent tissue differentiation, but suffers from poor gating efficiency, motion management and poor through-plane resolution, making high-resolution isotropic 3D imaging hard to achieve.

Our aim is to demonstrate the feasibility of creating 3D printed cardiovascular models derived from high-resolution free-breathing MR acquisition using recent motion-correction and super-resolution techniques. This will offer cardiologists, medical students and family members a better 3D anatomical understanding when planning interventions in patients with congenital heart malformations.

Acquisition

Subject 1: Ex vivo human heart (Super-Resolution only): 3D EFGRE pulse sequence with the following parameters: matrix size = 224 x 224 x 56 (reconstructed to 256 x 256 x 56), TR/TE = 18.3/8.9 ms, flip angle = 20°, slice thickness = 4 mm, resulting in a scanning time of about 3 x 3 min. The three acquired volumes had a resolution of 0.88 x 0.88 x 4 mm³.

Subject 2: Patient with Duchenne muscular dystrophy and major breathing difficulties (Motion-Correction + Super-Resolution): 2D multi-slice b-SSFP sequence with the following parameters: matrix size = 224 x 224, 20 frames, TE/TR = 1.8/4.1 ms, native in-plane resolution 1.6 x 1.6 mm², slice thickness =

²Presented in: Bustin A, Odille F, Menini A, Bonnemains L, Felblinger J, Hager G, Burschka D. MRI-Based 3D-Printed Cardiac Models using Super-Resolution Free-Breathing Reconstruction. BERTI Symposium Conference, Munich, January 2017

8 mm. The cine-GRICS motion correction step [cite] provided three cine volumes (with 32 cardiac phases). Acquisitions were performed during free-breathing and were ungated. Isotropic reconstruction was accomplished in mid-diastole.




Figure 7.7: Super-resolved MR reconstruction of the ex vivo human heart. The isotropic reconstruction provides noticeable anatomical details in every orientation with good contrast between blood and tissue.

Reconstruction and Processing

We reconstructed a single isotropic 3D volume with resolution $0.88 \times 0.88 \times 0.88 \text{ mm}^3$ (ex vivo, Figure 7.7) and $1.6 \times 1.6 \times 1.6 \text{ mm}^3$ (patient), using the super-resolution technique proposed in Section 7. The reconstructed high-resolution volume was then imported into a dedicated post-processing software (Seg3D, University of Utah), myocardium was segmented via thresholding and a 3D model was rendered prior to printing (Figure 7.8).

Results

Super-resolution techniques in free-breathing enables the reconstruction of 3D isotropic cine MRI in patients suffering with breathing difficulties. The 3D printed models (Figure 7.8 and Figure 7.9) illustrates the complex cardiovascular anatomy relevant for the morphology assessment: left and right ventricles, the ascending aorta and the




Figure 7.8: Volume rendering and slice selection prior to 3D printing.

brachiocephalic artery. The 3D model allows surgeons to identify the anatomy and possible congenital defects.

Conclusion

Three-dimensional cardiovascular models have the potential to improve the understanding of the morphology for cardiologists and surgeons before surgery. Patient safety and comfort may also be greatly improved by using MR acquisition in free-breathing combined with super-resolution. Additional comparative clinical experiments are necessary to evaluate the diagnostic accuracy of 3D printed models in more complex anatomy.




Figure 7.9: Physical 3D model of the ex vivo human heart and patient with Duchenne muscular dystrophy.

Chapter 8

3D Patch-Based Isotropic Reconstruction of MR Images¹

¹Submitted to: Bustin A, Voilliot D, Menini A, Felblinger J, Burschka D, Bonnemains L, Odille D. Isotropic Reconstruction of MR Images using 3D Patch-Based Self-Similarity Learning. Transactions on Medical Imaging, 2017

8.1 Introduction

Magnetic Resonance Imaging is a versatile imaging modality but its main limitation remains the need for choosing, at the acquisition stage, a compromise between signal/contrast-to-noise ratio, resolution and scan time. In clinical cardiac MRI, the scan time is most often constrained by the maximal breath-holding period achievable by the patient which is typically 20-25 seconds. As a result, acquiring 3D isotropic cardiac images with high spatial resolution (i.e. 1 to 1.5 mm voxel size) and with diagnostic quality is challenging. This would be particularly useful in late gadolinium enhancement (LGE) imaging which is the reference standard among all existing imaging modalities for assessing the location and extent of myocardial fibrosis in various cardiac diseases such as myocardial infarction. However, in the clinical practice, 3D LGE imaging has a high in-plane (1.5 mm) but low slice resolution (5 to 10 mm). Combining multiple such volumes (with low slice resolution) acquired in different orientations using a super-resolution reconstruction framework can overcome this limitation [90, 43].

Super-resolution (SR) reconstruction in MRI is now an established technique which has been successfully applied to diffusion-weighted imaging [116, 31], functional MRI [88], myocardial scar assessment [43, 7] and fetal brain MRI [48]. Super-resolution techniques can be classified into two main categories: geometrical methods and statistical methods.

Geometrical SR methods (also called multi-frame SR methods) make use of the physical coordinates of the acquired low-resolution (LR) data samples. These coordinates are varied between the multiple acquired datasets either by changing the volume orientation (the method that is used in this study) or by introducing a known shift (below the slice resolution) in the slice direction. These coordinates are integrated to the forward model of the acquisition system and the isotropic volume x is found by regularized inversion.

Statistical SR techniques (also called single-frame SR methods) aim at reconstructing a high-resolution (HR) image from a single acquired LR image using prior knowledge about the structure of the data, which can be obtained from the data

themselves or from an external database, e.g. by machine learning approaches. Such methods exploit redundancies in the image, e.g. relations between neighboring pixels, to estimate the missing data from the acquired data. Initially, such methods were built on the assumption that images were locally smooth or sparse in some transformed domain.

Recent research has been focused on the use of non-local similarity in the images [24, 73, 72]. Such techniques have been the subject of extensive work and applications, particularly in the area of dictionary learning [16, 104]. By using a learned database, patches in a LR image can be compared to external similar data, which may thus enforce sparsity within the similar-patches group and therefore overcome the shortcomings of local methods such as TV, i.e. over-smoothing in flat areas and staircasing. Although such techniques have shown great potential in MR reconstruction, dictionary learning is a large-scale and highly non-convex problem, which often requires high computational complexity and extensive prior training. Other than time complexity, similar patches are considered independently in the learning process, thus ignoring the relationships between similar patches within the image itself. Self-similarity using the well-known non-local means filter has recently been proposed for single-frame SR [31, 73, 72], where similar patches are extracted in a local neighborhood within an iterative framework. While those techniques have shown to be very powerful in single-frame SR, extraction and integration of self-similarity into a geometrical SR framework is not a trivial task and may require a more elaborate approach.

In this study, we present a novel technique that combines geometrical and statistical SR using a 3D self-similarity learning framework. The technique aims to achieve two goals: (i) merge information from multiple 3D anisotropic volumes acquired in different orientations into one single isotropic volume and (ii) recover sharp edges and thin anatomical structures using the self-similarity prior, which is hypothesized to be more robust in the presence of high noise levels and artifacts. These two goals are formulated as an optimization problem comprising a geometrical SR data consistency term and a low-complexity 3D patch-based regularization. The two corresponding sub-problems are solved iteratively into a simple but effective augmented Lagrangian

(AL) scheme. The method was evaluated in numerical simulations, and in actual MRI experiments with volunteer and patient data.

8.2 Theory

In this section, we first review the conventional geometrical SR reconstruction as well as dictionary learning frameworks. We then describe the detailed construction of the proposed 3D patch-grouping and its sparse representation. Then, an efficient cost functional is proposed which jointly estimates the group-sparse coefficients and the reconstructed isotropic volume within a dictionary learning framework. The proposed framework for joint self-similarity learning and geometrical SR reconstruction is illustrated in Figure 8.1.

8.2.1 Geometrical Super-Resolution

Given Z derived anisotropic measurements ρ_1, \dots, ρ_Z ($Z = 3$ in this study), representing the same object in arbitrary directions, affected by noise, and given the corresponding system matrices of the form $E_i = S_i B_i T_i$ ($i \in \{1, \dots, Z\}$), the isotropic volume x is found by solving the following regularized least squares optimization problem:

$$\min_x \sum_{i=1}^Z \|S_i B_i T_i x - \rho_i\|_2^2 + \lambda R(x) \quad (8.1)$$

Here T_i is a rigid image transformation that takes the SR image from the desired reconstructed orientation to the orientation of the i th acquisition (i.e. it is an interpolation operator that can describe an arbitrary orientation); $S_i B_i$ is a slice selection operator including a blurring operator B_i (i.e. a sum in the slice direction in the range of the slice thickness) and a downsampling operator S_i (in the slice direction). Here R is the chosen regularization function and λ is the corresponding nonnegative regularization parameter. The least squares term in (8.1) forces the reconstructed isotropic volume x to preserve a degree of fidelity with respect to the

acquired LR data, whereas the regularization term discourages aberrant values from appearing in the reconstructed volume. Previously studied regularization schemes include: Tikhonov, typically defined by $R_{TK}(x) = \|x\|^2$ or $R_{H^1}(x) = \|\nabla x\|^2$; edge-preserving regularizers such as TV regularization where $R_{TV}(x) = \|\nabla x\|$ or Beltrami, defined as $R_{BEL}(x) = \sqrt{1 + \beta^2 |\nabla x|^2}$, which offers a compromise between the smooth regularization of H^1 priors and the feature preservation of TV regularization.

8.2.2 Dictionary Learning

The general formulation of sparse and redundant representation considers a 3D signal x using the dictionary D and sparse vector α as the approximation $x \approx D\alpha$, satisfying $\|\alpha\|_0 \leq R$, where R is a predefined threshold and the l_0 -norm counts the number of nonzero elements in α . In other words, the signal x can be represented with a minimum number of coefficients α in the dictionary D . With sparsity prior and using the formulation (8.1), a volume x over dictionary D can be estimated by solving the l_0 minimization given as [1]

$$\min_{\alpha} \sum_{i=1}^Z \frac{1}{2} \|D\alpha - \rho_i\|^2 + \lambda \|\alpha\|_0 \quad (8.2)$$

In equation (8.2) the first term is known as the relation to measurements, and the other term represents the prior knowledge to the solution. The regularization parameter λ imposes the degree of closeness of the prior knowledge to the data. In existing SR models with sparsity prior, D is usually built from external HR trained dataset, making the assumption that a HR patch and its corresponding LR patches have the same sparse decomposition in their respective dictionaries, also called correlated dictionary [16, 125]. While being powerful, such techniques require high computational complexity as well as a rich dictionary to enable sparse patch approximations, and often ignore the relationships between patches, considering the reconstruction of each patch independently.

In this work, a novel 3D isotropic reconstruction algorithm is proposed, which exploits the redundancy of non-local 3D patches, in the acquired data itself, to build

the dictionary D . The isotropic reconstruction problem is seen as a joint low-rank and SR optimization, which is efficiently solved through simpler sub-problems using a variable splitting described hereafter.

8.2.3 3D Patch Representation

First, we define the 3D patch representation. In this work, we consider an unknown volume x of size $N = N_x \times N_y \times N_z$. A patch P_k^x , in the volume x , is defined as a small 3D block of size n^3 pixels around the pixel at indexes k . We define the operator $R_k : x \in R^{N_x \times N_y \times N_z} \mapsto R_k(x) = P_k^x \in R^{n^3}$ that extracts a 3D block from the volume x and its adjoint R_k^T such that $R_k^T P_k^x$ places back the patches at the original positions in x . Then, for a specific reference patch P_k^x , we search for all similar patches $P_{\bar{k}}^x$ in a large neighborhood H of size d^3 ($d > n$), with $\bar{k} \in \{k - d/2, \dots, k + d/2\}^3$. From the set of $(d - n + 1)^3$ 3D overlapped patches in H , we choose the L -most similar patches to the reference patch P_k^x . The similarity is measured by the l_2 -norm distance defined as $d(P_{\bar{k}}^x, P_k^x) = \|P_{\bar{k}}^x - P_k^x\|_2$ and the $L - 1$ patches $P_{k_l}^x$, $l = 1, \dots, L - 1$, with the smallest distance are selected. This selection, often called as block-matching, results in a sparse set of dimension 3D+L with lots of non-local similarities and can be analyzed in a very efficient manner with transformations such as multi-dimension Haar wavelet to produce sparsity in the case of 2D+L problems such as image denoising [34].

Here we choose to simplify the problem into an easier one by reducing the dimensionality of the similarity matrix by concatenating each vectorized 3D patch (representing one single column) into a 2D matrix (see Figure 8.1, Stage 1). This process is applied to each pixel in the volume x , therefore resulting in N groups x_k of size $(n^3 \times L)$, with

$$x_k = \left[P_k^x, P_{k_1}^x, \dots, P_{k_{L-1}}^x \right] \text{ for } k \in \{1, \dots, N\} \quad (8.3)$$

The 2D matrix x_k contains a high degree of non-local similarity and redundancy between intensity values and therefore exhibits a low-rank structure. A technique for low-rank and adaptive basis selection is singular value decomposition (SVD). The

SVD of the data matrix x_k is a decomposition of the form

$$x_k = \sum_{j=1}^L (\sigma_{x_k})_j (ux_k)_j (vx_k^T)_j = U_k S_k V_k^T \quad (8.4)$$

Where U_k is an orthogonal matrix of size $(n^3 \times n^3)$, V_k is an orthogonal matrix of size $(L \times L)$, with $U_k^T U_k = Id(n^3)$, $V_k^T V_k = Id(L)$ and S_k is a diagonal matrix of size $(n^3 \times L)$, with non-negative entries. Each element in dictionary D_k for group x_k is thus defined as $(d_k)_j = (ux_k)_j (vx_k^T)_j$, $j = 1, \dots, L$, and (σ_{x_k}) represents the sparse coefficients. Therefore, the element $(d_k)_j$ can be considered as an adaptive basis, and $(\sigma_{x_k})_j$ as the representation coefficient. The reconstructed group is then expressed by $x_k = D_k \alpha_k$ where D_k is the concatenation of all $(d_k)_j$ and α_k the concatenation of all $(\sigma_{x_k})_j$. Therefore, the volume x can be sparsely represented by the set of sparse coefficients $\{\alpha_k\}$:

$$x = D\alpha = \frac{\sum_{k=1}^L R_k^T (D_k \alpha_k)}{\sum_{k=1}^L R_k^T (I_k)} \quad (8.5)$$

Where D denotes the concatenation of all D_k , α denotes the concatenation of all α_k , and I_k is a matrix of the same size as $D_k \alpha_k$ with all elements being 1 (i.e. averaging matrix), for $k = 1, \dots, N$.

8.2.4 Proposed Joint Denoising and SR Cost Functional

Based on the observation that α_k is sparse when similar patches are grouped, we propose to reconstruct the MR volume by solving the following self-similarity SR problem in the group domain. Now, incorporating (8.5) into (8.1), our proposed scheme for volume SR using 3D Patch redundancy and SVD based self-dictionary is formulated as

$$\underset{\alpha}{\operatorname{argmin}} \sum_{i=1}^Z \frac{1}{2} \|ED\alpha - \rho_i\|_2^2 + \lambda \|\alpha\|_0 \quad (8.6)$$

By introducing a variable x , we transform the equation (8.6) into an equivalent constrained form

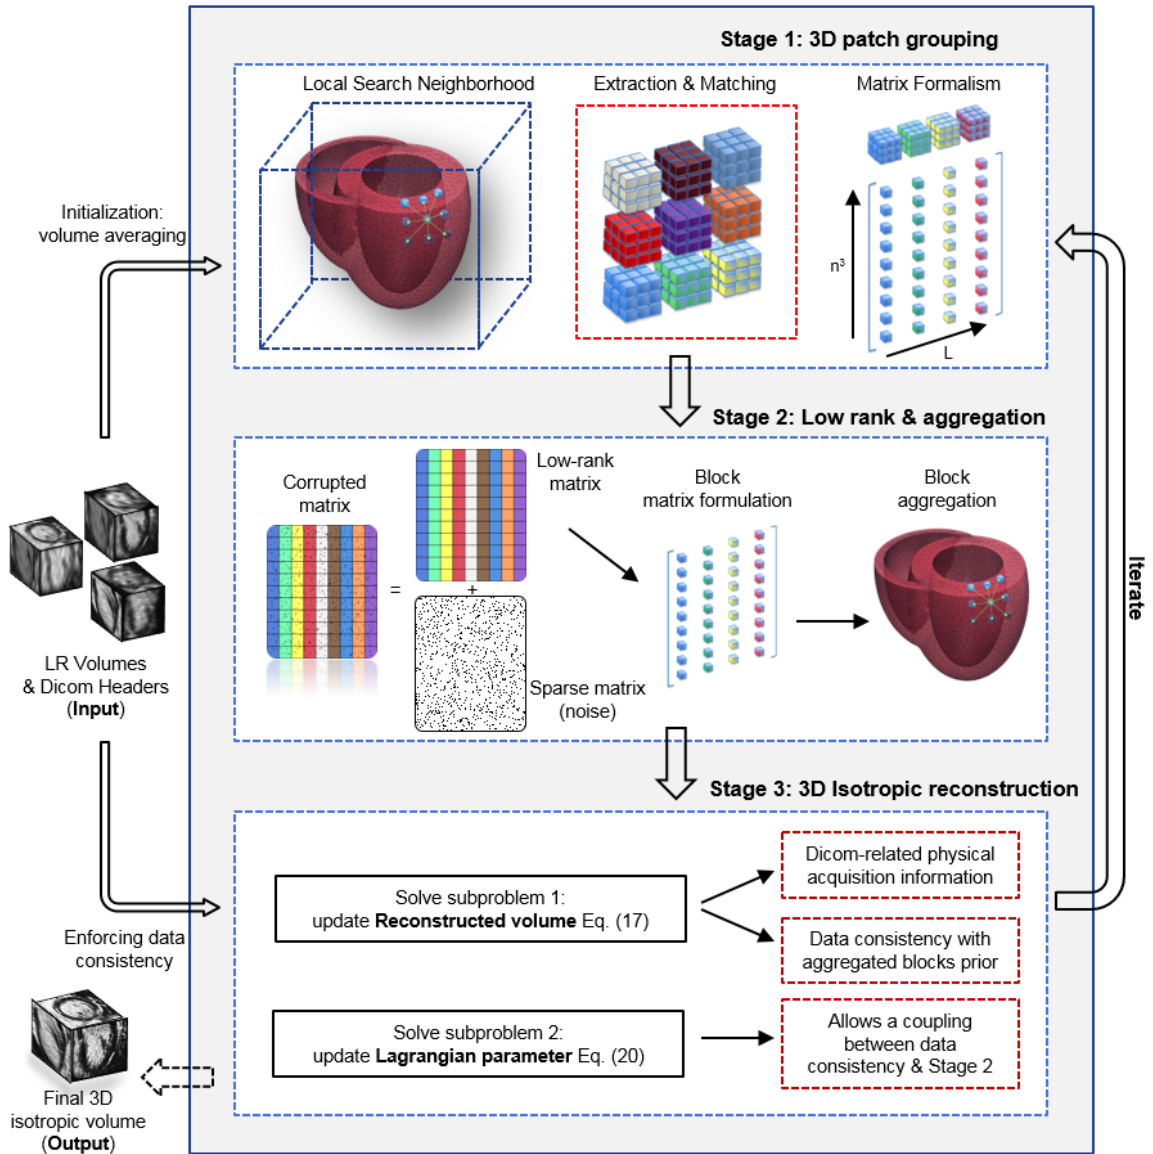


Figure 8.1: Flowchart of the proposed joint denoising and super-resolution reconstruction framework. The algorithm is initialized by averaging the acquired anisotropic volumes. The procedure starts with a self-similarity block-grouping followed by a low-rank step (denoising). The resulting matrix is used as a prior in a 3D super-resolution reconstruction. An augmented Lagrangian scheme is used to merge the denoising step and the isotropic reconstruction.

$$\operatorname{argmin}_{x,\alpha} \sum_{i=1}^Z \frac{1}{2} \|Ex - \rho_i\|_2^2 + \lambda \|\alpha\|_0 \text{ s.t. } x = D\alpha \quad (8.7)$$

The problem (8.7) is difficult to solve directly due to the constrained form of the functional. One strategy to solve this minimization problem is to approximate the equation (8.7) using its AL formulation. For a sake of simplicity and without loss of generality, we consider ρ as a concatenation of the acquired LR volumes $(\rho_i)_i, i = 1, \dots, Z$ and rewrite the problem (8.7) in its Lagrangian form [80] as

$$\mathcal{L}(x, \alpha) = \frac{1}{2} \|Ex - \rho\|_2^2 + \lambda \|\alpha\|_0 + \frac{\mu}{2} \|x - D\alpha - b\|_2^2 \quad (8.8)$$

Where E represents the acquisition model, D is the self-similarity dictionary and α the sparse coefficients. The parameter b denotes the AL multiplier of the original problem (8.7). The basic idea is now to solve the joint minimization problem efficiently through simpler sub-optimization problems by using a variable splitting. In our case, it consists in splitting the isotropic reconstruction term (minimization with regards to x) and the regularization term (minimization with regards to the sparse coefficients α).

8.2.5 Optimization

The augmented equation (8.8) is solved by alternating the optimization between the isotropic volume x and the sparse coefficients α , followed by an update of the augmented multiplier b .

Optimization 1: 3D Patch Based Denoising Update

The first sub-problem with regards to the variable α is

$$\mathcal{L}_1(\alpha) : \operatorname{argmin}_{\alpha} = \frac{\mu}{2} \|x - D\alpha - b\|_2^2 + \lambda \|\alpha\|_0 \quad (8.9)$$

It is obvious to see that (8.9) can be efficiently minimized by solving N sub-problems for all the groups. Each group-based sub-problem is thus formulated as

$$\underset{\alpha_k}{\operatorname{argmin}} = \frac{1}{2} \|x_k - b_k - D_k \alpha_k\|_2^2 + \lambda \|\alpha_k\|_0 \text{ for } k = 1, \dots, N \quad (8.10)$$

Setting $\tilde{x}_k = x_k - b_k = D_k \tilde{\alpha}_k$, where $\tilde{\alpha}_k$ represents the concatenation of the singular values of \tilde{x}_k , and using the unitary property of the SVD, we have:

$$(d_k)_j^T (d_k)_j = (v_{\alpha_k})_j (u_{\alpha_k}^T)_j (u_{\alpha_k})_j (v_{\alpha_k}^T)_j = (I_{\alpha_k})_j \quad (8.11)$$

and we can deduce:

$$\|\tilde{x}_k - D_k \alpha_k\|_2^2 = \|D_k \tilde{\alpha}_k - D_k \alpha_k\|_2^2 = \|\tilde{\alpha}_k - \alpha_k\|_2^2 \quad (8.12)$$

Equation (8.9) is thus equivalent to the following equation

$$\mathcal{L}_1^k(\alpha) : \alpha_k^* = \underset{\alpha_k}{\operatorname{argmin}} = \frac{1}{2} \|\tilde{\alpha}_k - \alpha_k\|_2^2 + \lambda \|\alpha_k\|_0 \quad (8.13)$$

Where the l_0 norm counts the number of non-zeros elements in α_k , and $\lambda > 0$ controls the strength of sparseness. The lower the parameter λ is, the more accurate the solution is, at the price of reducing the sparseness. The solution of (8.13) has a well-known closed form, and we get the optimum as

$$\alpha_k^* = H_{\sqrt{2\lambda}}(\tilde{\alpha}_k) \quad (8.14)$$

Where $H_\theta(\cdot)$ is the element-wise hard thresholding operator, defined for a scalar v as: $H_\theta(v) = \theta \cdot 1_{|v| > \theta}$. In other words, any singular value below $\sqrt{2\lambda}$ is set to 0. Note that only the singular values are modified, but the singular vectors are unperturbed. The value α_k^* is indirectly an estimate of the noise $x_k - b_k$, and thus the hard-thresholding step effectively shrinks the noise, i.e. small differences are accounted as features, and only large impulses are considered as noise.

Optimization 2: Super-Resolution Update

The second sub-problem with regards to the variable x is

$$\mathcal{L}_2(x) : \underset{\alpha}{\operatorname{argmin}} = \frac{1}{2} \|Ex - \rho\|_2^2 + \frac{\mu}{2} \|x - D\alpha - b\|_0 \quad (8.15)$$

Differentiating with respect to x , we find that the residual gradient step is

$$d = E^T Ex - E^T \rho + \mu(x - D\alpha - b) \quad (8.16)$$

Where the subscript T denotes the transposed conjugate. We use the gradient descent optimization method to iteratively update the reconstructed volume x :

$$x^* = x - \beta d \quad (8.17)$$

Where the relaxation parameter β can be updated iteratively or be set to a specific value (e.g. $\beta = 0.1$) to ensure convergence. Notice that this optimization step requires only one multiplication of $E^T \rho$, which can be determined during the initialization step, and only one multiplication $E^T Ex$ per iteration. In (8.16), $D\alpha$ represents a truncated SVD reconstruction where only the singular values above a specified threshold are included, leading to an estimated low-rank matrix. The Lagrangian parameter is then updated at iteration $(t + 1)$ as

$$b^{(t+1)} \leftarrow b^{(t)} + r(D\alpha^{(t+1)} - x^{(t+1)}) \quad (8.18)$$

which only involves an additional straightforward computation. The addition of a new variable b shares some similarities with the ROF model, proposed for image denoising in [86], and can be interpreted as iteratively adding back the residual noise. The accuracy of the patch grouping is highly affected by the noise in the volume $x^{(t)}$ at iteration t and can result in noise residual in the reconstructed volume $x^{(t+1)}$ at iteration $t + 1$. Therefore, optimization (8.13) and (8.15) are processed iteratively to improve the accuracy of the reconstructed volume. A flowchart illustrating all steps is shown in Figure 8.1 and the full algorithm is summarized in Algorithm 6.

Algorithm 6 3D isotropic MRI reconstruction using joint denoising and super-resolution optimization

```

1: Input:  $\rho, \mu, \lambda$ , outer loop iter, operator  $E$ 
2: Initialization:  $x^{(0)} = E^T(\rho); b^{(0)} = 0; \beta = 0.1;$ 
3: for  $t = 0, \dots, \textit{iter}$  do do
4:    $r^{(t)} = x^{(t)} - b^{(t)}$ 
5:   for  $k = 1, \dots, N$  do do % Solve  $\mathcal{L}_1^k$ 
6:      $r_k^{(t)} = \left[ P_k^{r(t)}, P_{s_1}^{r(t)}, \dots, P_{s_{L-1}}^{r(t)} \right]$ 
7:      $\left[ U_k^{(t)}, S_k^{(t)}, V_k^{(t)} \right] = \text{svd} \left( r_k^{(t)} \right)$ 
8:      $\omega_k^{(t+1)} = U_k^{(t)} H_{\sqrt{2\lambda}} \left( S_k^{(t)} \right) V_k^{T(t)}$ 
9:   end for
10:   $\omega^{t+1} = \frac{\sum_{k=1}^N R_k^T(\omega_k^{(t+1)})}{\sum_{k=1}^N R_k^T(I_k)}$ 
11:   $\tilde{x}_0^{(t)} = x^{(t)}$ 
12:  while Not Converged do do % Solve  $\mathcal{L}_2$ 
13:     $d^{(t)} = E^T E \left( \tilde{x}_i^{(t)} \right) - x^{(0)} + \mu \left( \tilde{x}_i^{(t)} - \omega^{(t+1)-b^{(t)}} \right)$ 
14:     $\tilde{x}_{i+1}^{(t)} = \tilde{x}_i^{(t)} - \beta d^{(t)}$ 
15:  end while
16:   $x^{(t+1)} = \tilde{x}^{(t)}$ 
17:   $b^{(t+1)} = b^{(t)} + \left( \omega^{(t+1)} + x^{(t+1)} \right)$ 
18: end for
19: Output: the isotropic 3D volume  $x$ 

```

8.3 Methods

8.3.1 Implementation

Reconstructions and evaluations were performed in Matlab (The MathWorks, Natick, MA) on a 32-core 3.2 GHz Intel Xeon E5-2667 v3 PC with 64 GB of RAM. The 3D patch grouping and singular value thresholding steps of the proposed method (stage 1 and 2 on Figure 8.1) were implemented in C to reduce the computational time.

Isotropic reconstruction from arbitrary volume orientations (i.e. with non-strictly orthogonal views) requires an accurate knowledge of patient position and system imaging configuration to avoid any alignment errors. This information was extracted from the images' DICOM file headers as recorded by the imaging system, including Spacing Between Slices, Image Orientation Patient (IOP) and Image Position Patient (IPP), and was used to construct the operator E . Patient motion for the cardiac MRI scans was handled by rigid registration as described hereafter in Section 8.3.4.

One of the key advantages of the proposed technique is the possibility of multithreading implementations for algorithm speedup. As seen in Section 8.2.4, the low-rank optimization is applied on each group x_k independently and therefore, multithreading can efficiently and easily be implemented for the patch search and hard-thresholding step. For our experiments, we divided the 3D volume into 32 sub-volumes, each of them being treated separately by one processor, thus dividing the processing time by a factor of 32. Besides the multithreading computing, patches overlapping can also be performed as a means of acceleration. Indeed, one can use a step d in the three dimensions and therefore perform the patch matching step from one reference patch P_k^x to the next P_{k+p}^x (located at a distance p pixels from the previous reference patch). In such a context, a step of $p = 4$ would theoretically divide the processing time by $4^3 = 64$, which is not negligible for large volumes. However, a large overlap step p would decrease the SNR and a tradeoff between acceleration and image quality/accuracy must be defined. Choosing the right values for the different parameters in the algorithm and their influence has to be discussed and is provided as supplementary material. It appears that the proposed method is almost parameter-free in the sense that by tuning the parameters in a specific range, the

isotropic reconstruction remains optimal.

In all experiments the proposed SR reconstruction was compared to that obtained with a conventional geometrical SR reconstruction using either $R_{TK}(x)$ or $R_{BEL}(x)$ regularization as defined in Section 8.2.1. For comparison with Beltrami regularization, we propose to solve (8.1) by adopting a primal-dual gradient-descent, as defined in Chapter 4 and performed in Chapter 5 and Chapter 6. In the case of Tikhonov regularization, the SR algorithm is a linear least-squares problem and is solved by using a conjugate gradient solver. For all the methods, the regularization parameter λ was tuned in order to obtain the best SNR in a selected region of interest. For clinical cases, where the ground truth is not available, we chose to visually optimize λ , but it is possible to numerically tune this parameter using more sophisticated techniques, such as the L-curve method [55]. The proposed AL algorithm is initialized by averaging the acquired volumes ρ_i , which is fast and practical since this averaging is then used at each iteration when solving the SR step \mathcal{L}_2 . This is accomplished using the transposed operator $E^T : x^{(0)} = E^T \rho = \frac{1}{Z} \sum_{i=1}^Z E^T \rho_i$.

8.3.2 Experiments on Numerical Phantom Data

In this section, we evaluated the proposed method on a synthetic phantom study. We first acquire a HR volume of a MR phantom with isotropic resolution $1 \times 1 \times 1 \text{ mm}^3$. A 3D enhanced fast gradient echo (3D EFGRE) sequence was used with the following parameters: matrix size = $256 \times 256 \times 256$, TR = 5.68 ms, TE = 2.20 ms, flip angle = 15° , acquisition time = 6.2 min, on a 3T GE Signa HDxt system (GE Healthcare, Milwaukee, WI). The simulation consisted of retrospectively downsampling the $100 \times 100 \times 100$ central cube of this HR volume into three LR orthogonal volumes with slice thickness in the range of 3 to 9 mm in the through-plane direction. Rician noise was generated and added to the anisotropic volumes at a level of $\sigma = (1, 3, 5, 7, 9) \%$ the maximum intensity. The proposed technique was then applied to produce a single isotropic 3D volume (resolution of $1 \times 1 \times 1 \text{ mm}^3$) with the following parameters: $n = 10 \times 10 \times 10$, $d = 4$, $L = 20$, $\lambda = 1$, $p = 4$, $iter = 4$.

The goals of the present experiment were threefold. First, we studied the ability of the proposed reconstruction technique to produce an isotropic 3D volume from very low-resolution volumes, even in the case of highly anisotropic voxels of 3 to 9 mm through-plane (versus 1 mm in-plane). Second, we investigated the robustness of the current approach against noise and local intensity variations, especially for high amplitude of noise σ . The impact of all the reconstruction parameters on image quality and SNR was studied, and optimal results were compared to the ground-truth HR volume, free of noise. Several measures were used to quantify the performances of the method. The first one is the intensity root mean square error (RMSE) defined as:

$$RMSE(HR_{ref}, HR_{rec}) = \sqrt{\frac{1}{N} \|HR_{ref} - HR_{rec}\|_2^2} \quad (8.19)$$

Where N is the number of pixels in the 3D volume and $\|\cdot\|_2$ stands for the usual l_2 norm. Additionally, the peak-signal-to-noise ratio (PSNR) between the reference HR and the reconstructed volume, evaluated in decibels, was computed as:

$$PSNR(HR_{ref}, HR_{rec}) = 20 \log_{10} \frac{\max(HR_{rec})}{RMSE(HR_{ref}, HR_{rec})} \quad (8.20)$$

Besides PSNR, which can accurately measure the local intensity difference between two volumes, we also employed another image quality measurement based on image sharpness, called Sharpness Index (SI) [18]. SI has been shown to provide a sensitivity to image sharpness with crucial information about the image geometry and image edges and therefore provides an interesting indicator of image sharpness:

$$SI(HR_{rec}) = -\log_{10} \phi \left(\frac{m - TV(HR_{rec})}{v} \right) \quad (8.21)$$

Where $m = \mathbb{E}[TV(HR_{rec})]$ is the expectation of the total variation of the reconstructed HR volume, and $v = Var[TV(HR_{rec})]$ is the corresponding variance. We define the sharpness index error (SIE) between the reference HR and the reconstructed volume as:

$$SIE(HR_{ref}, HR_{rec}) = |SI(HR_{ref}) - SI(HR_{rec})| \quad (8.22)$$

8.3.3 Experiments on Healthy Volunteer (Brain Imaging)

The proposed isotropic reconstruction was then used for brain MRI experiments to study the impact of the reconstruction on actual volunteer data. Three anisotropic acquisitions in axial, sagittal and coronal planes were sequentially collected in one subject with a EFGRE pulse sequence. The study was approved by an ethics committee (ClinicalTrials.gov Identifier: NCT02887053). One volunteer was scanned with a 3T Signa HDxt system (General Electrics, Milwaukee, USA) with an 8-channel head coil. The acquisition parameters were the following: matrix size = $224 \times 224 \times 56$ (reconstructed to $256 \times 256 \times 56$), TR = 10.2 ms, TE = 4.6 ms, flip angle = 20° , slice thickness = 4 mm, resulting in a scanning time of about 3×2 min. The three acquired volumes had a resolution of $1 \times 1 \times 4 \text{ mm}^3$ ($0.88 \times 0.88 \times 4 \text{ mm}^3$ after zero-filling) and thereby the resolution in the slice direction was more than 4 times higher than the in-plane resolution.

Impact of the denoising was evaluated by adding noise at a level of 5% the maximum intensity. Additionally, a HR volume of the same object was acquired at a resolution of $1 \times 1 \times 1 \text{ mm}^3$ ($0.88 \times 0.88 \times 0.88 \text{ mm}^3$ after zero-filling) and was used as a reference for visual comparison of brain structures. We reconstructed a single isotropic 3D volume with resolution $0.88 \times 0.88 \times 0.88 \text{ mm}^3$ using the proposed technique, with the same parameters as in the previous Section 8.3.2, as well as local Beltrami and Tikhonov reconstructions. In order to push imaging resolution beyond the practical limits of the system, we acquired three additional 3D volume of resolution $0.5 \times 0.5 \times 2 \text{ mm}^3$ (matrix size = $448 \times 448 \times 28$) located close to the brain cerebellum, resulting in a total scan time of $3 \times 3.8 \text{ min} = 11.4 \text{ min}$. An isotropic 3D volume was then reconstructed at a resolution of $0.5 \times 0.5 \times 0.5 \text{ mm}^3$. For visual comparison, we additionally reconstructed an isotropic volume using a simple averaging of the three acquired volumes, as it has been done in previous work [65]. Results were compared to the reference volume when available and the sharpness index was

used as an indicator of image sharpness.

8.3.4 Experiments on Patients (Cardiac Imaging)

In this study, we show how the developed reconstruction framework can be used for cardiac imaging and myocardial fibrosis assessment in LGE MRI.

Thirteen patients (14 men, mean age = 56 ± 13 years, range 37 - 78 years) with a remote (> 6 months) myocardial infarction underwent a cardiac MRI study for myocardial viability assessment before implantation of cardioverter-defibrillators. These data were acquired as part of a retrospective study from clinical MRI examinations with non-opposition of the patients for the use of the data for research purpose, which does not necessitate approval from an ethics committee according to the French law. Additionally, three patients with Duchenne muscular dystrophy were recruited and underwent conventional LGE imaging. These patients gave written informed consent and were included as part of a study which was approved by the local Institutional Review Board. All imaging studies were performed on a 1.5T Signa HDxt system with an 8-channel cardiac coil. Three conventional breath-held 3D LGE were performed in short-axis, horizontal long axis (HLA) and vertical long axis (VLA) views to delineate post-myocardial infarction scars. The scan parameters of the 3D LGE are provided in Table 8.1. Images were collected 15 minutes after Gadolinium injection with the following typical acquisition parameters: TR = 3.46 ms, TE = 1.28 ms, TI adjusted to cancel healthy myocardial signal, acquisition matrix size = $256 \times 256 \times 20$, acquisition time = ≈ 20 sec x 3 breath-holds, 1.25 mm in-plane resolution, 4.5 to 8 mm slice thickness. When needed, inconsistencies in heart position, due to different apnea amplitudes between the three acquisitions, were corrected using a manual rigid image registration framework. The registration was applied before the isotropic reconstruction. The three methods (3D Patch, Tikhonov and Beltrami) were used to reconstruct a single 3D isotropic volumes of resolution $1.25 \times 1.25 \times 1.25$ mm³. The same reconstruction parameters were used as in the previous experiments except for the patch size $n = 15$. The reconstructed volumes were evaluated for quality, contrast-to-noise ratio (CNR) and scar quality by two experienced cardiologists (D.V. and L.B.) blinded

to the reconstruction methods. Each observer viewed the three reconstructions for each patient in a random order and rated the images using the following ranking: each image was attributed the value of 1: Best, 2: Average, 3: Worst; this rating was exclusive, meaning that each rank (1-3) was assigned to a method with no repetition, therefore promoting the differentiation of each reconstruction. Cohen’s kappa coefficient was calculated to determine the inter-observer variability.

8.4 Results

8.4.1 Parameter Optimization

Like most image reconstruction algorithms, the performance of our technique relies on several parameters that needs to be carefully tuned in order to get the best reconstruction performance. In the following, we optimize, on simulation, the parameters that contribute significantly to the good performance of the method. Choosing the right values for the different parameters in the algorithm and their influence has to be discussed. The central parameters of interest are the size of patch n (default value: $10 \times 10 \times 10$), size of neighborhood window d (default value: 4), number of selected patches L (default value: 20), regularization parameter: λ (default value: 1), overlap: p (default value: 4), number of outer iterations: $iter$ (default value: 4). In each experiment, we let one parameter vary while remaining the others constant, with default values. We first evaluate the algorithm for the case of a slice thickness equal to 5 mm, and varying the noise level from $\sigma = 3\%$ to $\sigma = 11\%$ which correspond to typical values seen in clinical practice.

Influence of the size of patch n

The size of patch n is an important factor in our proposed isotropic 3D reconstruction and therefore a robust choice of this parameter would guarantee a high reconstruction performance. This parameter controls the degree of structural information within each patch. On one hand, a large value of n would capture the most geometric information and leads to a higher level of denoising, while a smaller value of n would

Table 8.1: Recommended 3D LGE protocol on a GE system.

Patient Position	
Patient position	supine
Patient entry	feet first
Coil	cardiac coil, torso coil
Series description	3DLGE hearty
Scanning Range	
Scan plan	Oblique
Frequency FOV	30-40 cm
Phase FOV	0.7-0.8
Slice thickness	2-4 mm
Spacing	0
Slab locations	40
Imaging Parameters	
Plane	Oblique
Mode	3D
Family	Gradient echo
Pulse sequence	Fast SPGR
Imaging options	Cardiac gating, IR prep, ASSET, Triggering, PG
Details	
TE	Minimum Full
Flip angle	10-15
Prep time	appropriate time from CINE IR
Acceleration factor	1.5
Frequency	224-256
Phase	192
Bandwidth	25 kHz (1.5T), 50 kHz (3.0T)
Shim	auto
Temporal resolution	≤ 150 ms
Cardiac Gating	
# RR interval	4
Trigger window	10
Trigger delay	appropriate delay from CINE scan

Influence of the size of the neighborhood window d

Figure 8.2b shows the influence of the size of the search neighborhood window on reconstruction quality for different values of noise. To investigate the effect of the parameter d on the PSNR, we conducted experiments by varying d from 1 to 8. Varying the size d does not seem to alter the reconstruction quality, with similar PSNR for d greater than 3. Increasing d would also increase the number of search patch and therefore increase the computation time. Accordingly, we decided to choose a value of d equal to 4 for all the experiments.

Influence of the regularization parameter λ

In this section, we analyze the performance of the proposed reconstruction algorithm when varying the similarity threshold λ . This parameter controls the sparsity of the 2D group x_k , and therefore the quality of the denoising. The noise attenuation is performed with the hard-thresholding step (16) by shrinkage of the singular values of x_k . Hence, when the noise amplitude is low, the grouping step is less affected by noise perturbation and therefore a low value of λ can be used for the hard-thresholding. On the other hand, for high level of noise, a larger regularization can be used to correct for any noise residual during the grouping stage. Figure 8.3 shows the influence of the regularization parameter when varying λ for different level of noise ($\sigma = 1 - 11\%$). As expected, for low level of noise ($\sigma = 1\%$), a low value of regularization ($\lambda = 0.01$) gives the best PSNR. As one can see, when the image is affected by more noise ($\sigma > 1\%$), equals to 1 gives the best performance in terms of PSNR. We therefore chose the value of $\lambda = 1$ for all the experiments presented in Sections 8.4.2-8.4.4.

Influence of the number of selected patches L

Here, we study the impact of the number of selected patches L on the reconstruction quality. We varied this parameter in a range of 1 and 100 for different levels of noise. The results of the experiments are shown in Figure 8.3b. As we can see in Figure 8.3, there is no optimal value that achieves the highest PSNR by varying the size of L for all levels of noise. That means that the proposed technique is not sensitive to the




Figure 8.3: Effect of the similarity threshold λ (a) and the number of selected patches L (b) on the PSNR for different values of noise ($\sigma = [1, 3, 5, 9, 11]$ %) for the isotropic reconstruction of the phantom data with retrospective slice thickness = 5 mm.

number of selected patches. Additionally, a large number of selected patches would increase the computation time. Accordingly, in the experiments, we empirically set the value of L equal to 20.

Influence of the overlap p

Another important parameter of the proposed isotropic reconstruction is the choice of the overlap parameter p . A particular choice of this parameter allows for a gain in computational speed, which is crucial for some real-time imaging applications, and has to be chosen based on the expected computational speed and reconstruction quality. To evaluate the impact of the overlap on PSNR, we applied the proposed reconstruction with reference patches separated from each other by an offset $p = 2, \dots, 10$ voxels in the three dimensions. Additionally, we recorded the consumed CPU time of the reconstruction for each evaluation. Figure 8.4 gives the results of the comparisons in terms of PSNR and CPU time over the overlap p . For each level of noise, we observe a very slow decrease of PSNR with respect to the offset p (Figure 8.4b). On the other hand, we observe a significant decrease of CPU time when the value of p grows (Figure 8.4a), suggesting that the small loss of SNR is negligible compared to the gain in reconstruction speed. As expected, the CPU time curves, shown in Figure 8.4a, follow a cubic decay of the form $p \mapsto 1/p^3$. Therefore, we




Figure 8.4: Effect of the overlap size p on elapsed CPU time (a) and PSNR for different values of noise ($\sigma = [1, 3, 5, 7, 9, 11]$ %) for the isotropic reconstruction of the phantom data with retrospective slice thickness = 5 mm.

demonstrate that there is a clear computational advantage of using an overlap for the reconstruction. As the loss in PSNR (Figure 8.4b) is small compared to the gain in acceleration, we chose a value of $p = 4$ in all experiments. The Figure 8.5a shows that the proposed methods take less than 3 minutes to reconstruct the isotropic 3D volume on this architecture, with default parameter values, $p = 4$, and matrix size = $100 \times 100 \times 100$.

Influence of the number of iterations $iter$

As mentioned in Section 8.2.5, the isotropic volume is iteratively reconstructed to account for the residual noise which propagates during the patch-grouping process. Here, we evaluate the impact of the number of iteration on the image quality for five different levels of noise, with a range of iterations from 1 to 10. Figure 8.5 shows the results of increasing the number of iterations on the PSNR. As expected, we observed an increase of the PSNR as function of the number of iterations. After the first two iterations, there is a significant gain in PSNR for each reconstruction (about 13 dB for $\sigma = 1\%$ and 18 dB for $\sigma = 11\%$). Furthermore, when $iter$ is greater than 6, we note that increasing the number of iterations in the reconstruction process does not seem to affect the PSNR, for all level of noise. The singular value shrinkage achieves a good estimation of the low-rank group by taking only the largest singular

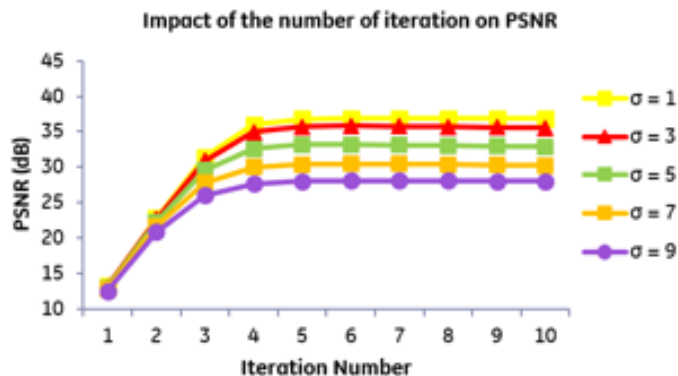


Figure 8.5: Effect of the number of outer iterations on the PSNR for different values of noise ($\sigma = [1, 3, 5, 7, 9, 11]$ %) for the isotropic reconstruction of the phantom data with retrospective slice thickness = 5 mm.

values, making most of the acquisition noise removed in the first few iterations. As a result, the number of iterations is independent of the noise level and a small number of iterations (between 4 and 6) is enough to get an accurate reconstruction and avoid useless computations. Accordingly, we choose a value of 4 iterations for all experiments as a good trade-off between reconstruction quality and computational speed.

8.4.2 Experiments on Phantom Data

Table 8.2 show the reconstruction performance of the proposed 3D patch-based reconstruction. These results show that the proposed reconstruction outperforms the existing techniques in terms of PSNR, RMSE and sharpness. For high level of noise, the average PSNR of the proposed technique is 26.21 dB, showing improvements of about 4.3 dB and 6.4 dB over Beltrami and Tikhonov respectively. Although the piecewise structure of the phantom seems better suited for gradient-based regularization, our 3D patch-based reconstruction always outperforms Beltrami, even for low level of noise. The 3D patch-based reconstruction improved the sharpness for different slice thickness values (slice thickness from 3 to 9 mm) and under different noise amplitudes. Visual comparisons of the three techniques, when varying the slice thickness

Table 8.2: Comparison of PSNR (P), RMSE (R) and Sharpness index (S) for the denoising of the numerical phantom (Simulation study) with different values of noise.

		Slice Thickness = 3 mm						Slice Thickness = 5 mm					
		$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 7$	$\sigma = 9$	$\sigma = 11$	$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 7$	$\sigma = 9$	$\sigma = 11$
NLM	P	35.47	28.08	24.69	22.43	20.77	19.49	33.97	27.55	24.44	22.40	20.87	19.68
	R	0.027	0.063	0.093	0.120	0.145	0.168	0.032	0.067	0.095	0.121	0.144	0.165
	S	8.16	14.15	16.92	18.92	20.32	20.97	8.06	14.5	17.05	18.6	20.27	20.06
BELTRAMI	P	37.98	29.99	26.83	24.56	23.39	21.60	36.48	29.99	26.69	24.54	23.46	21.62
	R	0.020	0.050	0.072	0.094	0.108	0.132	0.024	0.05	0.074	0.094	0.107	0.132
	S	5.39	12.25	14.51	16.93	18.36	19.83	6.39	11.59	14.71	16.74	18.53	19.42
3D PATCH	P	39.32	36.65	33.65	30.61	28.08	26.01	38.41	36.48	33.61	30.69	28.20	26.09
	R	0.017	0.023	0.033	0.047	0.063	0.080	0.019	0.024	0.033	0.046	0.062	0.079
	S	3.84	0.53	0.07	0.74	1.58	3.26	3.62	0.55	0.04	0.43	2.63	3.02
		Slice Thickness = 7 mm						Slice Thickness = 9 mm					
		$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 7$	$\sigma = 9$	$\sigma = 11$	$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 7$	$\sigma = 9$	$\sigma = 11$
NLM	P	33.65	27.44	24.53	22.57	21.11	19.95	33.51	27.61	24.77	22.85	21.44	20.30
	R	0.033	0.067	0.094	0.118	0.140	0.160	0.034	0.066	0.092	0.114	0.135	0.153
	S	8.25	14.52	17.03	18.68	20.25	20.06	8.17	13.69	15.78	18.90	19.90	20.24
BELTRAMI	P	35.74	30.00	26.91	24.80	23.73	21.91	35.23	30.28	27.30	25.25	22.99	22.34
	R	0.026	0.050	0.072	0.091	0.103	0.128	0.028	0.049	0.069	0.087	0.113	0.121
	S	6.32	11.03	14.13	16.61	18.28	18.49	6.54	10.89	12.98	16.05	18.63	18.66
3D PATCH	P	37.96	36.13	33.52	30.73	28.30	26.24	37.61	35.84	33.63	30.90	28.55	26.49
	R	0.020	0.025	0.033	0.046	0.061	0.077	0.021	0.026	0.033	0.045	0.059	0.075
	S	2.94	0.87	0.04	1.83	2.07	2.84	2.42	0.68	0.55	0.61	1.99	2.85

and keeping noise constant ($\sigma = 11\%$), are shown in Figure 8.6. For each method, the ground truth is shown on the right column (d). Interestingly, the reconstruction performances are always in favor of the proposed 3D patch-based reconstruction for the four simulated slice thicknesses. Beltrami reconstruction seemed efficient for small slice thicknesses (slice thickness < 5 mm), showing satisfactory noise reduction and edge preservation, but failed to recover a faithful image for low resolution experiments (slice thickness > 5 mm), significantly increasing blurring and residual artifacts. However, the proposed 3D patch-based reconstruction was able to minimize such artifacts, to reduce the noise and to reveal fine details and structures, even for large slice thicknesses (e.g. 9 mm), without exhibiting staircasing artefacts on flat areas.




Figure 8.6: Comparison of the proposed isotropic 3D reconstruction algorithm (c) with Tikhonov (a) and Beltrami (b) reconstructions applied to a high resolution isotropic phantom image (d) acquired on a 3T system and retrospectively undersampled with different slice thickness (3, 5, 7 and 9 mm) and a noise level $\sigma = 5\%$.

8.4.3 Experiments on Healthy Volunteer (Brain Imaging)

Figure 8.7 shows the visual and sharpness results of the healthy brain study. The sub-figures (a-c) are the three native acquisitions; sub-figures (d-f) are the reconstruction results by the averaging, Tikhonov and Beltrami methods, respectively; (g) is the proposed 3D patch-based reconstruction and (h) is the HR reference. The proposed reconstruction results in better visual image quality, recovers more details and enhances edges, while reducing the noise down to an acceptable level. Overall, the SI shows good results with the proposed reconstruction and correlates well with the visual human perception. Compared to Beltrami regularization, our technique improves the sharpness of the reconstructed volume by over 152% for the axial view (123% and 126% for sagittal and coronal views, respectively).

Figure 8.8 depicts the zoomed-in and cropped isotropic reconstructions on the brain ventricle by the different methods. Although Tikhonov is able to reconstruct a HR volume while reducing the blur observed in the averaged reconstruction, it fails to remove the noise of the acquired images and exhibits residual artifacts. On the other hand, Beltrami regularization is able to reduce the noise, at the cost of producing staircasing effects, often associated with l_1 -minimization problems, mainly due to the high level of noise. By exploiting the non-local similarities and geometries of the 3D volume, the proposed method clearly outperforms Tikhonov and Beltrami regularizations not only on flat areas (e.g. white matter) by removing the staircasing effect, but also on textured areas, with well-preserved structures, edges and image sharpness (e.g. brain ventricle). The results shown in Figure 8.9 for the brain cerebellum reconstruction are consistent with the simulation study, the patch-based reconstruction eliminates artifacts and suppresses a large amount of noise and possible local intensity variations, while preserving the main features and especially the small details (e.g. cerebellum sulcus).

8.4.4 Experiments on Patients (Cardiac Imaging)

In Figure 8.10, we compare the reconstructed isotropic ($1.25 \times 1.25 \times 1.25 \text{ mm}^3$) volume of a 20-year-old Duchenne patient with focal fibrosis, with the three native




Figure 8.7: Visual quality comparison of isotropic reconstruction on Brain imaging ($0.88 \times 0.88 \times 0.88 \text{ mm}^3$). From top to bottom: axial, coronal and sagittal views, reconstructed by simple averaging (d), Tikhonov (e), Beltrami (f), and the proposed 3D patch-based reconstruction (g). (h) Displays the acquired high resolution reference. Numbers in the top left corners represent the corresponding sharpness score SI.

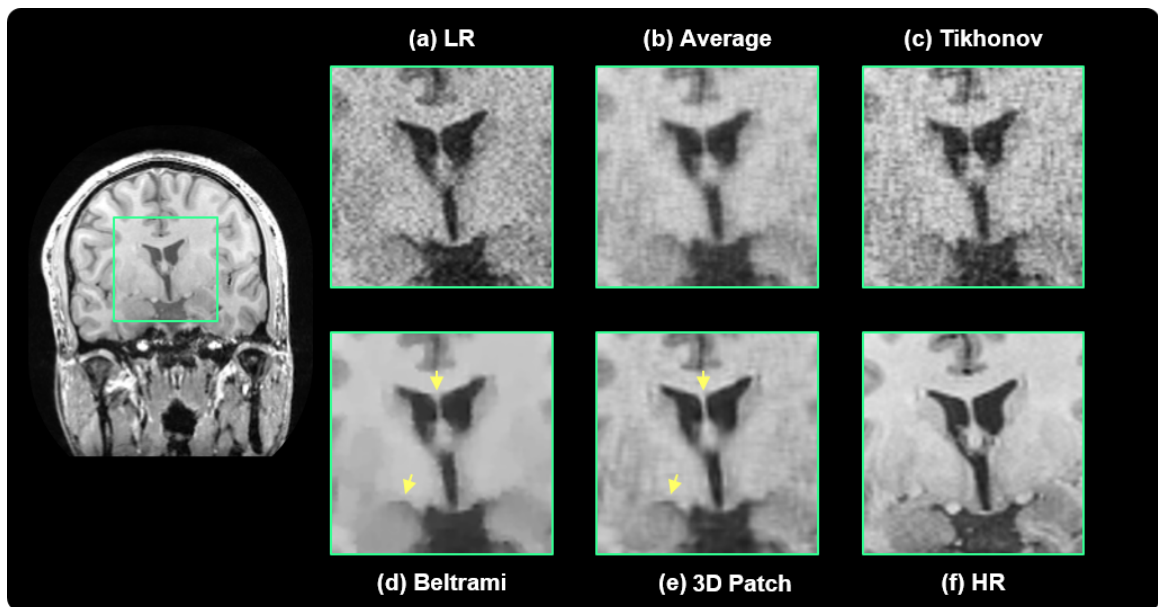


Figure 8.8: Comparison of the isotropic 3D reconstruction of the Brain ventricle ($0.88 \times 0.88 \times 0.88 \text{ mm}^3$) from LR acquisitions ($0.88 \times 0.88 \times 4 \text{ mm}^3$) with Averaging (b), Tikhonov (c), Beltrami (d) and the 3D Patch reconstruction (e).




Figure 8.9: Comparison of the isotropic 3D reconstruction of the Brain cerebellum ($0.5 \times 0.5 \times 0.5 \text{ mm}^3$) from LR acquisitions ($0.5 \times 0.5 \times 2 \text{ mm}^3$) with Averaging (b), Tikhonov (c), Beltrami (d) and the 3D Patch reconstruction (e).

acquisitions ($1.25 \times 1.25 \times 8 \text{ mm}^3$) in SA, HLA and VLA views. The reconstructed isotropic volume is visually better with a higher overall SNR and a satisfactory CNR. Several areas appear with hyper-enhanced signal on the mid-interventricular septum and in the infero-lateral wall on the native SA and HLA view, but the precise area and transmural extent are difficult to define. However, the isotropic 3D reconstruction has the ability to attenuate the partial volume effects and provides a fibrosis area that is consistent between the different views. The proposed isotropic reconstruction shows good robustness to noise as well as a good preserving of edges and fine details in every orientation.

Reconstructions from a 78-year-old patient with a large myocardial infarction scar are also shown in Figure 8.11. The LR data were acquired with resolution $1.25 \times 1.25 \times 4.5 \text{ mm}^3$ and reconstructed to $1.25 \times 1.25 \times 1.25 \text{ mm}^3$ for this patient. First, we observe a good agreement between the native acquisitions and the reconstructed HR volume. The proposed reconstruction technique still yields to significantly better overall SNR, reducing the noise in flat areas (e.g. blood pool) while preserving thin

anatomical structures such as fat, papillary muscles and myocardial scar.

In terms of scoring by the two cardiologists, the patch-based reconstruction had a higher ranking score than the other two methods and was ranked as the best image in 15 patients (94%) by the first expert and in 16 patients (100%) by the second expert. Overall, we noted a statistically significant difference between the methods (Wilcoxon signed rank test, $p < 0.05$). The inter-observer reproducibility of the visual assessment, as expressed by the Cohens kappa coefficient, was excellent between the experts ($\kappa = 0.875$).

8.4.5 Discussion

In this study, we have presented an advanced algorithm for isotropic 3D reconstruction of MR images based on a joint denoising and SR framework. The proposed method integrates non-local self-similarities information, by grouping 3D patches with similar structures, into a geometrical SR framework. Low-rank properties and sparsity of the group are enforced to reduce the noise of the reconstructed volume while the SR was used to recover an isotropic 3D volume. An augmented Lagrangian formulation was used to efficiently decompose the main cost function into two sub-problems that have straightforward solutions. Experimental results on simulation, brain imaging and clinical cardiac imaging indicate that the proposed technique outperforms classic state-of-the-art SR methods, including local-based regularizations, and can achieve higher SNR and image sharpness at any given noise level. The improved performance can be explained by the fact that MR images contain a rich amount of non-local correlated structures and therefore high sparsity degree can be achieved by merging this information.

Preliminary results in clinical cardiac MRI shows that the reconstructed isotropic 3D volume provides improved information to the cardiologist, at no cost for assessing the location and transmural extent of myocardial fibrosis. DVOI In cardiac MR, the proposed technique might also be used to improve the robustness of challenging tasks, such as heart cavities segmentation or scar segmentation.

The proposed method iterates between self-similarity extraction, rank reduction




Figure 8.10: Example of 3D isotropic reconstruction results on a 20-year-old male Duchenne patient with presence of focal fibrosis. The proposed 3D isotropic reconstruction enables the reconstruction of an isotropic ($1.25 \times 1.25 \times 1.25 \text{ mm}^3$) volume with good SNR and sharp edges, compared to the native anisotropic volumes ($1.25 \times 1.25 \times 8 \text{ mm}^3$).




Figure 8.11: Comparison of native anisotropic low resolution acquisition with the isotropic 3D volume reconstructed with the proposed method on a 78-year-old male with myocardial infarction. Isotropic reconstruction allows for a better visualization of the transmurular left-ventricle scar in every direction.

and SR reconstruction. Although, we have shown in the simulation (see supplementary material I) that only a few iterations, typically less than 6, are needed to obtain the best performances, its computation burden remains high. However, we have been able to accelerate the reconstruction using overlapping and multithreading, reducing the computational time by a factor of 128 (i.e. an overlapping step of 4 was used with 32 processors), without significantly compromising image quality. Additionally, higher acceleration factors could be obtained using an implementation of the patch optimization on a graphic processing unit (GPU).

Our approach is well suited for real clinical cases by taking into account the acquisition information such as IPP/IOP in the optimization itself, via the operator E . We were able to reconstruct an isotropic volume without necessarily having to acquire perfectly orthogonal LR volumes. For this reason, it was hard to compare our technique with other promising SR techniques, such as 3D non-local means [73, 72] or joint low-rank and local regularization technique [103], which have been mainly used for single-frame MR SR. Further comparison with recent promising work, such as reconstructions based on residual convolutional neural network [85], will be studied in the future. Ultimately, this technique can also be extended to dynamic 4D cine-MRI sequences by exploiting spatio-temporal sparsity and therefore improving resolution and quality of 3D+t MR images.

Finally, the proposed isotropic 3D patch reconstruction can play a major role in many applications, where the acquisition of isotropic volumes is a challenging task.

Chapter 9

Summary and Future Work

9.1 Summary of Contributions

The main goal of the work presented in this manuscript was the development of new reconstruction techniques for motion artifact reduction, acquisition noise reduction and resolution improvement in cardiovascular magnetic resonance imaging. A novel denoising algorithm has been described and evaluated for grey-scale image enhancement and then further improved for myocardial T_1 mapping imaging, with the ability to provide a saturation-based sequence that is accurate and precise. A new free-breathing reconstruction has been proposed in single-shot imaging, further improving image visualization and patient comfort. Finally, a technique for 3D isotropic imaging has been proposed, enabling substantial improvement in image resolution and noise reduction with the potential to better assess subendocardial areas of infarction in patient with ischemic heart disease. The main contributions presented in this thesis are summarized below:

- Design and implementation of an efficient recovery algorithm for grey-scale images by employing the Beltrami energy as regularity prior. The implementation makes use of a primal-dual projected gradient algorithm for Beltrami minimization problems with application to image denoising. A convex simplification of the color-version of the Beltrami energy is presented, and thereby render the algorithm applicable to multichannel images enhancement. Numerical results are

presented on different images to show the performance of the primal-dual projected gradient algorithm using the Beltrami framework, and compared against state-of-the-art total variation (TV) models. Results show that the proposed Beltrami algorithm converges faster, and slightly improves the quality of results with regards to TV. Beltrami regularization maintains TV's feature preservation but further reduces the staircasing effect in restored images.

- Proposal of a novel denoising method, using Beltrami primal-dual regularization, which exploits the spatio-temporal correlations in the space of T_1 weighted images while preserving sharp anatomical details. Validation was performed by investigating the impact of the proposed denoising method on accuracy and precision of myocardial T_1 mapping on simulation, ex vivo and in vivo experiments. The algorithm is shown to improve precision while preserving accuracy in myocardial T_1 mapping. Ultimately this technique could enable accurate and precise reconstruction of myocardial T_1 maps, with the potential to offer better visual image analysis and improved performance of post-processing procedures such as registration and segmentation.
- Development and implementation of an efficient motion-compensated reconstruction technique for free-breathing cardiac magnetic resonance imaging that allows high-quality images to be reconstructed from multiple undersampled single-shot acquisitions. The technique uses a joint image reconstruction and motion correction method consisting of several steps, including a non-rigid motion extraction and a motion-compensated reconstruction. The main optimization problem is solved using the proposed efficient Beltrami regularization which offers an ideal compromise between feature preservation and staircasing reduction. In simulation, phantom and volunteer experiments, the proposed approach exhibits visible quality improvement over previous methods while reconstructing sharper edges, even for high factors of acceleration. Ultimately, the proposed technique could enable accurate motion corrected reconstruction of single-shot images with higher spatial resolution than conventional methods, with the potential to offer a high-quality late gadolinium enhancement examination even

in challenging patients.

- Proposal and implementation of a motion-corrected super-resolution technique for high-resolution, 3D isotropic cine imaging of the heart. A general acquisition strategy is presented and consists of the acquisition of multiple 2D image stacks with different orientations using a multi-slice b-SSFP sequence in free-breathing. The reconstruction technique consists of a motion-compensated 3D cine reconstruction of each stack and a super-resolution technique combining the multiple cine volume stacks into a single isotropic cine volume. Results in healthy volunteers and patients with Duchenne muscular dystrophy and breathing difficulties show the feasibility of the proposed technique in a clinical setting.
- Design and development of a novel technique that combines geometrical and statistical super-resolution techniques using a 3D self-similarity learning framework. The proposed approach merges information from multiple 3D anisotropic volumes, acquired in different orientations, into one single isotropic volume using a low-complexity 3D patch-based regularization. Preliminary results are presented in numerical simulations, brain imaging and clinical cardiac MRI. Experiments show that the proposed technique outperforms current state-of-the-art techniques in terms of image sharpness, noise reduction, and isotropic reconstruction. The proposed method may be particularly useful for cardiac applications such as fibrosis assessment by late gadolinium enhancement magnetic resonance imaging.

9.2 Future Work

This section concludes the work presented in this thesis with some current and future research directions as well as some possible extensions of each method covered in the previous chapters.

9.2.1 Primal-Dual Optimization with Beltrami Regularization

- The proposed denoising technique formulates an optimization problem involving a dual variable to the primal (image) variable with the Beltrami regularization, offering an interesting compromise between smooth regularization prior and the feature preservation of total variation regularization. Beyond the simple denoising problem, we extended the primal dual projected gradient to a color-based image denoising technique using a simplified version of the color-Beltrami model. Future work will focus on convex relaxations of the true multichannel-Beltrami model.
- Beltrami regularization was applied locally, meaning that each pixel was reconstructed using the information provided around the pixel of interest. An interesting complication arises, when instead of the very local image intensity, a semi-local image patch is embedded in the optimization itself. The denoising result can therefore be improved by restoring an unknown pixel using other similar pixels in the image. This non-local generalization is however not direct since the noise-free original image does not necessarily have the same similarity structures within its noisy version. This non-local patch-based Beltrami extension is currently under investigation.

9.2.2 Myocardial T_1 Mapping Denoising

- The denoising approach successfully corrected the encountered noise in all subjects. In this study, the cardiac noise pattern was mainly studied in post-contrast, after injection of gadolinium, where signal-to-noise ratio is often found higher. The performance of the method in native myocardial T_1 mapping was not investigated and should be addressed in future work. Moreover, since the signal does not use signal variation a priori, the proposed strategy can potentially benefit to any parametric mapping applications such as T_2 or T_2^* where precision is influenced by noise.

- The reproducibility and spatial variability of post-contrast T_1 mapping was improved using the Beltrami denoising. The use of the anisotropic denoising on T_1 -weighted images resulted in higher reproducibility and spatial variability of myocardial T_1 mapping. This work raises some interest in the field of diffusion tensor imaging (DTI), which has the potential of defining structural remodeling of diseased myocardium fiber. Indeed, diffusion tensors describe how much pixel intensities locally diffuse along given orthogonal orientations, and for this reason, diffusion tensor are highly related to anisotropic behavior. More advanced comparisons with such techniques deserve further investigation.
- In the first study, we show that Beltrami regularization outperforms the classic total variation regularization in terms of speed and image quality. It would be potentially of interest to compare with recent state-of-the-art denoising method used in computer vision, such as BM3D or non-local reconstruction, where an assumption of local sparsity through similar patches is adopted. Such techniques would however require high computation time and would not satisfy fast reconstruction, often required by cardiologists.
- In this study, we established the precision and accuracy of the proposed method with respect to spin echo imaging. Additional phantom evaluation can be done by considering, for example, multi-shot acquisitions as gold-standard.
- The post-acquisition denoising method shows to improve precision of saturation-recovery myocardium T_1 mapping sequence. Precision could also be improved using for instance free-breathing acquisition and motion correction. This would allow multi shots acquisitions with an increased matrix size and could enable the acquisition of more samples on the T_1 recovery curve and not being limited to eight saturation times.
- Further investigations in cross-center and cross-vendor study with additional comparisons with inversion-recovery type methods (e.g. MOLLI) are warranted and currently investigated.

9.2.3 Motion-Correction in Free-Breathing Single-Shot MRI

- The feasibility of the proposed reconstruction has been evaluated in simulation, phantom and volunteer experiments. The method has been shown to allow non-rigid motion correction while efficiently recovering features, thanks to the Beltrami regularization scheme. The proposed non-rigid registration technique was applied independently for each frame. Additional information can be provided by considering the whole set of frames. In this way, a multi-frame non-rigid registration combined with a rank constraint and a total variation norm, should improve motion estimation accuracy on regions suffering from large undersampling artifacts. More advanced registration methods and non-cartesian sampling deserve further investigation as well.
- Meanwhile, it would be of great interest to adapt the local Beltrami algorithm to a patch-based non-local regularization which should perform better to recover local image information. Each region of the heart is thus reconstructed and weighted using all similar regions in the image. Besides, dynamic weights can be optimized for our Beltrami regularity by solving a specific minimization problem. This non-locally regularized reconstruction method associated with parallel imaging and motion correction will be investigated in future work.
- One interesting application of the proposed motion correction technique is for high-resolution isotropic 3D late gadolinium enhancement imaging of the heart, such as the one proposed recently in [43] for myocardial scar assessment. This will allow the reconstruction of isotropic motion corrected volumes while keeping the advantages of a 2D acquisition, i.e. high tissue and vessel contrast and short acquisition time.
- A limitation to the method is that potential through-plane motion cannot be corrected, although it remains small compared to the slice thickness. To overcome this problem, one could consider weighting the single-shot images according to their position in the extracted respiratory signal. The preliminary results presented in this study should be confirmed with further patient studies.

9.2.4 Super-Resolution Cardiac Cine MRI

- The proposed joint motion-correction and super-resolution technique provided reconstruction improvements in terms of image sharpness and image quality of the reconstructed cine volumes. The Beltrami constraint provided efficient denoising without altering the effective resolution. Alternative methods for super-resolution cine imaging include dictionary learning approaches which do not explicitly require motion correction.
- Despite its high efficiency, the b-SSFP sequence used in this study is sensitive to B_0 field inhomogeneity that may result in dark banding artifacts. Although the requirement in terms of B_0 field homogeneity is less important for the 2D b-SSFP than it is for its 3D version, it does rely on efficient shimming in the volume of interest. It should also be noted that both the in-flow effect and the location of the dark-band artifacts depend on the slice orientation and may result in outliers in the native images. This might be overcome with robust super-resolution algorithms.
- Another limitation is that the motion correction and the super-resolution steps were applied sequentially. The final reconstruction might still be improved by merging the two steps into a single optimization problem. Such an approach would consist of searching for the isotropic super-resolution image directly from the motion-corrupted k-space data.

9.2.5 3D Patch-Based Isotropic MR Reconstruction

- The proposed isotropic 3D reconstruction of MR images based on a joint denoising and super-resolution framework achieved improved performances in simulation, brain imaging and clinical cardiac imaging experiments. The technique iterates between self-similarity extraction, rank reduction and super-resolution reconstruction. Although, we have shown in the simulation experiments that only a few iterations, typically less than six, are needed to obtain the best

performances, its computation burden remains high. Computational time reduction can be obtained using an implementation of the patch optimization on a graphic processing unit (GPU).

- The technique reconstructs a single isotropic volume without necessarily having to acquire perfectly orthogonal low-resolution volumes. For this reason, it was hard to compare our technique with other promising super-resolution techniques, such as three-dimension non-local means [73, 72], or joint low-rank and local regularization technique [103], which have been mainly used for single-frame MR super-resolution. Further comparison with recent promising work, such as reconstructions based on residual convolutional neural network will be studied in the future.
- The proposed technique can also be extended to dynamic 4D cine-MRI sequences by exploiting spatio-temporal sparsity and therefore improving resolution and quality of 3D+t MR images.

List of Publications: Journal Articles and Proceedings

2017

Bustin A, Odille F. Isotropic Reconstruction of MR Images using 3D Patch-Based Self-Similarity Learning. United States Patent. 2017. (In Process)

Bustin A, Voilliot D, Menini A, Felblinger J, Burschka D, Bonnemains L, Odille F. Isotropic Reconstruction of MR Images using 3D Patch-Based Self-Similarity Learning. IEEE Transactions on Medical Imaging. 2017. (Submitted)

Bustin A, Ferry P, Codreanu A, Beaumont M, Liu Shufang, Burschka D, Felblinger J, Brau A, Menini A, Odille F. Impact of a Joint Anisotropic Denoising Technique on Precision and Accuracy of Myocardial T1 Mapping. Journal of Magnetic Resonance Imaging. 2017. (Submitted)

Liu S, **Bustin A**, Ferry P, Codreanu A, Burschka D, Menini A, Odille F. A Vectorized Levenberg-Marquardt Model Fitting Algorithm for Efficient Post-Processing of Cardiac T1 Mapping MRI. Journal of Medical Image Analysis. 2017. (Submitted)

Liu S, **Bustin A**, Burschka D, Menini A, Odille F. GPU Implementation of Levenberg-Marquardt Optimization for T1 Mapping in MRI. IEEE International Conference on Robotics and Automation. 2017. (Submitted)

Bustin A, Menini A, Liu S, Burschka D, Codreanu A, Al-Mallah M, Janich M, Brau A, Odille F. Saturation-based Myocardial T1 Mapping with Denoising: Initial Comparative Study with MOLLI. *Journal of Cardiovascular Magnetic Resonance*, x(Suppl x):Px, 2017.

Ferry P, Codreanu A, Liu S, **Bustin A**, Brau A, Beaumont M, Martin J, Glenn S. T1 Mapping in healthy subjects using SMART1Map at 3T: a Comparison with MOLLI. *Journal of Cardiovascular Magnetic Resonance*, x(Suppl x):Px, 2017.

2016

Bustin A, Menini A, Martin J, Burschka D, Felblinger J, Brau A, Odille F. Motion Estimated-Compensated Reconstruction with Preserved-Features in Free-Breathing Cardiac MRI. *Medical Image Computing and Computer Assisted Intervention*, 2016.

Bustin A, Ferry P, Codreanu A, Menini A, Odille F. Exploiting Structural Information Along the T1-saturation Recovery Curve for Efficient Cardiac T1 Mapping Denoising. In: *Proceedings of the 24th Annual Meeting of ISMRM*. Singapore; 2016.

Fisher A, Menini A, **Bustin A**, Johnson K, Francois C, Brau A. Fully Self-gated Motion Compensated Cine Reconstruction from Free-breathing 2D Radial Cardiac MRI Data. In: *Proceedings of the 24th Annual Meeting of ISMRM*, Singapore; 2016.

Bustin A, Odille F, Kudielka G, Janich M, Brau A, Menini A. A Fully Automated Binning Method for Improved SHARP Reconstruction of Free-Breathing Cardiac Images. *Journal of Cardiovascular Magnetic Resonance*, 18(Suppl 1):P267, 2016.

Fisher A, Menini A, **Bustin A**, Johnson K, Francois C, Brau A. Motion Compensated Reconstruction from Free-Breathing 2D Radial Cardiac MRI Data. *Journal of Cardiovascular Magnetic Resonance*, 18(Suppl 1):O110, 2016.

2015

Odille F, **Bustin A**, Bailiang C, Vuissoz PA, Felblinger J. Motion-Corrected. Super-Resolution Reconstruction for High-Resolution 3D Cardiac Cine MRI. *Medical Image Computing and Computer Assisted Intervention*, 2015.

Bustin A, Janich M, Brau A, Odille F, Wolff D, Shubayev O, Stanley D, Menini A. Joint Denoising and Motion Correction: Initial Application in Single-Shot Cardiac MRI. *Journal of Cardiovascular Magnetic Resonance*, 17(Suppl 1):Q29, 2015.

Rincon T, **Bustin A**, Sanchez A, Janich M, Burschka D, Menini A, Brau A. Optimal k-space Sampling in Multiple Accelerated Single Shot Acquisition for Motion Correction in MRI. In: *Proceedings of the Annual Meeting of ESMRMB*. Edinburgh, UK, 2015.

Bustin A, Menini A, Rincon T, Liu S, Burschka D, Janich M, Wolff S, Shubayev O, Stanley D, Odille F, Brau A. Motion Compensated Reconstruction in Accelerated Single-Shot Cardiac MRI. In: *Proceedings of the 23th Annual Meeting of ISMRM*, Toronto, Canada; 2015.

Bustin A, Menini A, Liu S, Burschka D, Janich M, Wolff S, Shubayev O, Stanley D, Odille F, Brau A. Reconstruction avec Compensation de Mouvement en IRM Cardiaque Single-Shot. In: *Proceedings of the Annual Meeting of SFRMBM*, Grenoble, France; 2015.

2014

Zosso D, **Bustin A**. A Primal-Dual Projected Gradient Algorithm for Efficient Beltrami Regularization. *UCLA CamReport*, 2014.

Zosso D, **Bustin A**. A Primal-Dual Projected Gradient Algorithm for Efficient Beltrami Regularization. In: Proceedings of the Gordon Research Conference on Image Science, Stonehill College, Easton, MA, USA: 2014.

Bibliography

- [1] M. Aharon, M. Elad, and A. Bruckstein. k-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006.
- [2] M. Akçakaya, H. Rayatzadeh, T. Basha, S. N. Hong, R. H. Chan, K. V. Kissinger, T. H. Hauser, M. E. Josephson, and W. J. Manning. Accelerated late gadolinium enhancement cardiac mr imaging with isotropic spatial resolution using compressed sensing: Initial experience. *Radiology*, 264(3):691–699, 2012.
- [3] A. Aldroubi and M. Unser. Wavelets in medicine and biology, 1996.
- [4] L. Alvarez, P.-L. Lions, and J.-M. Morel. Image Selective Smoothing and Edge Detection by Nonlinear Diffusion. II. *SIAM Journal on Numerical Analysis*, 29(3):845–866, 1992.
- [5] Y. Amano, M. Takayama, and S. Kumita. Contrast-enhanced myocardial T1-weighted scout (look-locker) imaging for the detection of myocardial damages in hypertrophic cardiomyopathy. *Journal of Magnetic Resonance Imaging*, 30(4):778–784, 2009.
- [6] K. Arrow, L. Hurwicz, and H. Uzawa. Studies in nonlinear programming. *Stanford University Press, Stanford, CA*, 1(9):5, 1958.
- [7] H. Ashikaga, H. L. Estner, D. A. Herzka, E. R. McVeigh, and H. R. Halperin. Quantitative assessment of single-image super-resolution in myocardial scar

- imaging. *IEEE Journal of Translational Engineering in Health and Medicine*, 2(1), 2014.
- [8] D. Atkinson and D. L. G. Hill. Reconstruction after rotational motion. *Magnetic Resonance in Medicine*, 49(1):183–187, 2003.
- [9] F. Balezeau, P.-A. Eliat, A. B. Cayamo, and H. Saint-Jalmes. Mapping of low flip angles in magnetic resonance. *Physics in medicine and biology*, 56(20):6635–47, 2011.
- [10] J. K. Barral, E. Gudmundson, N. Stikov, M. Etezadi-Amoli, P. Stoica, and D. G. Nishimura. A robust methodology for in vivo T1 mapping. *Magnetic Resonance in Medicine*, 64(4):1057–1067, 2010.
- [11] P. G. Batchelor, D. Atkinson, P. Irarrazaval, D. L. G. Hill, J. Hajnal, and D. Larkman. Matrix description of general motion correction applied to multi-shot images. *Magnetic Resonance in Medicine*, 54(5):1273–1280, 2005.
- [12] R. Ben-Ari and N. Sochen. A geometric framework and a new criterion in optical flow modeling. *J. Math. Imaging Vis.*, 33(2):178–194, nov 2009.
- [13] M. a. Bernstein, K. E. King, X. J. Zhou, and W. Fong. *Handbook of MRI Pulse Sequences*. Elsevier, 2004.
- [14] M. Bertero, T. A. Poggio, and V. Torre. Ill-Posed Problems in Early Vision. *Proceedings of the IEEE*, 76(8):869–889, 1988.
- [15] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and distributed computation*. Prentice Hall Inc., Old Tappan, NJ (USA), 1989.
- [16] K. K. Bhatia, A. N. Price, W. Shi, J. V. Hajnal, and D. Rueckert. Super-resolution reconstruction of cardiac MRI using coupled dictionary learning. *2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, pages 947–950, 2014.

- [17] J. Bigun, G. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE transactions on pattern analysis and machine intelligence*, 13(8):775 – 790, 1991.
- [18] G. Blanchet and L. Moisan. An explicit sharpness index related to global phase coherence. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pages 1065–1068, 2012.
- [19] I. Bogdanova, X. Bresson, J.-P. Thiran, and P. Vandergheynst. Scale space analysis and active contours for omnidirectional images. *IEEE Transactions on Image Processing*, 16(7):1888–1901, 2007.
- [20] P. A. Bottomley, C. J. Hardy, R. E. Argersinger, and G. Allen-Moore. A review of 1H nuclear magnetic resonance relaxation in pathology: are T1 and T2 diagnostic? *Med Phys*, 14(1):1–37, 1987.
- [21] X. Bresson and T. T. F. Chan. Fast dual minimization of the vectorial total variation norm and applications to color image processing. *Inverse Problems and Imaging*, Volume 2(Issue 4):455–484, 2008.
- [22] X. Bresson, S. Esedolu, S. Esedoglu, P. Vandergheynst, J.-P. Thiran, S. Osher, S. Esedoglu, P. Vandergheynst, J.-P. Thiran, and S. Osher. Fast Global Minimization of the Active Contour/Snake Model. *Journal of Mathematical Imaging and Vision*, 28(2):151–167, jul 2007.
- [23] X. Bresson, P. Vandergheynst, and J.-P. Thiran. Multiscale Active Contours. *Int. J. Comput. Vis.*, 70(3):197–211, 2006.
- [24] A. Buades, B. Coll, and J.-M. Morel. *A Non-Local Algorithm for Image Denoising*. IEEE, 2005.
- [25] J. Caballero, A. N. Price, D. Rueckert, and J. V. Hajnal. Dictionary learning and time sparsity for dynamic MR data reconstruction. *IEEE Transactions on Medical Imaging*, 33(4):979–994, 2014.

- [26] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
- [27] A. Chambolle. An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision*, 20:89–97, 2004.
- [28] A. Chambolle and T. Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [29] T. F. Chan, G. H. Golub, and P. Mulet. A Nonlinear Primal-Dual Method for Total Variation-Based Image Restoration. *SIAM Journal on Scientific Computing*, 20(6):1964–1977, 1999.
- [30] K. Chow, J. a. Flewitt, J. D. Green, J. J. Pagano, M. G. Friedrich, and R. B. Thompson. Saturation recovery single-shot acquisition (SASHA) for myocardial T 1 mapping. *Magnetic Resonance in Medicine*, 71(6):2082–2095, 2014.
- [31] P. Coupé. Collaborative patch-based super-resolution for diffusion-weighted images. *NeuroImage*, 83:245–261, 2013.
- [32] P. Coupé, P. Yger, and C. Barillot. Fast non local means denoising for 3D MR images. *Medical image computing and computer-assisted intervention : MIC-CAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, 9(Pt 2):33–40, 2006.
- [33] G. Cruz, D. Atkinson, C. Kolbitsch, T. Schaeffter, and C. Prieto. Comparison of Non-Rigid Motion Compensated Reconstructions for 3D Abdominal MRI. *ISMRM*, 22:886, 2014.
- [34] K. Dabov, A. Foi, and V. Katkovnik. Image denoising by sparse 3D transformation-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):1–16, 2007.

- [35] R. Damadian. Tumor detection by nuclear magnetic resonance. *Science*, 171(3976):1151–1153, 1971.
- [36] L. Dascal, G. Rosman, X.-C. Tai, and R. Kimmel. On Semi-implicit Splitting Schemes for the Beltrami Color Flow. In *Scale Space and Variational Methods in Computer Vision*, volume 5567 of *Lecture Notes in Computer Science*, pages 259–270. Springer Berlin / Heidelberg, 2009.
- [37] L. Dascal and N. A. Sochen. A Maximum Principle for Beltrami Color Flow. *SIAM J. Appl. Math.*, 65(5):1615–1632, 2005.
- [38] J. A. Derbyshire, H. Saybasili, L. Guo, O. Sayin, P. Kellman, R. J. Lederman, and D. A. Herzka. Golden-step phase encoding for flexible realtime Cardiac MRI. *Journal of Cardiovascular Magnetic Resonance*, 13(Suppl 1):P23, 2011.
- [39] M. Doneva, P. Börnert, H. Eggers, C. Stehning, J. S enegas, and A. Mertins. Compressed sensing reconstruction for magnetic resonance parameter mapping. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 64(4):1114–20, oct 2010.
- [40] D. L. Donoho. De-Noising by Soft-Thresholding. *IEEE Transactions on Information Theory*, 41(3):613–627, 1995.
- [41] D. L. Donoho, I. M. Johnstone, G. Kerkyacharian, and D. Picard. Wavelet Shrinkage: Asymptopia? *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(2):301–369, 1995.
- [42] D. L. Donoho and J. M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, 1994.
- [43] O. Dzyubachyk, Q. Tao, D. H. Poot, H. J. Lamb, K. Zeppenfeld, B. P. Lelieveldt, and R. J. van der Geest. Super-resolution reconstruction of late gadolinium-enhanced MRI for improved myocardial scar assessment. *Journal of Magnetic Resonance Imaging*, 42:160–167, 2015.

- [44] I. Ekeland, R. Temam, R. Témam, and R. Témam. *Convex Analysis and Variational Problems*. Society for Industrial and Applied Mathematics; 1 edition (January 1, 1987), 1999.
- [45] L. Feng, L. Axel, H. Chandarana, K. T. Block, D. K. Sodickson, and R. Otazo. XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing. *Magnetic Resonance in Medicine*, 75(2):775–788, 2016.
- [46] A. Fischer, S. Weick, C. O. Ritter, M. Beer, C. Wirth, H. Hebestreit, P. M. Jakob, D. Hahn, T. Bley, and H. Köstler. SELF-gated Non-Contrast-Enhanced FUNCTIONAL Lung imaging (SENCEFUL) using a quasi-random fast low-angle shot (FLASH) sequence and proton MRI. *NMR in biomedicine*, 27(8):907–17, 2014.
- [47] T. K. F. Foo, G. S. Slavin, D. a. Bluemke, M. Montequin, M. N. Hood, and V. B. Ho. Simultaneous myocardial and fat suppression in magnetic resonance myocardial delayed enhancement imaging. *Journal of magnetic resonance imaging : JMRI*, 26(4):927–33, 2007.
- [48] A. Gholipour, J. A. Estroff, and S. K. Warfield. Robust Super-Resolution Volume Reconstruction From Slice Acquisitions: Application to Fetal Brain MRI. *IEEE Transactions on Medical Imaging*, 29(10):1739–1758, 2010.
- [49] T. Goldstein, E. Esser, and R. G. Baraniuk. Adaptive Primal-Dual Hybrid Gradient Methods for Saddle-Point Problems. *arXiv preprint arXiv:1305.0546*, 2013.
- [50] H. Greenspan, S. Peled, G. Oz, and N. Kiryati. MRI inter-slice reconstruction using super-resolution. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2208:1204–1206, 2001.
- [51] H. Gudbjartsson and S. Patz. The rician distribution of noisy mri data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995.

- [52] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*, volume 82. Wiley-Liss New York, 1999.
- [53] W. W. W. Hager and H. Zhang. A Survey of Nonlinear Conjugate Gradient Methods. *Pacific journal of Optimization*, 2(1):35–58, 2006.
- [54] J. P. Haldar, V. J. Wedeen, M. Nezamzadeh, G. Dai, M. W. Weiner, N. Schuff, and Z. P. Liang. Improved diffusion imaging through SNR-enhancing joint reconstruction. *Magnetic Resonance in Medicine*, 69(1):277–289, 2013.
- [55] P. C. Hansen. Analysis of Discrete Ill-Posed Problems by Means of the L-Curve. *SIAM Review*, 34(4):561–580, 1992.
- [56] M. Hestenes and E. Stiefel. *Methods of conjugate gradients for solving linear systems*, 1952.
- [57] C. Huang, C. G. Graff, E. W. Clarkson, A. Bilgin, and M. I. Altbach. T2 mapping from highly undersampled data by reconstruction of principal component coefficient maps using compressed sensing. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 67(5):1355–66, 2012.
- [58] M. Karlsson and B. Nordell. Analysis of the Look-Locker T1 mapping sequence in dynamic contrast uptake studies: Simulation and in vivo validation. *Magnetic Resonance Imaging*, 18(8):947–954, 2000.
- [59] P. Kellman and M. S. Hansen. T1-mapping in the heart: accuracy and precision. *Journal of Cardiovascular Magnetic Resonance*, 16(1):2, 2014.
- [60] P. Kellman, A. C. Larson, L. Y. Hsu, Y. C. Chung, O. P. Simonetti, E. R. McVeigh, and A. E. Arai. Motion-corrected free-breathing delayed enhancement imaging of myocardial infarction. *Magnetic Resonance in Medicine*, 53(1):194–200, 2005.

- [61] P. Kellman, H. Xue, K. Chow, B. S. Spottiswoode, A. E. Arai, and R. B. Thompson. Optimized saturation recovery protocols for T1-mapping in the heart: influence of sampling strategies on precision. *Journal of Cardiovascular Magnetic Resonance*, 16(1):55, 2014.
- [62] R. Kimmel, R. Malladi, and N. Sochen. Images as Embedded Maps and Minimal Surfaces: Movies, Color, Texture, and Volumetric Medical Images. *Int. J. Comput. Vis.*, 39(2):111–129, 2000.
- [63] R. Kimmel, N. Sochen, and R. Malladi. From high energy physics to low level vision. *Scale-Space Theory in Computer Vision*, 1252:236–247, 1997.
- [64] F. Knoll, K. Bredies, T. Pock, and R. Stollberger. Second order total generalized variation (TGV) for MRI. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 65(2):480–91, feb 2011.
- [65] J. Koikkalainen, M. Pollari, J. Lötjönen, S. Kivistö, and K. Lauerma. Segmentation of cardiac structures simultaneously from short-and long-axis MR images. *Medical Image Computing and Computer-Assisted InterventionMICCAI 2004*, pages 427–434, 2004.
- [66] N. Komodakis and J.-C. Pesquet. Playing with Duality. *IEEE Signal Processing Magazine*, 32(6):31–54, 2015.
- [67] M. Kuklisova-Murgasova, G. Quaghebeur, M. A. Rutherford, J. V. Hajnal, and J. A. Schnabel. Reconstruction of fetal brain MRI with intensity matching and complete outlier removal. *Medical Image Analysis*, 16(8):1550–1564, 2012.
- [68] F. Lam, S. D. Babacan, J. P. Haldar, M. W. Weiner, N. Schuff, and Z. P. Liang. Denoising diffusion-weighted magnitude MR images using rank and edge constraints. *Magnetic Resonance in Medicine*, 71(3):1272–1284, 2014.
- [69] P. C. Lauterbur. Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance, 1973.

- [70] D. Liang, B. Liu, and L. Ying. Accelerating sensitivity encoding using compressed sensing. In *IEEE Engineering in Medicine and Biology Society*, pages 1667–1670, 2008.
- [71] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic resonance in medicine*, 58(6):1182–1195, 2007.
- [72] J. V. Manjón, P. Coupé, A. Buades, D. L. Collins, and M. Robles. MRI superresolution using self-similarity and image priors. *International Journal of Biomedical Imaging*, 2010, 2010.
- [73] J. V. Manjón, P. Coupé, A. Buades, V. Fonov, D. Louis Collins, and M. Robles. Non-local MRI upsampling. *Medical Image Analysis*, 14(6):784–792, 2010.
- [74] J. V. Manjón, P. Coupé, L. Martí-Bonmatí, D. L. Collins, and M. Robles. Adaptive non-local means denoising of MR images with spatially varying noise levels. *Journal of Magnetic Resonance Imaging*, 31(1):192–203, 2010.
- [75] P. Mansfield and A. A. Maudsley. Planar spin imaging by NMR. *Journal of Magnetic Resonance (1969)*, 27(1):101–119, 1977.
- [76] D. W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters, 1963.
- [77] A. Menini, G. S. Slavin, J. a. Stainsby, P. Ferry, J. Felblinger, and F. Odille. Motion correction of multi-contrast images applied to T1 and T 2 quantification in cardiac MRI. *Magma (New York, N.Y.)*, pages 1–18, mar 2014.
- [78] D. R. Messroghli, S. Plein, D. M. Higgins, K. Walters, T. R. Jones, J. P. Ridgway, and M. U. Sivananthan. Human myocardium: single-breath-hold MR T1 mapping with high spatial resolution–reproducibility study. *Radiology*, 238(3):1004–1012, 2006.

- [79] D. R. Messroghli, A. Radjenovic, S. Kozerke, D. M. Higgins, M. U. Sivananthan, and J. P. Ridgway. Modified Look-Locker inversion recovery (MOLLI) for high-resolution T1 mapping of the heart. *Magnetic Resonance in Medicine*, 52(1):141–146, 2004.
- [80] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer, 2nd edition, 1999.
- [81] F. Odille, A. Bustin, B. Chen, P.-A. Vuissoz, and J. Felblinger. Motion-Corrected, Super-Resolution Reconstruction for High-Resolution 3D Cardiac Cine MRI. *Medical image computing and computer-assisted intervention*, 9351:435–442, 2015.
- [82] F. Odille, J. M. Escanyé, D. Atkinson, L. Bonnemains, and J. Felblinger. Non-rigid registration improves MRI T2 quantification in heart transplant patient follow-up. *Journal of Magnetic Resonance Imaging*, 42(1):168–174, 2015.
- [83] F. Odille, A. Menini, J.-M. Escanye, P.-A. Vuissoz, P.-y. Marie, M. Beaumont, and J. Felblinger. Joint Reconstruction of Multiple Images and Motion in MRI: Application to Free-Breathing Myocardial T2 Quantification. *IEEE Transactions on Medical Imaging*, 35(1):197–207, 2016.
- [84] F. Odille, P. A. Vuissoz, J. Felblinger, and D. Atkinson. Generalized Reconstruction by Inversion of Coupled Systems (GRICS) applied to parallel MRI. *Magnetic Resonance in Medicine*, 60(1):146–157, 2008.
- [85] O. Oktay, W. Bai, M. Lee, R. Guerrero, K. Kamnitsas, J. Caballero, A. de Marvao, S. Cook, D. O’Regan, and D. Rueckert. Multi-input cardiac image super-resolution using convolutional neural networks. *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2016.
- [86] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin. An Iterative Regularization Method for Total Variation-Based Image Restoration. *Multiscale Modeling & Simulation*, 4(2):460–489, 2005.

- [87] W. Pauli. Zur Frage der theoretischen Deutung der Satelliten einiger Spektrallinien und ihrer Beeinflussung durch magnetische Felder. *Naturwissenschaften*, 12(37):741–743, 1924.
- [88] R. R. Peeters, P. Kornprobst, M. Nikolova, S. Sunaert, T. Vieville, G. Mandain, R. Deriche, O. Faugeras, M. Ng, and P. Van Hecke. The use of super-resolution techniques to reduce slice thickness in functional MRI. *International Journal of Imaging Systems and Technology*, 14(3):131–138, 2004.
- [89] F. D. Pineda, M. Medved, X. Fan, and G. S. Karczmar. B1 and T1 mapping of the breast with a reference tissue method. *Magnetic Resonance in Medicine*, 75(4):1565–1573, 2016.
- [90] E. Plenge, D. H. J. Poot, M. Bernsen, G. Kotek, G. Houston, P. Wielopolski, L. Van Der Weerd, W. J. Niessen, and E. Meijering. Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magnetic Resonance in Medicine*, 68(6):1983–1993, 2012.
- [91] A. M. Polyakov. Quantum geometry of bosonic strings. *Physics Letters B*, 103(3):207–210, 1981.
- [92] W. G. Proctor and F. C. Yu. The dependence of a nuclear magnetic resonance frequency upon chemical compound. *Physical Review*, 77(5):717, 1950.
- [93] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 42(5):952–62, nov 1999.
- [94] I. I. Rabi, S. Millman, P. Kusch, and J. R. Zacharias. The molecular beam resonance method for measuring nuclear magnetic moments. *Physical Review*, 55(6):526–535, 1939.
- [95] G. Rosman, X.-C. Tai, L. Dascal, and R. Kimmel. Augmented Lagrangian for Polyakov Action Minimization in Color Images. In *Aip Conference Proceedings*, volume 1281, pages 1018–1021, 2010.

- [96] G. Rosman, X.-C. Tai, L. Dascal, and R. Kimmel. Polyakov Action for Efficient Color Image Processing. Technical Report 623, Technion CIS-2010-04, 2010.
- [97] G. Rosman, X. C. Tai, L. Dascal, and R. Kimmel. *Polyakov action minimization for efficient color image processing*. Springer, 2010.
- [98] S. Roujol, M. Foppa, S. Weingärtner, W. J. Manning, and R. Nezafat. Adaptive registration of varying contrast-weighted images for improved tissue characterization (ARCTIC): Application to T1 mapping. *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, 00:1–14, 2014.
- [99] F. Rousseau, O. A. Glenn, B. Iordanova, C. Rodriguez-Carranza, D. B. Vigneron, J. A. Barkovich, and C. Studholme. Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images. *Academic Radiology*, 13(9):1072–1081, 2006.
- [100] L. Rudin, S. Osher, and E. Fatemi. Nonlinear Total Variation based Noise Removal Algorithms. *Physica D*, 60(1-4):259–268, 1992.
- [101] K. Scheffler and S. Lehnhardt. Principles and applications of balanced SSFP techniques. *European Radiology*, 13(11):2409–2418, 2003.
- [102] A. D. Scott, J. Keegan, and D. N. Firmin. Motion in cardiovascular MR imaging. *Radiology*, 250(2):331–351, 2009.
- [103] F. Shi, J. Cheng, L. Wang, P.-T. Yap, and D. Shen. LRTV: MR Image Super-Resolution with Low-Rank and Total Variation Regularizations. *IEEE Transactions on Medical Imaging*, 0062(c):1–1, 2015.
- [104] W. Shi, J. Caballero, C. Ledig, X. Zhuang, W. Bai, K. Bhatia, A. M. S. M. de Marvao, T. Dawes, D. O’Regan, and D. Rueckert. Cardiac Image Super-Resolution with Global Correspondence Using Multi-Atlas PatchMatch. *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, 8151:9–16, 2013.

- [105] G. S. Slavin. Breath-Held Myocardial T1 Mapping Using Multiple Single-Point Saturation Recovery. *Proceedings of the ISMRM*, 20(6):1244, 2012.
- [106] G. S. Slavin and J. A. Stainsby. True T1 mapping with SMART1Map (saturation method using adaptive recovery times for cardiac T 1 mapping): a comparison with MOLLI. *Journal of Cardiovascular Magnetic Resonance*, 15(Suppl 1):P3, 2013.
- [107] N. Sochen, R. Kimmel, and R. Malladi. A general framework for low level vision. *IEEE Transactions on Image Processing*, 7(3):310–318, 1998.
- [108] J. A. Stainsby and G. S. Slavin. Myocardial T1 mapping using SMART1Map : initial in vivo experience. *Journal of Cardiovascular Magnetic Resonance*, 15(Suppl 1):P13, 2013.
- [109] J. C. Strikwerda. *Finite difference schemes and partial differential equations*. SIAM, 2004.
- [110] D. M. Strong, P. Blomgren, and T. F. Chan. Spatially adaptive local-feature-driven total variation minimizing image restoration. In *Proceedings of SPIE 3167*, volume 3167, pages 222–233. SPIE, 1997.
- [111] X. C. Tai, S. Borok, and J. Hahn. Image denoising using TV-Stokes equation with an orientation-matching minimization. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 490–501. Springer, 2009.
- [112] A. N. Tichonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics*, 4:1035–1038, 1963.
- [113] S. Tourbier, X. Bresson, P. Haggmann, J. P. Thiran, R. Meuli, and M. B. Cuadra. Efficient total variation algorithm for fetal brain MRI reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 252–259. Springer, 2014.

- [114] A. Tristán-Vega, V. García-Pérez, S. Aja-Fernández, and C. F. Westin. Efficient and robust nonlocal means denoising of MR data based on salient features matching. *Computer Methods and Programs in Biomedicine*, 105(2):131–144, 2012.
- [115] 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*, 70(2):504–516, 2013.
- [116] G. Van Steenkiste, B. Jeurissen, J. Veraart, A. J. den Dekker, P. M. Parizel, D. H. J. Poot, and J. Sijbers. Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations. *Magnetic Resonance in Medicine*, 195(January 2015):181–195, 2015.
- [117] L. Vese. A Study in the BV Space of a Denoising—Deblurring Variational Problem. *Applied Mathematics and Optimization*, 44(2):131–161, 2001.
- [118] P.-A. Vuissoz, F. Odille, B. Fernandez, M. Lohezic, A. Benhadid, D. Mandry, and J. Felblinger. Free-breathing imaging of the heart using 2D cine-GRICS (generalized reconstruction by inversion of coupled systems) with assessment of ventricular volumes and function. *Journal of magnetic resonance imaging : JMRI*, 35(2):340–51, feb 2012.
- [119] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, 13(4):600–12, apr 2004.
- [120] J. B. M. Warntjes, O. Dahlqvist Leinhard, J. West, and P. Lundberg. Rapid magnetic resonance quantification on the brain: Optimization for clinical usage. *Magnetic Resonance in Medicine*, 60(2):320–329, 2008.

- [121] J. B. Weaver, Y. Xu, D. M. Healy, and L. D. Cromwell. Filtering noise from images with wavelet transforms. *Magnetic Resonance in Medicine*, 21(2):288–295, 1991.
- [122] J. Weickert, B. M. ter Haar Romeny, and M. A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing*, 7(3):398–410, 1998.
- [123] M. Westwood, L. J. Anderson, D. N. Firmin, P. D. Gatehouse, C. C. Charrier, B. Wonke, and D. J. Pennell. A single breath-hold multiecho T2* cardiovascular magnetic resonance technique for diagnosis of myocardial iron overload. *Journal of Magnetic Resonance Imaging*, 18(1):33–39, 2003.
- [124] S. Winkelmann, T. Schaeffter, T. Koehler, H. Eggers, and O. Doessel. An optimal radial profile order based on the golden ratio for time-resolved MRI. *IEEE Transactions on Medical Imaging*, 26(1):68–76, 2007.
- [125] J. Yang, J. Wright, T. S. Huang, Y. Ma, J. Yang, Y. Ma, S. Member, J. Wright, and T. S. Huang. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010.
- [126] T. Zhang, J. M. Pauly, and I. R. Levesque. Accelerating parameter mapping with a locally low rank constraint. *Magnetic Resonance in Medicine*, 73(2):655–661, 2015.
- [127] B. Zhao, W. Lu, T. K. Hitchens, F. Lam, C. Ho, and Z. P. Liang. Accelerated MR parameter mapping with low-rank and sparsity constraints. *Magnetic Resonance in Medicine*, 74(2):489–498, 2015.
- [128] B. Zhao, W. Lu, and Z.-P. Liang. Highly Accelerated Parameter Mapping with Joint Partial Separability and Sparsity Constraints. *Proceedings of the ISMRM*, 20:2233, 2012.
- [129] M. Zhu and T. Chan. An efficient primal-dual hybrid gradient algorithm for total variation image restoration. *UCLA CAM Report*, pages 8–34, 2008.

- [130] M. Zhu, S. J. Wright, and T. F. Chan. Duality-based algorithms for total-variation-regularized image restoration. *Computational Optimization and Applications*, 47(3):377–400, 2010.
- [131] M. Zibulevsky and M. Elad. L1-L2 optimization in signal and image processing. *IEEE Signal Processing Magazine*, 27(3):76–88, 2010.
- [132] D. Zosso and J.-P. Thiran. A scale-space of cortical feature maps. *IEEE Signal Processing Letters*, 16(10):873–876, 2009.