Technological breakthroughs magnetic resonance imaging towards motion-robust super-resolution quantitative for improved detection of brain diseases

Quinten Beirinckx

Ռ

Technological breakthroughs towards motion-robust super-resolution quantitative magnetic resonance imaging for improved detection of brain diseases

Quinten Beirinckx



Supervisors Prof. Dr. Jan Sijbers | Dr. Ir. Arnold J. den Dekker | Prof. Dr. Marleen Verhoye

Thesis submitted for the degree of Doctor of Science: Physics Faculty of Science | Department of Physics | Antwerp, 2024





Faculteit Wetenschappen Departement Fysica

Technological breakthroughs towards motion-robust super-resolution quantitative magnetic resonance imaging for improved detection of brain diseases

Technologische doorbraken richting bewegingsrobuuste superresolutie kwantitatieve magnetische resonantiebeeldvorming voor verbeterde detectie van hersenziekten

Proefschrift voorgelegd tot het behalen van de graad van

Doctor in de Wetenschappen: Fysica

aan de Universiteit Antwerpen, te verdedigen door

Quinten BEIRINCKX

Promotoren: Prof. Dr. Jan Sijbers Dr. Ir. Arnold J. den Dekker Prof. Dr. Marleen Verhoye

Antwerpen, 2024

Doctoral Committee

Chairperson

Prof. Dr. Sabine Van Doorslaer - University of Antwerp, Antwerp, Belgium

Members

Prof. Dr. Jan Sijbers - University of Antwerp, Antwerp, Belgium Dr. Ir. Arnold J. den Dekker - University of Antwerp, Antwerp, Belgium Prof. Dr. Marleen Verhoye - University of Antwerp, Antwerp, Belgium Prof. Dr. Joris Dirckx - University of Antwerp, Antwerp, Belgium

External members

Prof. Dr. Ir. Aleksandra Pizurica - Ghent University, Ghent, Belgium Dr. Ir. Matthan Caan - Amsterdam University Medical Center, University of Amsterdam, Amsterdam, The Netherlands

Contact information

Quinten Beirinckx
 imec - Vision Lab, Department of Physics
 University of Antwerp (CDE)
 Universiteitsplein 1, Building N.1.16
 B-2610 Wilrijk - Antwerpen, Belgium

- **a** +32 (0)3 265 24 58
- 🖗 quinten.beirinckx@uantwerpen.be
- 🕸 qbeirinckx@gmail.com
- https://visielab.uantwerpen.be/people/quinten-beirinckx

Copyright © 2024 by Quinten Beirinckx

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, broadcasting or by any other information storage and retrieval system without written permission from the copyright owner.

Cover illustration

High-resolution quantitative brain T1 mapping from low-resolution T1-weighted MRI scans using model-based super-resolution reconstruction.

The research presented in this thesis was performed at the imec-Vision Lab (Dept. of Physics) of the University of Antwerp, in close collaboration with the Antwerp University Hospital (UZA, Dept. of Radiology) for acquisition of the in vivo MR relaxometry data, and the C.J. Gorter MRI Center (Dept. of Radiology) of the Leiden University Medical Center (Leiden, The Netherlands) for acquisition of the in vivo Arterial Spin Labeling MRI data.

Dankwoord

Het onderzoek in dit proefschrift zou nooit mogelijk zijn geweest zonder de hulp en steun van collega's, vrienden en familie. Daarom wil ik graag mijn oprechte dank uitspreken aan de volgende mensen:

Aan Jan Sijbers en Arjan den Dekker, mijn supervisors en promotoren bij imec-Vision Lab. Bedankt dat jullie mij de kans hebben gegeven om als onderzoeker onder jullie vleugels te starten. Bedankt voor jullie niet-aflatende steun, waardevolle begeleiding en eindeloze aanmoedigingen tijdens mijn doctoraatstraject, van start tot finish. Bedankt voor de vele uren die jullie hebben besteed aan het proeflezen van mijn werk, het verbeteren van mijn manuscripten en het aanreiken van nieuwe inzichten tijdens mijn uitgebreid PhD-traject. Jullie complementaire expertise, geduld, oog voor detail en constructieve feedback hebben een belangrijke rol gespeeld bij het vormgeven van de onderzoeksbijdragen in dit proefschrift en bij mijn groei als onderzoeker en ook als mens. Ik ben oprecht dankbaar dat ik de kans heb gekregen om onder jullie beider mentorschap te mogen samenwerken.

Aan Ben Jeurissen, wie ik graag als mede-mentor van mijn onderzoekswerk aanstip. Bedankt dat je steeds beschikbaar was met vernuftige inzichten en nuttige antwoorden op technische en wetenschappelijke vragen omtrent mijn implementaties en modelleringen. Dankzij jou ben ik heel wat nieuwe programmeervaardigheden en -tips rijker. Het was steeds erg aangenaam om met jou samen te werken en te breinstormen. Dank ook voor de ontspannende en spontane momenten naast het werk, op conferenties, al fietsend, of tijdens de gezellige etentjes onder collega's.

Aan Dirk Poot, die mij aan het begin van mijn doctoraat heeft geholpen met enkele specifieke onderzoeksvragen over superresolutie-reconstructie, die onze gezamenlijke publicaties altijd zorgvuldig heeft nagelezen, met vaak gevatte en meerwaarde biedende feedback. Ook oprecht bedankt voor de twee korte onderzoeksbezoeken aan het Erasmus Medisch Centrum in Rotterdam, resulterend in een vruchtbare wetenschappelijke samenwerking.

Aan Marleen Verhoye, voor mijn eerste leerrijke stappen in de wereld van MRI-scanners tijdens de lessen over biomedische beeldvormingstechnieken in mijn master fysica. Jouw passie voor het MRI-domein en als docent heeft mij toen diepgaand geïnspireerd. Bedankt ook voor de uitgebreide ondersteuning in de aanloop naar mijn IOF-project en tijdens de verschillende feedbackmomenten omtrent mijn onderzoek onderweg, evenals voor het grondig proeflezen van onze gezamenlijke publicaties.

Aan Piet Bladt, Gabriel Ramos Llordén, Diana Giraldo Franco, en Jan Morez, met wie ik het genoegen had om succesvol samen te werken en meerdere leerrijke buitenlandse conferentiereizen te ondernemen. Aan Piet, voor onze gezamenlijke wetenschappelijke krachttoeren in ASL MRI, voor onze inzichtrijke gesprekken over werk en daarbuiten, voor de Messi poncho en Joan Miró herinneringen. Aan Gabriel, die me met Spaanse warmte en gastvrijheid leerde om te 'jaja'-en in MR relaxometrie en jezelf als onderzoeker te profileren. Aan Diana, voor onze samenwerking rond multiple sclerosis, de fijne concert- en after-workmomenten, en voor de spontane Zuid-Amerikaanse verzen. Aan Jan Morez, voor de pittige pepers, gelijkgestemde nuchterheid, en kunstzinnige insteken.

Aan mijn collega's Michele Nicastro en Céline Smekens, voor de vriendschappelijke samenwerking en vele laatavondse scanmomenten samen, voor de abstracte superresolutie discussies en verjaardagsdrinks, voor het steun- en toeverlaatschap langsheen de PhD weg. Aan Quinten Collier, Annelinde Buikema, Daniel Frenkel, Nicholas Franken, Pieter-Jan Vanthienen, met wie ik de voorbije jaren een kantoor deelde. Bedankt voor alle inspiratievolle interdisciplinaire gesprekken die we hadden, de dagelijkse dartscompetitie, de onderlinge steun naar elkaar, de soms diepe en leuke levensinzichten en vertelsels, de gezamenlijk aan het scherm gekluisterde concentratiemomenten in onze zoektocht naar eureka-momenten, de wederzijdse motivatiespeeches onderweg. Alle succes in jullie toekomstige verwezenlijkingen!

Aan de vele andere collega's bij imec-Vision Lab. Ik ben jullie dankbaar voor het creëren van een stimulerende en collaboratieve omgeving gedurende al deze jaren, waar wetenschappelijke ideeën en kennis ten volle mogen groeien en kruisbestuiven. Van harte mijn dank voor de kameraadschap en de vele fijne herinneringen. Ook mijn dank voor de vele aangename teambuildingmomenten die we naast het werk hebben gedeeld. Ik denk dan aan de Schotse Highland-games, de gezellige pub crawls, de sportieve padel- of voetbalwedstrijden, de fietsestafette in het Zoniënwoud. Mijn dank ook voor de flexibele samenwerking en wederzijdse team-spirit wanneer we werden uitgedaagd door cryptolocker-virus-aanvallen of de covid pandemie. Dank aan Jens Renders en Mateo Sierens, voor de nauwgezette samenwerking rond recurrent inference machines (RIMs). Aan Alice Presenti, voor de Italiaanse anekdotes en concertmomenten. Aan Banafshe en Vincenzo, voor de aangename babbels en goedlachsheid doorheen dit PhD avontuur. Aan de vele collega's van het B-Q MINDED team: Roberto, Maíra, Emanoel, Dimitrios, Marco, Riwaj, Seon Yeong, en Dennis, bedankt! Speciale dank aan Christine Waerebeek voor de ondersteuning tijdens de leerrijke Summer Schools omtrent kwantitatieve MRI. Aan de collega's met wie ik voor het leslokaal stond in de 1ste Bachelor Farmaceutische en Biomedische Wetenschappen: Jonathan, Jeroen, Femke, Lars-Paul, Nathanaël, Jannes, Caroline, Ben, en Eline. Dankjewel voor de fijne momenten samen, zowel tijdens het werk als daarnaast. Een speciale vermelding ook voor Paul, Vincent, Bikram, Anh-Tuan, Brian, Vera, Shabab, Kristina, Björn, Tim, Pavel, Zhihua, Marina, Mina, Domenico, Arpad, Van, Jian, Neam, en Maxime voor onze hartelijke lunch- en koffiepauzegesprekken.

Veel van het werk in dit proefschrift zou niet mogelijk zijn geweest zonder de MRI-datasets die mij ter beschikking werden gesteld of die ik kon verwerven met de hulp van vakbekwame medecollega's. Veel dank gaat uit naar de afdeling radiologie van het Universitair Ziekenhuis Antwerpen, in het bijzonder naar Floris Vanhevel, Céline Smekens en Thomas Janssens van Siemens Healthineers, en naar Michele Nicastro voor hun hulp bij het verwerven van de MR-relaxometriedatasets. Dank ook aan Thijs Van Osch en zijn team bij het C.J. Gorter MRI Centrum van het Leids Universitair Medisch Centrum, voor het verstrekken van de arteriële spin labeling data voor de laatste contributie van dit proefschrift, en voor de delicate feedback en het proeflezen in de review fase van onze publicatie in NeuroImage.

Ook mijn dankbaarheid aan de Universiteit Antwerpen om mij een Industrieel Onderzoeksfonds (IOF) toe te kennen om een deel van het onderzoek in dit proefschrift financieel te ondersteunen. Ik dank ook het μ neuro-consortium voor hun ondersteuning.

Aan alle andere mensen en collega-onderzoekers die op de een of andere manier aan dit PhD-avontuur hebben bijgedragen, mijn oprechte dank. Aan Jelle Veraart, voor onze verhelderende discussies over ultra-low field MRI en voor de hulp bij het schrijven en verdedigen van een postdoctorale financieringsaanvraag voor de Belgian American Educational Foundation (BAEF). Aan Tim Van De Looverbosch en Winnok De Vos, voor de inspirerende breinstormsessies over het combineren van superresolutie met light-sheet fluorescentiemicroscopie, een onderzoekslijn die ongetwijfeld nog van zich zal laten horen. Aan Pieter-Jan Guns en Wouter Van Putte, voor de hulp bij het initiëren van de IOF- en B-Q MINDED-projecten rond kwantitatieve MRI aan de Universiteit Antwerpen. Een speciale vermelding ook voor Jan en Arjan, Liesbeth Vanherp, Elisabeth Levrau, Bram Verbinnen, Pieter Van Dyck en Thomas Janssens, voor de hulp bij het schrijven en verdedigen van een postdoctoraal projectvoorstel voor VLAIO (Vlaams Agentschap Innoveren & Ondernemen). Het schrijven van zulk projectvoorstel en de vindingrijke ondernemerszin tijdens de diverse team-gebonden vergaderingen en gesprekken waren een openbaring en hebben me mede gesterkt in mijn ontwikkeling als wetenschapper. Aan Falk Mielke en Joaquim Sanctorum, voor de doorgedreven training met de X-stralen-tomografie opstelling van de Flex-CT scanner bij het imec-Vision Lab en de prettige samenwerking tijdens de regelmatige scansessies. De batches met Javavinken, varkenslongen of menselijke longkwabben die moesten worden gescand, vormden een atypische afleiding van mijn proefschrift.

Daarnaast wil ik een aantal belangrijke mensen en vrienden uit mijn directe omgeving bedanken. Mijn zaalvoetbalvrienden van DC United, dankzij wie ik wekelijks de wetenschap en het computerschermstaren kon vergeten tijdens intense sportmomenten en discussies over niet-wetenschappelijke aspecten van het leven. Dorothee, voor de toegewijde ondersteuning naar elkaar, jouw ondernemersmentaliteit, onze verhelderende gesprekken en artistiek gelach gevuld met de nodige porties levensinzichten. De liefde, voor zijn nemen en geven van energie die soms ontbrak en dan weer door de wolken heen scheen.

Veel dankbaarheid gaat ook uit naar mijn geliefde familie. In het bijzonder mijn ouders, Jef en Caroline, voor de onvoorwaardelijke kansen en vrijheid die ze mij voortdurend gaven om mijn eigen pad te vinden en mijn dromen na te streven, zowel in de sport als in de studie. Zonder jullie jarenlange liefde en steun had ik dit doctoraatswerk nooit kunnen voltooien. Het doet me deugd om te weten dat er een warme thuis is waar ik op elk moment terecht kan, om welke reden dan ook. Ook mijn broer en zus, Mattijs en Anna, verdienen een woord van dank voor hun steun en aanmoedigingen doorheen deze doctoraatsjaren en voor het steeds beschikbaar zijn voor interessante vragen over mijn onderzoek, met kruisbestuiving naar de biologie of het kunstzinnige. Ik besef hoe gelukkig ik mezelf mag prijzen met jullie als familie. Ook een woord van dank voor Flor, voor zijn gezelschaphoudende aaibaarheid.

Tot slot gaat mijn grote dank uit naar jou, Machteld, voor je steun en begrip tijdens de laatste eindsprint van het schrijven van dit proefschrift. Jouw energie en onnavolgbare glimlach vormden een krachtige motivatiebron om door te zetten. Ik ben oprecht dankbaar en gelukkig dat je als liefdevolle compagnon-de-route in mijn leven bent gewandeld.

Bedankt allemaal!

A)DuiiSx

Quinten Beirinckx, Antwerpen, April 2024

Contents

Su	ımma	ry		xv
Sa	menv	atting		xix
I	Pro	ologue		1
1	Prol 1.1 1.2 1.3 1.4	ogue Introdu The ne Toward The ur	uction	3 3 4 5 6
	1.5 1.6 Refe	Thesis Thesis rences	contributions	7 10 13
11	Ba	ickgro	und	15
2	Mag	jnetic R	Resonance Imaging: The basics	17
	2.1	Introdu	uction	17
	2.2	A brief	[•] history of MRI	18
	2.3	NMR	signal generation and detection	19
		2.3.1	Physical principles, nuclear spin and magnetic moment	19
		2.3.2	Macroscopic effect of the static magnetic field	21
		2.3.3		22
	0.4	2.3.4	Relaxation and readout strategies	24
	2.4			20
		2.4.1	2D readout	20
	25	Z.4.2 The M	IRI trade-off	-32 33
	2.5	251	Scan time	34
		2.5.2	Spatial resolution	35
		2.5.3	Signal-to-noise ratio	38
	Refe	rences	~ 	39

3	The	Advent	of Quantitative MRI		43
	3.1	Introdu			43
	3.2	Basic o	concept of quantitative MRI:		44
	3.3	Clinica	I interests of quantitative MRI		46
	3.4	MR Re			48
		3.4.1			48
		3.4.2	Biophysical basis of relaxation		48
		3.4.3	Methods for T_1 and T_2 measurements		49
		3.4.4	Quantitative relaxometry mapping		52
		3.4.5	Clinical applications		53
		3.4.6	Existing problems in MR relaxometry		55
			3.4.6.1 Motion		55
			3.4.6.2 Specific absorption rate		57
			3.4.6.3 Acquisition time constraints		57
			3.4.6.4 Slice profile variation		58
			3.4.6.5 Need for protocol standardization and metrology		58
	3.5	Arteria	I Spin Labeling		60
		3.5.1	Introduction		60
		3.5.2	Basic principle		61
		3.5.3	Recommended implementation		63
		0.0.0	3.5.3.1 Signal generation		63
			3.5.3.2 Perfusion parameter quantification		68
		351	Clinical applications		73
		3.5.4 3.5.5	Existing problems in ASI		7/
	36	Dorom	Existing problems in ASE		76
	5.0		Statistical parameter estimators		70
		3.0.1	2.6.1.1 Accuracy and precision		70
			3.6.1.1 Accuracy and precision		70
			3.6.1.2 (Root) mean squared error		11
		3.6.2			78
		3.6.3	Least squares estimators		79
		3.6.4	Bayesian estimators		80
			3.6.4.1 Bayes theorem		80
			3.6.4.2 Maximum a posteriori estimator		81
	Refe	rences .			82
4	Supe	er-resoli	ution reconstruction as prime protagonist for accelerated	(q)MRI	93
	4.1	Introdu	lction		94
		4.1.1	Resolution challenges in MRI		94
		4.1.2	Resolution enhancement techniques		96
		4.1.3	Super-resolution reconstruction		97
	4.2	Super-	resolution reconstruction applied to MRI		98
		4.2.1	Concept and definition		99
		4.2.2	Historical misconceptions		99
		4.2.3	Key components		100
		2.0	4 2 3 1 Acquisition strategies		100
			4 2 3 2 Imaging model		103
			4233 Reconstruction		106
	42	The no	ed for robust motion compensation		111
	T.J Rofo	rences			112

III Contributions

117

5	Joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework : a simulation					
	stud	У		1	19	
	5.1	Introdu	ction		21	
	5.2	Theory			23	
		5.2.1	MR imaging model		23	
		5.2.2	The joint Maximum Likelihood estimator		25	
		5.2.3	Implementation		26	
	5.3	Simula	ion experiments		28	
	5.4	Results	and discussion		33	
	5.5	Conclu	sion		38	
	Appe	endices			39	
	5.A	Downs	impling and upsampling operators		39	
	Refe	rences .			40	
6	Мос	lel-base	l super-resolution reconstruction with jo	int motion estimation for		
	impr	oved qu	antitative MRI parameter mapping	14	43	
	6.1	Introdu	ction	14	45	
	6.2	Theory		14	47	
		6.2.1	Forward model	14	47	
		6.2.2	Joint Bayesian estimation framework	14	48	
			6.2.2.1 Bayes theorem	14	48	
			6.2.2.2 Maximum a posteriori estimator	14	48	
			6.2.2.3 Likelihood function	14	48	
			6.2.2.4 Prior distributions	14	49	
			6.2.2.5 Alternating minimization	1!	50	
			6.2.2.6 Regularization parameter selection	on 1!	50	
	6.3	Materials and Methods				
		6.3.1	Whole Brain Simulations		53	
		6.3.2	In Vivo Data		54	
		6.3.3	Quantitative Image Analysis	1!	58	
		6.3.4	Implementation	1!	59	
	6.4	Results		10	60	
		6.4.1	Whole brain simulations	10	б0	
		6.4.2	In Vivo Data	10	63	
	6.5	Discuss	ion	1	72	
	6.6	Conclu	sion	1	74	
	Appe	endices		1	74	
	6.A	Analyti	cal derivatives for joint MAP optimization	1	74	
		6.A.1	MAP estimation of motion parameters .	1	75	
		6.A.2	MAP estimation of tissue parameters	1	76	
	6.B	Model	operators: implementation and derivatives	1	78	
		6.B.1	Warping operators and derivatives	1	78	
			6.B.1.1 Extended implementation details	;	80	
			6.B.1.2 Computational complexity analys	sis	81	
			6.B.1.3 Alternative image warping impler	mentations 18	82	

		6.B.2	Blurring operator	183
		6.B.3	Resampling operator	184
			6.B.3.1 Kernel extensions	185
		6.B.4	Operator validation	186
			6.B.4.1 Matrix transpose test	186
			6.B.4.2 Dot-product test	186
	6.C	Applica	ations in musculoskeletal imaging	188
		6.C.1	Anatomical knee MRI	188
		6.C.2	Quantitative knee MRI	192
	Refe	rences .		195
7	A su	per-res	olution reconstruction framework for quantitative brain perfusion	
	map	ping usi	ing pseudo-continuous Arterial Spin Labeling	201
	7.1	Introdu	uction	203
	7.2	Theory		204
		7.2.1	Super-resolution reconstruction forward model	204
		7.2.2	Single-PLD pCASL signal model	205
		7.2.3	Joint Bayesian estimation framework	206
			7.2.3.1 Maximum a posteriori estimator	206
			7.2.3.2 Optimization	207
			7.2.3.3 Implementation	208
	7.3	Metho	ds	208
		7.3.1	Simulation experiments	209
			7.3.1.1 Synthetic data generation	210
			7.3.1.2 Parameter estimation	212
			7.3.1.3 Performance analysis	215
		732	Real data experiment	216
	74	Results	s	218
		741	Simulation experiments	218
			7 4 1 1 Simulation experiments without motion	218
			7 4 1 2 Simulation experiments with motion	210
		742	Real data experiment	221
	75	Discuss		236
	1.5	7 5 1	Differences with existing techniques	236
		752	Improved CBE quantification from single-PLD pCASL data	236
		753	loint estimation of CBE and motion parameters	230
		754	Model assumptions and prospective extensions	230
		7.5. 4 7.5.5		233
		7.5.5	Anticipated clinical impact and future perspectives	241
	76	7.5.0 Conclu	sion	242 012
	7.0 Anno	Conciu	SIOIT	243
	Appe		anandant DI D	243
	1.A	Slice-di		243
	1.B	васкдг		244
	1.C	Linear	Torward model	245
	1.D	Extra s	Simulation experiment	247
	(.E	Notion	i parameter estimates	247
	Kete	rences .		250

8	Conclusions and Future Perspectives References	253 257		
List of Abbreviations				
Aca	ademic overview	263		

Summary

This dissertation deals with the use of model-based super-resolution reconstruction (SRR) with joint patient motion estimation for the purpose of improved quantitative magnetic resonance imaging (qMRI), which can provide accurate, precise, and accessible biomarkers for clear numerical differentiation of brain disease states. The manuscript consists of eight chapters, which are divided over three main parts. In **Part I** (Prologue), the main **motivation and existing challenges** of qMRI as a medical imaging technique are introduced. **Part II** (Background) provides the necessary **background material** to the research areas in which the contributions of this work are situated, namely, magnetic resonance imaging (MRI), the extension of MRI to qMRI, and SRR as an advanced imaging tool. **Part III** (Contributions) then provides an overview of the main **contributions** of this thesis.

Prologue

In **Chapter 1**, a comprehensive overview is provided regarding the positioning of this thesis research within the realm of modern medical imaging applications, particularly focusing on brain MRI. With society experiencing rapid growth and aging, there is a heightened prevalence of neurodegenerative disorders and age-related diseases. This accentuates the urgent need for methods that enable timely disease detection, monitor disease progression, and evaluate the effectiveness of new therapies. This urgency underscores the need for reliable and easily accessible quantitative biomarkers capable of identifying diseases prior to the onset of clinical symptoms. While MRI is renowned for its exceptional soft tissue contrast and inherent patient safety, its widespread utilization as a biomarker detection tool encounters a critical challenge. Conventional MRI relies on gualitative image contrast evaluation, complicating the quantitative comparison of tissue properties within and between scans or subjects. Transitioning to quantitative MRI (qMRI) is imperative to overcome these limitations, enabling absolute quantification of tissue characteristics independent of experimental design, thereby enhancing diagnostics. Unfortunately, the dissemination of gMRI faces challenges such as low spatial resolution, low signal-to-noise ratio (SNR), and long scan times. These extended scan times, required to compensate for the low SNR and spatial resolution, can impact patient comfort and compliance, increase the risk of motion artifacts, and reduce patient throughput. To address the need for rapid MRI techniques without compromising spatial resolution or SNR, this thesis explores the use of model-based SRR. The application of SRR for qMRI is challenging, however, and a number of technical hurdles remain, which this thesis aims to tackle.

Background

A solid grasp of the basic principles of MRI is crucial for understanding the contributions presented in this thesis. Starting from an overview of the historical advancements in MRI, **Chapter 2** therefore briefly explains some fundamental principles and key concepts about MRI. In particular, the phenomenon of nuclear magnetic resonance is considered, which is rooted in quantum mechanical considerations yet interpretable through a classical lens when analyzing the interaction of a large number of hydrogen atoms in and width an external magnetic field, radio frequency (RF) waves and magnetic field gradients. The chapter also introduces crucial concepts such as excitation and relaxation of magnetic spins, with a focus on parameters vital for quantitative neuroimaging research, including T1 and T2 relaxation times. Furthermore, it outlines the process of signal generation and detection in MRI, highlighting elements pertinent to the modalities and methodologies employed in this thesis. In particular, the use of 2D multi-slice and 3D acquisition strategies is described, pinpointing their respective advantages and disadvantages. Finally, the chapter provides a detailed list of the specific MRI pulse sequences used for the acquisition of the *in vivo* whole brain data in the research contributions.

Chapter 3 outlines the significance of quantitative MRI (qMRI) in response to shortcomings of and as an addition to conventional MRI. First, the chapter provides a succinct overview of the fundamental concept of qMRI, focusing on T_1 relaxation parameter mapping. Subsequently, it delves into several key clinical advantages of gMRI, including enhanced tissue and pathology characterization, early detection of pathophysiological changes, longitudinal patient monitoring, and multi-centric assessment. Moreover, the chapter examines MR relaxometry and arterial spin labeling (ASL) MRI as two pivotal clinical applications of qMRI, used to quantify relaxation times and tissue perfusion, respectively. Each application's underlying principle is thoroughly elucidated, providing essential insights into both the acquisition of contrast-weighted data for qMRI and the quantification of resultant qMRI parameters. Moreover, existing limitations of MR relaxometry and ASL are explained, with this thesis aiming to propose viable solutions. Additionally, the chapter underscores the clinical relevance of each application. Finally, recognizing the centrality of parameter estimation in qMRI, the chapter concludes by introducing basic principles of parameter estimation, essential for the accurate and precise derivation of biophysical quantitative parameters from contrast-weighted MRI scans.

In **Chapter 4**, the role of **super-resolution reconstruction** is elaborated. Starting with a general discussion on the challenges of resolution in MRI and the consequent demand for resolution-enhancing methods, the general concept of SRR is introduced. In addition to the fundamental principle of this technique, the key components are highlighted necessary for SRR application in MRI. Particularly, the choice of an appropriate acquisition strategy, the selection of an accurate imaging model that realistically models both the acquisition and noise, and finally, the model-based reconstruction process addressing technical aspects such as regularization and hyperparameter selection. At the end of this chapter, the need for robust motion compensation is also briefly touched upon, which is indispensable for ensuring the spatial alignment of the low-resolution image set on which SRR is performed.

Contributions

Chapter 5 presents an **extensive Monte Carlo simulation study** on small-image 'checkerboard'like phantoms in which for the first time model-based SRR for qMRI is augmented with joint motion estimation, leveraging **T1 mapping** as an MR relaxometry model of choice. A **joint Maximum Likelihood estimation** framework is employed, optimizing motion and relaxometry parameters estimates alternately. Additionally, this chapter combines downsampling and blurring with a rectangular slice profile in one and the same forward operator to streamline computational efficiency. The proposed SRR framework, featuring joint motion estimation, is extensively compared against a framework without motion estimation and a previously documented SRR T1 mapping method employing a motion pre-registration strategy. This chapter serves as an initial proof-of-concept for model-based SRR with joint motion estimation in qMRI applications, setting the stage for further development of a large-scale and more computationally intensive model-based SRR framework tailored for whole-brain *in vivo* images, as detailed in the subsequent chapter.

Chapter 6 expands the previous simulation study and proof-of-concept demonstration further towards a widely applicable framework for model-based super-resolution reconstruction with joint estimation of motion parameters and isotropic, high-resolution 3D quantitative MRI parameter maps from motion-corrupted, low-resolution 2D multi-slice MRI scans. Representing the pinnacle of this dissertation's contributions, this chapter combines all important technological developments into one polyvalent and robust estimation framework. Central to this framework is a Bayesian approach that leverages prior knowledge of the tissue and noise statistics, in which special attention is given to realistic modeling of noise and data distributions by use of a Rician probability distribution in the likelihood function of the magnitude MR data. The framework's potential is demonstrated in both simulations and real data experiments, using **T1 and T2 mapping** as carrying examples. Attention is also given to the experimental design of the acquisition protocols for relaxometry mapping, in which contrast weighting and (geometric) distribution of the individual images in the quantitative image series is subject to restrictions inherent to the MR acquisition sequence of interest. Specifically, this chapter explores the use of a multi-echo spin-echo (MESE) sequence for acquisition of a 2D multi-slice super-resolution image series for T2 mapping, while also acknowledging limitations of such a sequence. It should be highlighted that the chapter's sections and formulas are written with modularity in mind, facilitating substitution of alternative quantitative signal and/or noise models. In fact, this modularity in signal and noise model was demonstrated in the subsequent chapter, where a perfusion model with a Gaussian noise model is combined within the framework. The appendices of this second contribution chapter compile essential mathematical components, a.k.a 'building blocks', and tools that were developed for the implementation of the proposed motion-robust super-resolution reconstruction framework for gMRI. Key aspects include implementing linear operators that constitute the super-resolution forward model, qMRI signal models describing the relationship between signal intensity and underlying tissue parameters in a voxel, and practical considerations for solving a large-scale parameter estimation problem with joint estimation of qMRI and motion parameters. Finally, the chapter showcases some use cases of anatomical and quantitative SRR in musculoskeletal MRI, reconstructed with the proposed framework. As opposed to brain MRI, musculoskeletal (MSK) MRI focuses specifically on joint structures, including wrists, ankles, knees, etc. It is demonstrated that SRR can also play an important role to improve existing 3D resolutions of MSK MRI without significantly increasing the scan time of clinical protocols.

Chapter 7 presents a novel super-resolution reconstruction framework to estimate 3D isotropic high-resolution quantitative cerebral blood flow (CBF) maps from a series of single-post-labeling-delay (single-PLD) pseudo-continuous Arterial Spin Labeling (ASL) control-label image pairs, each acquired with low through-plane resolution and rotated sliceencoding direction in a 2D multi-slice readout scheme. Building upon the SRR framework introduced in Chapter 6, motion between control and label images was jointly estimated in a Bayesian estimation framework, enabling accurate and precise CBF quantification without propagation of pre-registration errors, while optimally exploiting prior knowledge of tissue properties and noise statistics. The rotation of the slice-encoding direction for each control-label image pair as well as the lower through-plane resolution ensures a more uniform distribution of the PLD throughout the brain and increases the effectiveness of background suppression. Combined, this significantly improves the SNR compared to conventional 2D multi-slice readout with direct high through-plane resolution, where CBF quantification has traditionally been hampered by detrimental perfusion SNR slice dependence in sequentially acquired slices. The proposed method was validated both qualitatively and quantitatively in synthetic whole brain simulations and on *in vivo* human brain data. It has been demonstrated that the framework provided superior CBF estimation in terms of root-mean-square error compared to a state-of-the-art approach using a conventional 2D multi-slice readout strategy with ascending slice order and isotropic resolution in the same scan time, even when additional hardware acceleration techniques like multiband are applied in the latter.

Finally, **Chapter 8** concludes the thesis and outlines potential future research directions.

Samenvatting

Deze dissertatie behandelt het gebruik van modelgebaseerde superresolutie-reconstructie (SRR) met gezamenlijke schatting van patiëntenbeweging, voor de verbetering van kwantitatieve magnetische resonantiebeeldvorming (qMRI). Deze qMRI techniek kan nauwkeurige, precieze en toegankelijke biomarkers genereren voor een duidelijke numerieke differentiatie van de toestand van hersenziekten. Het manuscript bestaat uit acht hoofdstukken, die verdeeld zijn over drie delen. In **Deel I** (Proloog) worden de belangrijkste **motivatie en bestaande uitdagingen** van qMRI als medische beeldvormingstechniek geïntroduceerd. **Deel II** (Achtergrond) biedt het noodzakelijke **achtergrondmateriaal** voor de onderzoeksgebieden waarin de bijdragen van dit werk zich situeren, namelijk magnetische resonantie beeldvorming (MRI), de uitbreiding van MRI naar qMRI, en SRR als geavanceerde beeldvormingstechniek. **Deel III** (Bijdragen) geeft vervolgens een overzicht van de belangrijkste **bijdragen** van dit PhD onderzoek.

Proloog

In **Hoofdstuk 1** wordt een uitgebreid overzicht gegeven van de positionering van het onderzoek van dit proefschrift binnen het domein van moderne medische beeldvormingstoepassingen, met specifieke focus op MRI van de hersenen. Door de snelle groei en veroudering van de samenleving neemt de prevalentie van neurodegeneratieve aandoeningen en leeftijdsgebonden ziekten toe. Dit onderstreept de dringende noodzaak voor methoden die niet alleen vroegtijdige ziektedetectie mogelijk maken, maar ook het verloop van ziekten kunnen monitoren en de effectiviteit van nieuwe therapieën kunnen evalueren. Deze urgentie benadrukt de behoefte aan betrouwbare en gemakkelijk toegankelijke kwantitatieve biomarkers die in staat zijn om ziekten te identificeren vóór aanvang van klinische symptomen. Hoewel MRI bekend staat om zijn uitzonderlijke contrast in zacht weefsel en intrinsieke veiligheid voor patiënten, stuit het wijdverbreide gebruik ervan als een biomarkerdetectietool op een kritieke uitdaging. Conventionele MRI vertrouwt op kwalitatieve evaluatie van beeldcontrast, wat de kwantitatieve vergelijking van weefseleigenschappen binnen en tussen scans of proefpersonen compliceert. De overgang naar kwantitatieve MRI (gMRI) is essentieel om deze beperkingen te overwinnen, zodat absolute kwantificering van weefselkenmerken onafhankelijk van experimenteel ontwerp mogelijk is, en daarmee de diagnostiek wordt verbeterd. Helaas wordt het gebruik van qMRI geconfronteerd met uitdagingen zoals lage spatiale resolutie, lage signaal-ruisverhouding (SNR), en lange scantijden. Deze verlengde scantijden, nodig om de lage SNR en spatiale resolutie te compenseren, kunnen het patiëntencomfort en -conformiteit beïnvloeden, het risico op bewegingsartefacten vergroten, en de dagelijkse doorstroom van patiënten verminderen. Om te voldoen aan de behoefte aan snelle MRI-technieken zonder de spatiale resolutie of SNR in gevaar te brengen, onderzoekt dit proefschrift het gebruik van modelgebaseerde SRR. De toepassing van SRR voor qMRI is echter niet vanzelfsprekend, en een aantal technische uitdagingen blijven bestaan, welke dit proefschrift beoogt op te lossen.

Achtergrond

Het begrijpen van de **basisprincipes van MRI** is een voorwaarde om de bijdragen in dit proefschrift te bestuderen. Vanuit een overzicht van de historische vooruitgang in MRI, verklaart Hoofdstuk 2 kort enkele fundamentele principes en sleutelconcepten over MRI. In het bijzonder wordt het fenomeen van nucleaire magnetische resonantie beschouwd, dat zijn basis vindt in kwantummechanische overwegingen maar interpreteerbaar is vanuit een klassieke benadering wanneer de interactie van een groot aantal waterstofatomen in en met een extern magnetisch veld, radiofreguentie (RF) golven en magnetische veldgradiënten wordt beschouwd. Het hoofdstuk introduceert ook cruciale concepten zoals excitatie en relaxatie van magnetische spins, met een focus op parameters die essentieel zijn voor kwantitatief neuroimaging-onderzoek, waaronder T1- en T2-relaxatietijden. Verder schetst het hoofdstuk de processen van signaalgeneratie en -detectie in MRI, waarbij elementen worden benadrukt die relevant zijn voor de modaliteiten en methodologieën die in dit proefschrift worden gebruikt. In het bijzonder wordt de toepassing van 2D multi-slice en 3D-acquisitiestrategieën beschreven, waarbij hun respectievelijke voor- en nadelen worden benadrukt. Tot slot biedt het hoofdstuk een gedetailleerde lijst van de specifieke MRI-pulssequenties die zijn gebruikt voor de opnames van de *in vivo* data van het brein in de onderzoeksbijdragen.

Hoofdstuk 3 schetst het belang van kwantitatieve MRI (qMRI) als reactie op tekortkomingen van en als aanvulling op conventionele MRI. Allereerst biedt het hoofdstuk een beknopt overzicht van het fundamentele concept van qMRI, met de nadruk op T_1 relaxatieparameter mapping. Vervolgens gaat het in op verschillende belangrijke klinische voordelen van qMRI, waaronder verbeterde weefsel- en pathologiekarakterisering, vroegtijdige detectie van pathofysiologische veranderingen, longitudinale patiëntenmonitoring en multi-center evaluatie. Bovendien onderzoekt het hoofdstuk MR relaxometrie en arteriële spin labeling (ASL) MRI als twee cruciale klinische toepassingen van qMRI, die respectievelijk worden gebruikt om relaxatietijden en perfusie te kwantificeren. Het onderliggende principe van elke toepassing wordt grondig toegelicht, waardoor essentiële inzichten worden verkregen in zowel de verwerving van contrast-gewogen scans voor qMRI als de kwantificering van resulterende qMRI parameters. Bovendien worden bestaande beperkingen van MR relaxometrie en ASL toegelicht, waarbij dit proefschrift streeft naar het voorstellen van haalbare oplossingen. Daarnaast benadrukt het hoofdstuk de klinische relevantie van elke toepassing. Ten slotte, in erkenning van de centrale rol van parameter schatting in gMRI, concludeert het hoofdstuk door de **basisprincipes van** parameterschatting te introduceren, essentieel voor de nauwkeurige en precieze schatting van biofysische kwantitatieve parameters uit contrast-gewogen MRI-scans.

In **Hoofdstuk 4** wordt de rol van **superresolutie-reconstructie** uitgewerkt. Beginnend met een algemene discussie over de uitdagingen van beeldresolutie in MRI en de daaruit voortvloeiende vraag naar resolutieverbeterende methoden, wordt het algemene concept van SRR geïntroduceerd. Naast het fundamentele principe van deze techniek worden de belangrijkste componenten belicht die nodig zijn voor de toepassing van SRR in MRI. Met name de keuze van een geschikte acquisitiestrategie, de selectie van een nauwkeurig beeldvormingsmodel dat zowel het acquisitieproces als de ruis realistisch modelleert, en tot slot het modelgebaseerde reconstructieproces dat technische aspecten zoals regularisatie en hyperparameterselectie behandelt. Aan het einde van dit hoofdstuk wordt ook kort ingegaan op de noodzaak van robuuste bewegingscompensatie, die onmisbaar is voor het waarborgen van de spatiale uitlijning van de lage-resolutie beeldenset waarop SRR wordt uitgevoerd.

Bijdragen

Hoofdstuk 5 presenteert een **uitgebreide Monte Carlo-simulatiestudie** op kleine 'schaakbord'achtige fantomen waarin voor het eerst modelgebaseerde SRR voor qMRI wordt uitgebreid met gezamenlijke bewegingsschatting, waarbij gekozen wordt voor **T1-mapping** als MRrelaxometriemodel. Een **gezamenlijk Maximum Likelihood-schattingsraamwerk** wordt toegepast, waarbij afwisselend bewegings- en relaxometrieparameters worden geoptimaliseerd. Daarnaast combineert dit hoofdstuk downsampling en blurring met een rechthoekig sliceprofiel in één en dezelfde voorwaartse operator om de computationele efficiëntie te verbeteren. Het voorgestelde SRR-raamwerk, met gezamenlijke bewegingsschatting, wordt uitgebreid vergeleken met een raamwerk zonder bewegingsschatting en een eerder beschreven SRR T1-mappingmethode die een pre-registratiestrategie gebruikt. Dit hoofdstuk dient als een eerste proof-of-concept voor modelgebaseerde SRR met gezamenlijke bewegingsschatting in qMRI-toepassingen, en vormt de basis voor verdere ontwikkeling van een grootschalig en meer rekenintensief modelgebaseerd SRR-raamwerk dat is afgestemd op *in vivo* beelden van een volledig brein, zoals gedetailleerd beschreven in het volgende hoofdstuk.

Hoofdstuk 6 breidt de voorgaande simulatiestudie en proof-of-concept-demonstratie verder uit naar een breed toepasbaar raamwerk voor modelgebaseerde superresolutie-reconstructie met gezamenlijke schatting van bewegingsparameters en isotropische, hoge-resolutie 3Dkwantitatieve MRI-parametermappen uit door beweging aangetaste, lage-resolutie 2D-multislice MRI-scans. Als hoogtepunt van de bijdragen van deze dissertatie combineert dit hoofdstuk alle belangrijke technologische ontwikkelingen in één polyvalent en robuust schattingsraamwerk. Centraal in dit raamwerk staat een Bayesiaanse benadering die gebruik maakt van voorkennis van de weefsel- en ruisstatistieken, waarbij speciale aandacht wordt besteed aan realistische modellering van ruis- en dataverdelingen door gebruik te maken van een Riciaanse kansverdeling in de waarschijnlijkheidsfunctie van de magnitude MRI data. Het potentieel van het raamwerk wordt gedemonstreerd in zowel simulaties als in experimenten met echte hersendata, waarbij T1- en T2-mapping als belangrijke voorbeelden dienen. Er wordt tevens aandacht besteed aan het experimenteel ontwerp van de acquisitieprotocollen voor relaxometrische mapping, waarbij contrastweging en (geometrische) verdeling van de individuele beelden in de kwantitatieve beeldenreeks onderworpen zijn aan beperkingen inherent voor de gekozen MR-pulssequentie. Specifiek onderzoekt dit hoofdstuk het gebruik van een multi-echo spin-echo (MESE)-sequentie voor de acquisitie van een 2D-multi-slice superresolutie-beeldreeks voor T2-mapping, terwijl ook beperkingen van een dergelijke sequentie worden erkend. Het is belangrijk om te vermelden dat de tekst en formules van dit hoofdstuk zijn geschreven vanuit een modulair oogpunt. Hierdoor is de methodologie ook eenvoudig transfereerbaar naar alternatieve kwantitatieve signaal- en/of ruismodellen. Meer nog, deze modulariteit in signaal- en ruismodel wordt gedemonstreerd in het volgende hoofdstuk, waarin een perfusiemodel met een Gaussisch ruismodel binnen het raamwerk wordt gecombineerd. De appendices van dit tweede Hoofdstuk 6 bundelen essentiële wiskundige componenten, zogenaamde 'bouwstenen', en tools die zijn ontwikkeld voor de implementatie van het voorgestelde bewegingsrobuuste superresolutie-reconstructieraamwerk voor qMRI. Belangrijke aspecten zijn onder meer de implementatie van lineaire operatoren die het superresolutie-voorwaartsmodel vormen, gMRI-signaalmodellen die de relatie tussen signaalintensiteit en onderliggende weefselparameters in een voxel beschrijven, en praktische overwegingen voor het oplossen van een grootschalig parameterschattingsprobleem met gezamenlijke schatting van qMRI- en bewegingsparameters. Tot slot toont het hoofdstuk

enkele specifieke toepassingen van het gebruik van anatomische en kwantitatieve SRR in musculoskeletale MRI, gereconstrueerd met het gepresenteerde raamwerk. In tegenstelling tot hersen-MRI richt musculoskeletale (MSK) MRI zich specifiek op gewrichtsstructuren, inclusief polsen, enkels, knieën, enz. Het wordt aangetoond dat SRR ook een belangrijke rol kan spelen om de bestaande 3D-resoluties van MSK MRI te verbeteren zonder de scantijd van klinische protocollen aanzienlijk te verhogen.

Hoofdstuk 7 presenteert een nieuw superresolutie-reconstructieraamwerk om 3D-isotrope hoge-resolutie kwantitatieve cerebrale bloedstroom (CBF) mappen te schatten uit een reeks post-labeling-delay (single-PLD) pseudo-continue Arteriële Spin Labeling (ASL) controle-label beeldparen, elk verworven met een lage through-plane resolutie en gedraaide slice richting in een 2D-multislice uitleesschema. Voortbouwend op het SRR-raamwerk geïntroduceerd in Hoofdstuk 6, wordt de beweging tussen controle- en labelbeelden gezamenlijk geschat in een Bayesiaans schattingsraamwerk, waardoor nauwkeurige en precieze CBFkwantificatie mogelijk wordt zonder propagatie van pre-registratiefouten, terwijl optimaal gebruik wordt gemaakt van voorkennis van weefseleigenschappen en ruisstatistieken. De rotatie van de slice richting voor elk controle-label beeldpaar, evenals de lagere through-plane resolutie, zorgt voor een meer uniforme verdeling van de PLD over het hersengebied en verhoogt de effectiviteit van achtergrondonderdrukking. Gecombineerd verbetert dit aanzienlijk de SNR in vergelijking met conventionele 2D-multislice uitlezing met directe hoge through-plane resolutie, waar CBF-kwantificatie traditioneel wordt gehinderd door nadelige perfusie-SNR-slice-afhankelijkheid in seguentieel verworven slices. De gepresenteerde methode werd zowel kwalitatief als kwantitatief gevalideerd in synthetische volledige hersensimulaties en op in vivo menselijke hersendata. Er wordt aangetoond dat het raamwerk superieure CBF-schatting biedt in termen van root-mean-square error in vergelijking met een state-of-the-art aanpak die een conventionele 2D-multislice uitleestrategie hanteert met oplopende slicevolgorde en isotrope resolutie in dezelfde scantijd, zelfs wanneer aanvullende hardwareversnellende technieken zoals multiband worden toegepast in de laatstgenoemde methode.

Tot slot worden in **Hoofdstuk 8** de conclusies van deze thesis gebundeld en worden potentiële toekomstige onderzoekspaden aangestipt.

Part I

Prologue

Prologue

Contents

1.1 Introduction	3
1.2 The need for biomarkers for neurodegenerative diseases	4
1.3 Towards absolute quantification of biomarkers	5
1.4 The unmet need for technical development	6
1.5 Thesis contributions	7
1.6 Thesis organisation	10
References	13

1.1 Introduction

With nearly 100 billion neurons and 100 trillion connections, **the human brain** remains one of the most intriguing mysteries in science and one of the greatest challenges in medicine. Through multiple studies, science is trying to unravel the highly complex anatomy and functions of both healthy and diseased brains. Whereas originally information about the neural architecture and functionality of the human brain could only be obtained through histological post-mortem studies, it is now possible to study the human brain completely *in vivo*. Thanks in part to advanced developments in medical imaging techniques, with magnetic resonance imaging (MRI) playing a prominent role, the brain puzzle has been systematically further unravelled. However, the medical world is not standing still and new insights into brain disorders remain necessary, coupled with new expectations for technological progress. Hence, this prologue briefly outlines a number of important components in contemporary MRI research, and provides a general view of the overarching motivation for a physicist, scientist,

and modern-day PhD researcher to dive into the field of MRI. Part of that motivation undeniably stems from the conviction to answer stringent questions in a society and world in which the fraction of the elderly population rises significantly year after year, and where the call for large-scale screening and early detection of age-related diseases is getting increasingly louder. Healthcare organization needs to be thoroughly rethought in the upcoming decades, and MRI can play a central role in this transition. However, some critical technological hurdles have to be overcome, which will be touched upon in the next paragraphs. At the end of this introductory consideration, the main contributions of this thesis are summarized, and the structure of the thesis manuscript is briefly explained.

1.2 The need for biomarkers for neurodegenerative diseases

We are protagonists in an aging era. While we are facing historical challenges of COVID-19, we are also facing significant challenges of a **fast-growing aging society**. According to a United Nations report, 1 in 10 people was over 65 in 2022 and by 2050 the ratio will be almost doubled to 1 in 6 (United Nations, Department of Economic and Social Affairs, Population Division, 2022). As the global population ages, the prevalence of neurodegenerative disorders and age-related diseases (NDAD) is fast increasing. Examples of NDAD include Alzheimer's disease and related dementias, Parkinson's disease, and motor neuron diseases (Deuschl et al., 2020; Feigin et al., 2020). Moreover, neurological disorders are the leading cause of disability worldwide, and the second leading cause of death globally, accounting for approximately 9 million deaths per year (World Health Organization, 2022). The World Health Organization estimates that half of the worldwide economic impact of disability will be due specifically to brain-related conditions by 2030 if we do not change this trajectory (Mathers & Loncar, 2006). In low- and middle-income countries, which have a higher population growth rate, age-related diseases and disorders will pose an even more severe threat to their development (Winkler, 2020). The impact of age-related diseases on individuals living with a disease and their caregivers, families, and the society at large cannot be underestimated given its physical, psychological, social, and economical burden.

Timely detection of NDAD before the manifestation of clinical symptoms is paramount to prevent or delay their progression. Considering that prevention, delay, and/or treatment is more likely to be successful for patients in the earliest phases of their disease, it is important to discover reliable and accessible biomarkers that can detect NDAD prior to the manifestation of the disease. Indeed, an early detection is necessary to maximize the therapeutic window or to enroll such patients in clinical trials to promote the steady progression of treatment development and evaluation (Hansson, 2021). In recent years, large efforts have been made to discover biomarkers that identify neurodegenerative diseases earlier, more easily, and more accurately. Here, the term biomarker refers to 'objective characteristics that are measured as an indicator of normal biological processes, pathogenic processes, or responses to an exposure or intervention, including therapeutic interventions' (Hockings et al., 2020). However, current diagnostic biomarkers for NDAD, including MRI markers, are often invasive, require specialized personal or expensive hardware (Hansson, 2021; Teunissen et al., 2022). These constraints, together with financial and logistical issues limit broad-based implementation of these biomarkers for wide application for screening in primary care settings. The growing responsibility of primary care physicians and care teams

in the screening and detection of diseases is imperative in this era of an aging population. Therefore, to promote the early detection of at-risk individuals, there is a need to identify accessible and scalable biomarkers of brain health that can be obtained regularly in the general population at point-of-care facilities.

1.3 Towards absolute quantification of biomarkers

Magnetic resonance imaging (MRI) is a powerful, noninvasive, medical imaging modality distinguished by its ability to provide excellent image contrast in soft tissues. Consequently, MRI is prominently used in diagnostic medicine and biomedical research. Unlike other medical imaging modalities such as computed tomography (CT) or positron emission tomography (PET) imaging, MRI does not require the use of ionizing radiation, making it a relatively safe imaging modality for human study. As the name implies, MRI scanners rely on the use of strong magnetic fields, both static and dynamic, together with radio frequency (RF) waves, to generate images. Due to its harmless nature, MRI is used extensively as a diagnostic tool for scientific research. It is therefore not surprising that MRI scanners can be found in almost all major hospitals in the Western world. MRI is also a very versatile technique. By using dedicated acquisition sequences, MRI is able to display a variety of physical phenomena and provide good contrast for some specific organs or tissue types. Examples include: imaging of brain activity (functional MRI), blood vessels (MR angiography), metabolic changes (MR spectroscopy), directional information of tissue structures (diffusion MRI), biochemical structure specific relaxation properties of tissue (MR relaxometry), and blood perfusion (**perfusion MRI**). The latter two are of particular interest in this thesis.

Conventional, anatomical MRI consists of the qualitative evaluation of image contrast – i.e. relative local signal intensity differences in images. This image contrast depends on many different factors, including not only the underlying biophysical tissue properties of interest, but also the MRI hardware and software. This does not hinder visual inspection of anatomy, but it makes conventional MRI subjective: the expertise of radiologists plays a key role in the evaluation of contrast-weighted images for disease diagnosis and monitoring. Moreover, conventional MRI hampers quantitative comparison of tissue properties within a scan, between successive scans, and between subjects.

A long-standing goal in the magnetic resonance (MR) community has been quantitative imaging, where properties of interest are quantitatively mapped, and image interpretation is both anatomical and numerical. **Quantitative MRI (qMRI)** enables absolute quantification and mapping of biophysical tissue characteristics, completely independent of experimental design (Deoni, 2010). These characteristics, i.e. *parameters*, which are extracted from a set of MR images with varying contrast settings, can be expressed as numbers with absolute physical units. As such, this allows for a more objective comparison (across subjects, protocols, sites, or over time). For this reason, **qMRI has the potential to make a great clinical impact on diagnostics by providing quantitative biomarkers for clear numerical differentiation of NDAD disease states**, complementing or replacing invasive biopsies. qMRI can enable earlier detection of disease, and increases the quality of information available to artificial intelligence algorithms for predicting prognosis or therapeutic response (Tsehaie et al., 2017; Keenan et al., 2019).

1.4 The unmet need for technical development

Unfortunately, clinical adoption and dissemination of qMRI is lagging behind, which is mainly due to its long scan time requirements. As will be more extensively discussed in Chapter 3, the root cause of these long scan times for qMRI is the acquisition of multiple contrast weighted images that are required to extract quantitative parametric maps with adequate accuracy, precision, and spatial resolution. Alas, long scan times are disadvantageous for medical MRI:

- From a diagnostic perspective, long scan times increase the likelihood of patient motion during the MRI scan which typically leads to loss of spatial resolution and/or artefacts in the resulting MRI image, and which in turn are detrimental for accurate diagnosis.
- From an economical perspective, long scan times reduce the patient throughput, i.e. the number of patients that can be scanned in a hospital or point-of-care facility per day. Ideally, scan time is short, such that waiting times for a scan are almost nonexistent and large-scale regular screening of patients becomes feasible.
- *From a patients perspective*, long scan times lead to problems with patient comfort and compliance. Hence, it is important that scan time is limited for routine clinical use.

Given the **need for rapid imaging techniques**, the development of accelerated MRI methods has sparked in recent years. For example, methods have been proposed that enable reconstruction from highly under-sampled images and hence speed up image acquisition, such as model-based reconstruction (Maier et al., 2019), low-rank approaches (Zhang et al., 2015), or the imposition of sparsity constraints (Zhao et al., 2012). However, these methods generally come at the cost of either a lower precision or a lower spatial resolution of the reconstructed (parametric) MR images. Therefore, MRI research has been focused on the development of innovative technologies to optimize this trade-off between SNR, spatial resolution, and scan time. The contributions and research in this thesis are specifically focused on the use of so-called model-based super-resolution reconstruction (SRR) techniques, that directly estimate high-resolution (HR) images or HR quantitative parameter maps from sets of low-resolution (LR) contrast-weighted MR images (Poot et al., 2010; Plenge et al., 2012; Poot et al., 2013; Van Steenkiste et al., 2016, 2017). As will be shown throughout this thesis, the use of SRR techniques in the context of MRI offers great potential to balance the existing trade-off between SNR, spatial resolution, and scan time. More elaborate background information on the theory behind SRR, which has also proven its strengths in many other imaging applications (Park et al., 2003), is provided in Chapter 4.

Furthermore, as will be discussed in Chapter 3, qMRI extends upon conventional anatomical MRI by introducing biophysical, often nonlinear, signal models in addition to conventional image reconstruction models. These signal models typically describe a voxel-wise relation between the HR image and several biophysical tissue parameters of interest (e.g., T1 and T2 relaxation times). Often these signal models also depend on a set of acquisition parameters (e.g., inversion times, echo times, ...) to be chosen in optimal accordance with the MRI experiment at hand. As a result, qMRI is no longer a mere imaging problem that tries to estimate voxel intensities of the 3D HR image, but it develops into a **complicated multi-parameter estimation problem** in which it must be examined how underlying tissue parameters can be accurately and precisely estimated from a series of MR images. As such, qMRI also requires a well-structured statistical estimation framework for large-scale parameter estimation based on MR input data acquired within a preferably short scan time.

As mentioned, MRI is also subject to external factors that can complicate the imaging process. These factors include e.g. the presence of noise, magnetic field inhomogeneity issues inherent to some specific MR sequence types, or RF nonuniformity of particular head coils. In addition, **unwanted patient motion** can also have a major impact on the final outcome of an MRI scan, as it results in unwanted ghosting and blurring artefacts (Zaitsev et al., 2015). As such, there exists a need for motion compensation techniques that can correct motion artifacts either prospectively, i.e. by obtaining real-time tracking data of the position and orientation of the subject during a scan, or retrospectively, i.e. by modification of the MR image data during reconstruction of a scan (Godenschweger et al., 2016). As will be explained in Chapter 3, the use of motion correction strategies in combination with qMRI parameter estimation problems requires careful consideration. Conventional qMRI methods usually correct for motion by performing image registration routines as a pre-processing step, prior to the estimation of the HR parameter maps (Studler et al., 2010; Bron et al., 2013; Guyader et al., 2015; Van Steenkiste et al., 2016, 2017), where the latter step is often preceded by an intermediate step of HR image reconstruction (Scherrer et al., 2012; Poot et al., 2013). A downside to such multi-step approaches is the lack of a feedback mechanism that connects the motion compensation routine with the final estimation of the HR parameter maps. As a result, registration errors may propagate into the parameter estimation step, possibly leading to inaccurate (i.e., biased) estimates which do not reflect the underlying tissue (Nachmani et al., 2019). To avoid error propagation, image registration needs to be integrated in a joint motion/qMRI parameter estimation framework. That is, by providing an explicit model for the patient's motion during scanning, the corresponding motion parameters of that model can be estimated simultaneously with the qMRI parameters.

Now, it becomes even more challenging when you add up all these components - SRR, qMRI, and joint retrospective motion correction - in one widely-applicable MRI framework. This combination is not trivial, and while building a combined framework the level of complexity of the framework will gradually increase. Therefore, a crucial aim of this thesis is to study to what extent a combined framework for motion-robust quantitative MRI using super-resolution reconstruction can be developed, without compromising any of the individual components. As such, the innovations in this work can help with bringing accelerated motion-robust quantitative MRI within reach of patients to overcome some of the existing technological barriers in gMRI.

1.5 Thesis contributions

The primary objective of this dissertation is:

To provide a model-based super-resolution reconstruction framework with joint estimation of inter-scan patient motion, that can be applied in a variety of clinical applications in which quantitative MRI and motion correction are both desirable. Specifically, this framework is intended for medical imaging applications where the use of thick-slice contrast-weighted MRI data is current clinical practice due to scan time limitations or other inherent requirements of the imaging modality.

Several challenges needed to be overcome in this context, which mainly revolved around the **three key technical components** mentioned in the previous section:

- 1. Robust and accurate patient motion correction.
- 2. Accurate and precise estimation of quantitative MRI parameters.
- 3. Robust model-based super-resolution reconstruction to balance the trade-off between resolution, SNR, and scan time.

As such, the framework will be capable of estimating high-fidelity, high-resolution 3D quantitative parameter maps from a set of contrast-weighted low-resolution MRI scans.

The highlights of the main contributions presented in this thesis are summarized hereafter. Note that the contributions in this thesis were arranged in the order in which their developments fit the timeline of this PhD. This means that the gradual expansion of the SRR framework, by adding more accurate physical models or introducing more efficient implementations, is visible throughout the chapters. In particular, as a proof-of-concept study the estimation of motion and quantitative MRI parameters was first applied to small-image phantoms using T1 mapping as an MR relaxometry model of choice (Chapter 5), followed by more advanced reconstructions using both real-size whole brain phantoms and *in vivo* brain data while also exploiting prior knowledge about the tissue and noise statistics (Chapter 6). Finally, the general applicability of the super-resolution framework for a more advanced MRI modality such as Arterial Spin Labeling (ASL) MRI was demonstrated (Chapter 7).

Contribution 1: Joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework: a simulation study

A super-resolution framework for joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images is proposed, which is tested by means of an extensive Monte Carlo simulation study.

Highlights of this contribution:

- Proof-of-concept study in which extensive Monte Carlo simulations are performed on T1-weighted small-image phantoms to demonstrate the potential of augmenting model-based SRR for quantitative T1 mapping with joint inter-image motion estimation.
- A joint maximum likelihood estimator is used to optimally exploit knowledge about the data distribution of the low-resolution images. The measured low-resolution images are assumed to be Gaussian distributed.
- The proposed SRR method is benchmarked against three alternative approaches, including SRR without motion estimation, SRR using (multi-level) mutual information based registration as a preprocessing step, and a previously reported SRR T1 mapping approach using a loop-wise pre-registration scheme.
- Performance is analysed for different SNR values of the input data sets, and measured in terms of relative bias, relative standard deviation, relative root-mean-square error, and motion component root-(mean)-mean-square error. Histograms of the voxel data distribution of the reconstructed T1 parameter maps are also visually compared.
- The extension of super-resolution reconstruction with simultaneous motion estimation yields more accurate T1 maps compared to a previously reported SRR-based T1 mapping approach in which motion registration is applied as a preprocessing step.

Contribution 2: Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping

A general model-based super-resolution reconstruction framework with joint motion estimation for improved quantitative MRI parameter mapping is proposed, that was specifically used to study T1 and T2 brain mapping on synthetic whole brain phantoms and *in vivo* brain data.

Highlights of this contribution:

- A general model-based super-resolution reconstruction framework with joint estimation of motion parameters and isotropic, high-resolution 3D quantitative MRI parameter maps from motion-corrupted, low-resolution 2D multi-slice MRI scans.
- A Bayesian estimation approach to maximally exploit prior knowledge of the tissue and noise statistics, introducing an upwind Total Variation prior on the high resolution parameter maps to be estimated. The measured low-resolution images are assumed to be Rician distributed.
- The framework is validated and benchmarked using whole brain Monte Carlo simulations with realistic spatially varying noise and magnitude MR specific data distributions (Rician distribution).
- Demonstration of the clinical potential of the presented reconstruction framework for two MR relaxometry quantitative mapping protocols to directly map high-resolution 3D T1 or T2 relaxation maps from whole brain *in vivo* MRI data with low through-plane resolution, i.e. with thick slices.
- Superior accuracy compared to quantitative MRI with motion pre-compensation is demonstrated.

Contribution 3: A super-resolution reconstruction framework for quantitative brain perfusion mapping using pseudo-continuous Arterial Spin Labeling

This contribution introduces a model-based super-resolution reconstruction framework for single post-labeling delay (single-PLD) pseudo-Continuous Arterial Spin Labeling (pCASL) MRI, building on a joint Bayesian estimation framework that aims to estimate motion-corrected 3D isotropic HR quantitative cerebral blood flow (CBF) maps from a set of 2D multi-slice pCASL control-label image pairs acquired with low through-plane resolution and rotated slice-encoding direction. By improving upon existing disadvantages of 2D multi-slice readout for pCASL, the proposed framework provides a promising alternative to the recommended segmented 3D readout schemes, which to date remain sensitive to inter-shot motion and through-plane blurring due to T2 decay along the long echo trains.

Highlights of this contribution:

- Isotropic, high-resolution 3D quantitative cerebral blood flow (CBF) mapping from 2D multi-slice single-post-labeling-delay (single-PLD) pseudo-continuous ASL data.
- Joint estimation of CBF and inter-image motion between ASL control-label images.

- A Bayesian estimation approach to exploit prior knowledge of the tissue and noise statistics, in which the estimation of CBF parameter maps is formulated as an efficient linear estimation problem.
- Validation and benchmarking using whole brain Monte Carlo simulations and *in vivo* brain data.
- More optimal background suppression and post-labeling delay compared to conventional 2D multi-slice readout.

Software contributions

In order to support the research presented in this thesis and to facilitate collaboration with fellow researchers within or outside academia, the developed concepts for motion-robust model-based super-resolution reconstruction for multi-slice (q)MRI were bundled in two software packages. A first package (https://github.com/qbeirinckx/Super-Resolution-Reconstruction) contains all the MATLAB code that was used for the simulation experiments and real data reconstructions in this thesis. Particularly, the software package includes individual coding modules for the forward and transpose operators (warping, blurring, resampling) in the SRR forward model introduced in contribution chapter 6, as well as implementations for three-dimensional regularization/prior terms on image or tissue parameter maps (both Total Variation and Laplacian regularization are currently implemented, cf. contribution chapters 5-7). Further, different log-likelihood functions for magnitude MRI data (assuming e.g. Gaussian or Rician distributed data) can be selected during the estimation process to create Maximum Likelihood or Bayesian estimators. Finally, the user can select a preferred signal model of choice (T1, T2, diffusion, arterial spin labeling, twoor many-parameter models, etc.), making the code versatile and multi-usable for gMRI reconstruction. A second package, called **STORM**, acronym for **S**uper-resolution **To**mographic Reconstruction for MRI (https://github.com/qbeirinckx/STORM), provides a general Python toolkit for anatomical super-resolution reconstruction with joint motion estimation. This package provides an efficient warping operator implementation and has already been used to improve multiple sclerosis lesion segmentation from retrospective data (Giraldo et al., 2023), while it also enables the use of SRR in follow-up Artificial Intelligence (AI) and learning-based method development where Python is the preferred software language.

1.6 Thesis organisation

This thesis is structured in 8 chapters. In the **current chapter** (Prologue), the motivation and existing technical challenges behind the work in this thesis were briefly introduced. The need for discovery of reliable and accessible biomarkers for detection of neurodegenerative diseases in an aging era was pointed out, and how qMRI as a medical imaging modality can be a crucial protagonist by providing quantitative biomarkers for clear numerical differentiation of disease states. The different technical needs that come with bringing qMRI to the clinic were highlighted. In particular, the combination of innovative technologies that break the traditional MRI trade-off between resolution, precision, and acquisition time, while also providing robust motion/qMRI parameter estimation.

Understanding the basic principles of MRI is a prerequisite to studying the contributions in this thesis. Starting from an overview of the historical advancements in MRI, the **second chapter**

(Magnetic Resonance Imaging: The basics) therefore briefly explains some fundamental principles and concepts about MRI. In particular, the phenomenon of nuclear magnetic resonance is considered, which is grounded in quantum mechanical considerations, but which benefits from a classical interpretation when the interaction of a large number of hydrogen atoms in an external magnetic field is considered. In addition, fundamental concepts such as excitation and relaxation of magnetic spins are introduced. Since the relaxation process is typically characterized by parameters of great importance for quantitative neuroimaging research, e.g. T1 and T2 relaxation times, extra attention is paid to these phenomena. Next, a short overview of signal generation and detection in MRI is given, highlighting aspects that are of particular interest to the MRI modalities and methods used in this thesis. In particular, the use of 2D multi-slice and 3D acquisition strategies is described, pinpointing some of the advantages and disadvantages of both strategies. Finally, the chapter provides a detailed list of the specific MRI pulse sequences used for the acquisition of the *in vivo* whole brain data in the research contributions.

The **third chapter** (The Advent of Quantitative MRI) outlines the importance of quantitative MRI in response to shortcomings of and as an addition to conventional MRI. MR relaxometry and arterial spin labeling MRI are presented as two important clinical applications of qMRI, for which this thesis provides new methods to answer existing technical needs. Finally, in the context of qMRI, the basic principles of parameter estimation are introduced.

In the **fourth chapter** (Super-resolution reconstruction as prime protagonist for accelerated (q)MRI), the role of super-resolution reconstruction is elaborated. Starting with a general discussion on the challenges of resolution in MRI and the consequent demand for resolutionenhancing methods, the general concept of SRR is introduced. In addition to the fundamental principle of this technique, the key components are highlighted necessary for SRR application in MRI. Particularly, the choice of an appropriate acquisition strategy, the selection of an accurate imaging model that realistically models both the acquisition and noise, and finally, the model-based reconstruction process addressing technical aspects such as regularization and hyperparameter selection. At the end of this chapter, the need for robust motion compensation is also briefly touched upon, which is indispensable for ensuring the spatial alignment of the SRR image set.

Next, the remaining chapters discuss the main contributions of this thesis, where the **fifth chapter** (Joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework : a simulation study) describes an extensive Monte Carlo simulation study on small-image 'checkerboard'-like phantoms in which for the first time model-based SRR for qMRI is augmented with joint motion estimation, using T1 mapping as an MR relaxometry model of choice. A joint Maximum Likelihood estimation framework is used in which motion and relaxometry parameters are optimized alternately. Also in this contribution, downsampling and blurring with a rectangular slice profile are combined in one and the same forward operator to ease the computational burden. The proposed SRR framework with joint motion estimation is extensively benchmarked against a framework without motion estimation and a previously reported SRR T1 mapping approach with a motion pre-registration strategy. This contribution chapter provides a first proof-of-concept study for model-based SRR with joint motion estimation for qMRI applications, and lays the foundations for the further development of the large-scale and computationally more demanding SRR framework for application to whole brain in vivo images, as described in more detail in the subsequent chapter.

The **sixth chapter** (Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping) expands the simulation study and proofof-concept demonstration further towards a widely applicable framework for model-based super-resolution reconstruction with joint estimation of motion parameters and isotropic, high-resolution 3D quantitative MRI parameter maps from motion-corrupted, low-resolution 2D multi-slice MRI scans. This chapter describes the *magnum opus of this dissertation* in which all important technological developments are combined in one polyvalent and robust estimation framework. Among other things, a Bayesian estimation framework is used to maximally exploit prior knowledge of the tissue and noise statistics, in which special attention is given to realistic modeling of noise and data distributions by use of a Rician probability distribution in the likelihood function of the magnitude MR data. The framework's potential is demonstrated in both simulations and real data experiments, using T1 and T2 mapping as carrying examples. Attention is also given to the experimental design of the acquisition protocols for relaxometry mapping, in which contrast weighting and (geometric) distribution of the individual images in the quantitative image series is subject to restrictions inherent to the MR acquisition sequence of interest. In particular, it is investigated how a multi-echo spin-echo (MESE) sequence can be used for acquisition of a 2D multi-slice super-resolution image series for T2 mapping, while also discussing limitations of such a sequence. It should be highlighted that the sections and formulae in this chapter have been written in such a way that it should be straightforward to substitute other quantitative signal and/or noise models into the framework. In fact, this modularity in signal and noise model was illustrated in the following chapter in which a perfusion model with a Gaussian noise model is combined under these considerations. The appendices of this second contribution chapter bundle the different mathematical components, a.k.a 'building blocks', and tools that were developed in the reconstruction of the motion-robust super-resolution reconstruction framework for qMRI. An important part is dedicated to the implementation of the linear operators that constitute the super-resolution forward model, and the qMRI signal models that describe the relation between the signal intensity and the underlying biophysical tissue parameters in a voxel. In addition, some practical considerations are discussed for solving a large-scale parameter estimation problem with joint estimation of qMRI and motion parameters. Finally, some examples of use cases for anatomical and guantitative SRR applied in musculoskeletal MRI are presented, which were reconstructed with the developed framework. As opposed to brain MRI, musculoskeletal (MSK) MRI focuses specifically on joint structures, including wrists, ankles, knees, etc. It is demonstrated that SRR can also play an important role to improve existing 3D resolutions of MSK MRI without significantly increasing the scan time of clinical protocols.

In the **seventh chapter** (A super-resolution reconstruction framework for quantitative brain perfusion mapping using pseudo-continuous Arterial Spin Labeling), a super-resolution reconstruction framework was proposed that estimates 3D isotropic high-resolution quantitative cerebral blood flow (CBF) maps from a series of single-post-labeling-delay (single-PLD) pseudo-continuous Arterial Spin Labeling (ASL) control-label image pairs, each acquired with low through-plane resolution and rotated slice-encoding direction in a 2D multi-slice readout scheme. Building upon the SRR framework introduced in Chapter 6, motion between control and label images was jointly estimated in a Bayesian estimation framework, enabling accurate and precise CBF quantification without propagation of pre-registration errors, while optimally exploiting prior knowledge of tissue properties and noise statistics. The rotation of the slice-encoding direction for each control-label image pair as well as the lower through-plane

resolution ensures a more uniform distribution of the PLD throughout the brain and increases the effectiveness of background suppression. Combined, this significantly improved the SNR compared to conventional 2D multi-slice readout with direct high through-plane resolution, where CBF quantification has traditionally been hampered by detrimental perfusion SNR slice dependence in sequentially acquired slices. The proposed method was validated both qualitatively and quantitatively in synthetic whole brain simulations and on *in vivo* human brain data. It has been demonstrated that the framework provided superior CBF estimation in terms of root-mean-square error compared to a state-of-the-art approach using a conventional 2D multi-slice readout strategy with ascending slice order and isotropic resolution in the same scan time, even when additional hardware acceleration techniques like multiband are applied in the latter.

Finally, the **eighth chapter** (Conclusions and Future Perspectives) concludes the thesis and outlines potential future research directions.

At the end of the thesis manuscript, a list of abbreviations and a concise academic CV at the time of writing are also provided.

References

- Bron, E. E., van Tiel, J., Smit, H., Poot, D. H. J., Niessen, W. J., Krestin, G. P., Weinans, H., Oei, E. H. G., Kotek, G., & Klein, S. (2013). Image registration improves human knee cartilage T1 mapping with delayed gadolinium-enhanced MRI of cartilage (dGEMRIC). *European Radiology*, 23(1), 246–252.
- Deoni, S. C. L. (2010). Quantitative relaxometry of the brain. *Top. Magn. Reson. Imaging*, 21(2), 101–113.
- Deuschl, G., et al. (2020). The burden of neurological diseases in europe: an analysis for the global burden of disease study 2017. *The Lancet Public Health*, *5*(10), e551–e567.
- Feigin, V. L., et al. (2020). The global burden of neurological disorders: translating evidence into policy. The Lancet Neurology, 19(3), 255–265.
- Giraldo, D. L., Beirinckx, Q., den Dekker, A. J., Jeurissen, B., & Sijbers, J. (2023). Super-Resolution Reconstruction of Multi-Slice T2-W FLAIR MRI Improves Multiple Sclerosis Lesion Segmentation. *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2023, 1–4.
- Godenschweger, F., et al. (2016). Motion correction in MRI of the brain. *Phys. Med. Biol.*, *61*(5), R32–R56.
- Guyader, J.-M., Bernardin, L., Douglas, N. H. M., Poot, D. H. J., Niessen, W. J., & Klein, S. (2015). Influence of image registration on apparent diffusion coefficient images computed from free-breathing diffusion mr images of the abdomen. *J. Magn. Reson. Imaging*, *42*(2), 315–330.

Hansson, O. (2021). Biomarkers for neurodegenerative diseases. Nature Medicine, 27(6), 954-963.

- Hockings, P., et al. (2020). MRI Biomarkers. In *Quantitative Magnetic Resonance Imaging*, vol. 1 of *Advances in Magnetic Resonance Technology and Applications*, (pp. liii–lxxxvi). Academic Press.
- Keenan, K. E., et al. (2019). Recommendations towards standards for quantitative MRI (qMRI) and outstanding needs. *Journal of Magnetic Resonance Imaging*, 49(7), e26–e39.
- Maier, O., et al. (2019). Rapid T1 quantification from high resolution 3D data with model-based reconstruction. *Magn. Reson. Med.*, *81*(3), 2072–2089.
- Mathers, C. D., & Loncar, D. (2006). Projections of global mortality and burden of disease from 2002 to 2030. *PLOS Medicine*, *3*(11), 1–20.
- Nachmani, A., Schurr, R., Joskowicz, L., & Mezer, A. A. (2019). The effect of motion correction interpolation on quantitative T1 mapping with MRI. *Med. Image Anal.*, *52*, 119–127.
- Park, S. C., Park, M. K., & Kang, M. G. (2003). Super-resolution image reconstruction: a technical overview. IEEE Signal Process. Mag., 20(3), 21–36.
- Plenge, E., et al. (2012). Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magn. Reson. Med.*, *68*(6), 1983–1993.
- Poot, D. H. J., Jeurissen, B., Bastiaensen, Y., Veraart, J., Van Hecke, W., Parizel, P. M., & Sijbers, J. (2013). Super-resolution for multislice diffusion tensor imaging. *Magn. Reson. Med.*, 69(1), 103–113.
- Poot, D. H. J., Van Meir, V., & Sijbers, J. (2010). General and efficient super-resolution method for multi-slice MRI. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, vol. 6361, (pp. 615–622).
- Scherrer, B., Gholipour, A., & Warfield, S. K. (2012). Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Med. Image Anal.*, 16(7), 1465–1476.
- Studler, U., White, L. M., Andreisek, G., Luu, S., Cheng, H.-L. M., & Sussman, M. S. (2010). Impact of motion on T1 mapping acquired with inversion recovery fast spin echo and rapid spoiled gradient recalled-echo pulse sequences for delayed gadolinium-enhanced MRI of cartilage (dGEMRIC) in volunteers. J. Magn. Reson. Imaging, 32(2), 394–398.
- Teunissen, C. E., et al. (2022). Blood-based biomarkers for Alzheimer's disease: towards clinical implementation. Lancet Neurology, 21(1), 66–77.
- Tsehaie, J., et al. (2017). Value of quantitative MRI parameters in predicting and evaluating clinical outcome in conservatively treated patients with chronic midportion Achilles tendinopathy: A prospective study. *Journal of Science and Medicine in Sport*, 20(7), 633–637.
- United Nations, Department of Economic and Social Affairs, Population Division (2022). World Population Prospects 2022: Ten Key Messages. The United Nations.
- Van Steenkiste, G., et al. (2016). Super-resolution reconstruction of diffusion parameters from diffusionweighted images with different slice orientations. *Magn. Reson. Med.*, 75(1), 181–195.
- Van Steenkiste, G., et al. (2017). Super-resolution T1 estimation: Quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. *Magn. Reson. Med.*, 77(5), 1818–1830.
- Winkler, A. S. (2020). The growing burden of neurological disorders in low-income and middle-income countries: priorities for policy making. *The Lancet Neurology*, *19*(3), 200–202.
- World Health Organization (2022). *Optimizing brain health across the life course: WHO position paper*. Geneva: World Health Organization.
- Zaitsev, M., Maclaren, J., & Herbst, M. (2015). Motion artifacts in MRI: A complex problem with many partial solutions. J. Magn. Reson. Imaging, 42(4), 887–901.
- Zhang, T., Pauly, J. M., & Levesque, I. R. (2015). Accelerating parameter mapping with a locally low rank constraint. *Magn. Reson. Med.*, 73(2), 655–661.
- Zhao, B., Haldar, J. P., Christodoulou, A. G., & Liang, Z. (2012). Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints. *IEEE Trans. Med. Imag.*, *31*(9), 1809–1820.

Part II

Background



Magnetic Resonance Imaging: The basics

Contents

2.1	Introduction	17
2.2	A brief history of MRI	18
2.3	NMR signal generation and detection	19
2.4	MR image formation and readout strategies	26
2.5	The MRI trade-off	33
References		

2.1 Introduction

Since this thesis contains several contributions on MRI, it is worthwhile to clearly state the general idea behind this versatile imaging technique. Starting from an overview of the historical advancements in MRI, this chapter briefly explains some fundamental principles and concepts about MRI. In particular, the phenomenon of nuclear magnetic resonance (NMR) is considered, which is grounded in guantum mechanical considerations, but which, as will be emphasized, also benefits from a classical physical interpretation when the interaction of a large number of hydrogen atoms in an external magnetic field is considered. In addition, fundamental concepts such as excitation and relaxation of magnetic spins are introduced. Since the relaxation process is typically characterized by parameters of great importance in a clinical context, especially T1 and T2 relaxation times, extra attention is paid to these phenomena. In addition to generating and detecting the MR signal, this section also explains how an image is created through a well-thought-out spatial coding of the MR signal. Furthermore, the use of 2D multi-slice and 3D acquisition strategies is described, also pinpointing some of the advantages and disadvantages of both strategies. Finally, a detailed list is provided of the particular MRI pulse sequences used in this thesis to obtain the in vivo whole brain data sets.

2.2 A brief history of MRI

The historic journey of MRI began in the 1930s when the American physicist Isodor Isaac Rabi first described the phenomena of *nuclear magnetic resonance* (NMR) (Rabi et al., 1938). He showed that one can manipulate and identify atomic nuclei, which behave like spinning tops whose orientation axes are aligned with an externally magnetic field, by exposing them to radio-waves. Recognition for this pioneering work followed in 1944, when he was awarded the Nobel Prize in Physics.

Apart from Rabi's work, several key researchers made foundational discoveries. The Serbian-Croatian scientist Nikola Tesla first observed the principles of magnetic resonance as early as 1896 (Roguin, 2004), laying the groundwork for future explorations into electromagnetic fields and resonance phenomena. In the early 1940s, Soviet physicist Yevgeny Zavoisky working at Kazan State University made significant contributions by detecting electron paramagnetic resonance (Zavoisky, 1945), a technique closely related to NMR. Dutch physicist Cornelis Gorter, whose experiments on paramagnetic relaxation in solid paramagnetic salts inspired Zavoisky, also made notable contributions (Gorter, 1936). Additionally, Dutch scientist Pieter Zeeman played a crucial role in understanding the magnetic properties of materials. Zeeman's discovery of the spectral line splitting (see section 2.3.1), for which he received the Nobel Prize in 1902, was pivotal in developing spectroscopic techniques (Zeeman, 1897).

In the aftermath of the Second World War, when developments in radar and electronic technologies were extensively explored, the groups of Felix Bloch and Edward Mills Purcell independently demonstrated that any solid or liquid can be placed in a magnetic field to identify its specific atoms, without affecting it in any perceptible way using the NMR phenomena (Bloch, 1946; Purcell et al., 1946). They were jointly awarded the Nobel Prize in Physics in 1952. During the next decades, NMR grew into a widely used application for structural analysis of materials. However, it was not until the 1970s that NMR signals could be used to generate two-dimensional (2D) images. Paul Lauterbur expanded upon the work of Herman Carr to develop spatial information encoding principles (Carr & Purcell, 1954; Lauterbur, 1973). Peter Mansfield developed a method, currently known as 'echo planar imaging' (EPI) to acquire such 2D images in only a few seconds (Mansfield, 1977). Both scientists received the Nobel Prize in Physiology and Medicine in 2003 for their seminal contributions, which led to the applications of magnetic resonance in medical imaging. However, the Nobel Prize ensued some controversy (Dreizen, 2004). Why was Raymond Damadian not honored by the Nobel Prize for his contribution to MRI in medicine? The Armenian-American medical doctor reported that differences among normal tissues and between normal and cancer tissues can be distinguished in vivo by NMR (Damadian, 1971). Moreover, he was the first to achieve human whole-body MR images (granted a patent in 1974, (Damadian, 1974)). Finally, another critical contribution that forever changed the way MRI was done, was made by Richard R. Ernst. Inspired by the Belgian scientist Jean Jeener (Jeener, 1971; Jeener et al., 1979), Ernst and coworkers were the first to use the Fourier transform to reconstruct two-dimensional (2D) NMR images, using switched magnetic field gradients in the time domain for spatially encoding (Ernst & Anderson, 1966; Kumar et al., 1975). Ernst was awarded a Nobel Prize in Chemistry in 1991 for his contributions.

Since those early discoveries, modern-day MRI technology has undergone many changes. For one thing, early marketers decided to drop the word 'nuclear' from 'nuclear magnetic resonance,' reasoning that this would allay people's fears about radiation. What initially

started with a single MRI scanner and a magnetic field strength of 0.1 Tesla (T) (Damadian, 1974) has grown into an important biomedical technology with more than 65,000 MRI scanner units used worldwide, which together perform more than 150 million exams per year (OECD, 2021), operating at different magnetic field strengths; for clinical routine (\leq 3T), and/or research purposes (\leq 10.5T) (Vaughan et al., 2006; Moser et al., 2012).

However, as already pointed out in the Prologue, MRI also remains an inaccessible imaging technology, cost-prohibitive given the need for specialized hardware, installation in dedicated RF-shielded hospital spaces, and highly-trained operators and radiologists. Access is particularly limited in rural areas, and in low- and middle-income countries - thereby contributing to the current challenges of health disparities. In a world where we are today facing significant challenges of a fast-growing aging society, with the prevalence of age-related diseases and disorders rising year after year, recent MRI research is therefore also focusing on improving the accessibility of MRI and making it more cost-friendly to ease the burden on global healthcare costs. As an example, the development of low-cost ultra-low-field MRI scanners ($\sim mT$ range) that can be used for rapid and regular screening of patients in the general population at point-of-care facilities (Liu et al., 2021), has recently regained interest from the MRI research community.

2.3 NMR signal generation and detection

2.3.1 Physical principles, nuclear spin and magnetic moment

As mentioned in the previous section, the NMR phenomenon was first described by F. Bloch and E. Purcell. In short, the NMR phenomenon describes the absorption and subsequent re-emission of electromagnetic (EM) radiation by a system of nuclei with an odd number of protons or neutrons in a static magnetic field, when perturbed by a second oscillating magnetic field with a specifically selected frequency.

If an atomic nucleus possesses an odd number of protons or neutrons, it has an intrinsic angular momentum $J = \hbar I$, with I the intrinsic spin, a dimensionless vector, and \hbar the reduced constant of Planck $(1.05 \cdot 10^{-34} J \cdot s)$. Since a nucleus is charged, the intrinsic angular momentum J is coupled with a *magnetic dipole moment* μ :

$$\boldsymbol{\mu} = \boldsymbol{\gamma} \boldsymbol{\hbar} \boldsymbol{I} \tag{2.1}$$

where γ denotes the gyro-magnetic ratio of the nucleus. In quantum mechanics, the spin angular momentum operator $\hat{J} = \hbar \hat{I}$ has eigenvalues $\hbar \sqrt{I(I+1)}$ with *I* the spin quantum number. This spin quantum number is an intrinsic property of the nucleus, which is an integer or a half integer. In this thesis, and in the vast majority of clinical MRI exams, the considered nucleus is the *hydrogen proton* (¹H). This proton has a high natural abundance in the human body in the form of water molecules. The spin quantum number of ¹H is $I = \frac{1}{2}$, while its gyro-magnetic ratio is $\gamma = 2.675 \cdot 10^8 rad/s/T$. When a proton is placed in an external magnetic field B_0 directed along the *z*-axis, the component of \hat{I} parallel with the magnetic field, \hat{I}_z , has eigenvalues I_z that can take 2I + 1 values: $-I, -I + 1, \ldots, I$. These eigenvalues I_z are the possible outcomes of a measurement of the angular momentum along the *z*-axis. In the case of ¹H, there are two possible values: *spin up*, $I_z = +1/2$, or *spin down*, $I_z = -1/2$. Considering the linear relation between the intrinsic angular momentum I and the magnetic moment μ , the proton magnetic moment has only two possible states: $+\frac{1}{2}\gamma\hbar$

or $-\frac{1}{2}\gamma\hbar$. Those two discrete magnetic moments of the proton possess opposite potential energy in an external magnetic field B_0 :

$$E = -\mu \cdot B_0 = -\mu_z B_0 = \begin{cases} E_{\downarrow} = +\frac{1}{2}\gamma\hbar B_0 & \text{spin down} \quad (I_z = -\frac{1}{2}), \\ E_{\uparrow} = -\frac{1}{2}\gamma\hbar B_0 & \text{spin up} \quad (I_z = +\frac{1}{2}), \end{cases}$$
(2.2)

where '·' denotes the scalar product, and $|B_0| = B_0$. This discrete difference between both energy levels is referred to as *nuclear Zeeman-splitting* (Fig. 2.1). The lower energy level, spin up, corresponds to the *z*-component of the magnetic moment oriented parallel with B_0 , while the higher energy level, spin down, is compatible with the magnetic moment oriented anti-parallel with B_0 . The energy difference between both states is given by:

$$\Delta E = E_{\downarrow} - E_{\uparrow} = \gamma \hbar B_0 = h\nu_L = \hbar \omega_L. \tag{2.3}$$

The above equation is referred to as the **resonance condition**. It states that energy transitions from one energy level to another are possible by absorption or emission of a photon with energy $\Delta E = \gamma \hbar B_0$. Such photons are characterized by an angular frequency $\omega_L = \gamma B_0$, commonly referred to as the **Larmor frequency** (Larmor, 1897). This Larmor frequency corresponds to the precession frequency of hydrogen spins in a magnetic field, meaning it is the rate at which the magnetic moments of hydrogen nuclei precess around the axis of the external magnetic field B_0 . In a clinical setting, magnetic field strength ranges typically from 1.5T to 3T, which corresponds to a Larmor resonance frequency in the radio frequency (RF) part of the EM spectrum. Therefore, RF shielding (cf. Faraday cage principle) of the MRI scanner is mandatory to prevent external EM radiation from contaminating/distorting the MR signal, and to prevent EM radiation generated by the MR scanner from causing interference with medical devices nearby.



Figure 2.1: Nuclear Zeeman effect for a ¹H proton $(I = \frac{1}{2})$ in an external magnetic field $B_0 = B_0 e_z$, with e_z the unit vector along the z-axis. In the presence of a magnetic field $(B_0 \neq 0)$ a splitting occurs in two separate energy levels, corresponding to the two nuclear spin states of a $I = \frac{1}{2}$ particle. The energy split is proportional to the magnitude of the applied magnetic field B_0 .

2.3.2 Macroscopic effect of the static magnetic field

In practice, the matter being investigated in the MRI scanner consists of a large number of nuclei, in close proximity of one another. While spin of a single nucleus is a quantum effect, the presence of an *ensemble* of nuclei allows to study magnetic resonance as a classical phenomenon (Hanson, 2008). This is justified by the *correspondence principle*, which states that the behaviour of systems described by quantum mechanics follows classical physics in the limit of large quantum numbers (Bohr, 1976). Since $E_{\uparrow} < E_{\downarrow}$, the equilibrium population of spins in the spin up state (higher stability), N_{\uparrow} , exceeds the population of spins in the spin up state of spins and the spin statistics, the ratio between both population numbers is given by

$$P = \frac{N_{\uparrow}}{N_{\downarrow}} = \exp\left(\frac{\Delta E}{k_B T}\right) > 1, \qquad (2.4)$$

with k_B the Boltzmann constant and T the absolute temperature. Since sending EM waves at the resonance frequency will cause both stimulated absorption and stimulated emission between the energy states, only a net absorption/emission can be detected. At physiological temperature (310.15K) and a typical magnetic field strength of 3T, Eq. (2.4) results in P = 1.0000197. In other words, only about $2 \cdot 10^{-5}$ proton spins will contribute to the NMR signal. Hence, due to this low sensitivity, a large concentration of protons is necessary to create an NMR signal. Fortunately, for a typical voxel volume in MRI, there are about 10^{21} protons (Bushberg et al., 2012), so there are $2 \cdot 10^{-5} \cdot 10^{21}$, or approximately $2 \cdot 10^{16}$ more protons in the low-energy state, producing an observable net absorption of RF energy on a macroscopic scale. In addition, it can be shown, by first-order approximation of Eq. (2.4), that the difference in occupation of both spin states is given by

$$\frac{N_{\uparrow} - N_{\downarrow}}{N} \approx \frac{\gamma \hbar B_0}{2k_B T} = 9.84 \cdot 10^{-6} \text{ at } 3\text{T and } 310.15\text{K}, \tag{2.5}$$

with $N = N_{\uparrow} + N_{\downarrow}$ equal to the total number of protons. Again, although the difference in occupation is small, it is sufficient to generate an observable macroscopic magnetization vector M. Being able to treat the behavior of all spins in the ensemble in terms of a net magnetization vector M allows a classical description of NMR.

The resulting **bulk magnetization**, which is the vector sum of all the microscopic magnetic moments in the object, i.e. $M = \sum_{n=1}^{N} \mu_n$, can be further decomposed in an *x*, *y*, and *z*-component,

$$\boldsymbol{M} = M_{x}\boldsymbol{e}_{x} + M_{y}\boldsymbol{e}_{y} + M_{z}\boldsymbol{e}_{z}$$
$$= \left(\sum_{n=1}^{N} \mu_{x,n}\right)\boldsymbol{e}_{x} + \left(\sum_{n=1}^{N} \mu_{y,n}\right)\boldsymbol{e}_{y} + \left(\sum_{n=1}^{N} \mu_{z,n}\right)\boldsymbol{e}_{z}, \qquad (2.6)$$

where e_x , e_y , and e_z denote the unit vectors along the *x*, *y*, and *z*-axis, respectively. At equilibrium, both transverse components in Eq. (2.6) will be zero, i.e. $M_x = 0$ and $M_y = 0$, because of the random phase of the individual magnetic dipole moments when they precess around the $B_0 = B_0 e_z$ field. The *z*-component, on the other hand, is nonzero:

$$\boldsymbol{M} = M_{z}\boldsymbol{e}_{z} = M_{0}\boldsymbol{e}_{z} = \left(\sum_{n=1}^{N_{\uparrow}} \frac{1}{2}\gamma\hbar - \sum_{n=1}^{N_{\downarrow}} \frac{1}{2}\gamma\hbar\right)\boldsymbol{e}_{z} = \frac{1}{2}(N_{\uparrow} - N_{\downarrow})\gamma\hbar\boldsymbol{e}_{z}, \qquad (2.7)$$

21

where we have substituted the two possible states of the proton magnetic moment. Furthermore, by combining Eq. (2.7) with Eq. (2.5), it follows that the magnitude of the bulk magnetization vector at equilibrium is equal to

$$|\mathbf{M}| = M_0 = \frac{\gamma^2 \hbar^2 B_0 N}{4k_B T}.$$
(2.8)

Equation (2.8) indicates how the magnitude of M is directly proportional to the strength of the B_0 field and the total number of spins N. The former is characteristic to the object being imaged and can be used as image contrast to create proton density weighted images, while γ , \hbar , and k_B are constants. Therefore, only B_0 or the absolute temperature T are controllable parameters that can influence |M|. Given that MRI experiments are typically carried out at room temperature to maximize patient's comfort, one is limited to changing the magnetic field strength B_0 for an overall increase of the bulk magnetization and NMR signal. This explains why high magnetic field strengths result in better signal-to-noise ratio (SNR) of the scans, and why scanning at ultra-low field (ULF) strengths inherently suffers from low SNR, demanding increased scan times to acquire enough NMR signal.

2.3.3 Excitation

In the previous section, it was stated that the external static magnetic field B_0 leads to a net magnetization vector M that is parallel to the direction of B_0 . However, this phenomenon in itself does not allow to detect an NMR signal. To detect the magnetization, an electrically conducting receiver coil is placed around the subject, perpendicular to the transverse plane, such that a time-varying rotating transverse magnetization component $(M_{x,v})$ will induce an (alternating) voltage in the coil (cf. Faraday's law of induction). The amplitude of the voltage, ϵ_{i} is proportional to the negative of the rate of temporal variation of the flux¹ ϕ_{i} which in turn is proportional to the magnitude of the transverse magnetization normal to the coil ($\phi \propto M_{x,y}$), i.e. $\epsilon \propto -\frac{d\phi}{dt}$. So, to generate measurable signals in such a receiver coil, the magnetization vector $oldsymbol{M}$ should be tilted from the equilibrium position into the transverse xy-plane. To enable the tilting, an additional time varying magnetic field B_1 is briefly turned on, which is perpendicular to B_0 , and oscillates with $\omega_1 = \omega_0$. This B_1 field is referred to as the radio frequency (RF) pulse (Rabi et al., 1938), because of its short-lived effect. Note that B_1 is typically about 5 orders of magnitude weaker than B_0 . In Fig. 2.2, the motion of the magnetization M is shown when a resonant RF field B_1 is applied, both in the reference laboratory² frame (Fig. 2.2(a)) and in a reference frame rotating² at $\omega_0 = \omega_1$ along with M and B_1 (Fig. 2.2(b)). In the laboratory frame, the magnetization spirals down

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} : \mathbf{R}(t) = \begin{pmatrix} \cos(\omega_1 t) & \sin(\omega_1 t) & 0 \\ -\sin(\omega_1 t) & \cos(\omega_1 t) & 0 \\ 0 & 0 & 1 \end{pmatrix},$$
(2.9)

with $\mathbf{R}(t)$ being a rotation matrix, i.e. $\mathbf{R}(t)\mathbf{R}^{\mathsf{T}}(t) = \mathbf{I}$ and $\det(\mathbf{R}(t)) = 1$.

¹At higher magnetic field strengths, the magnetization precesses at a higher frequency, and so the value of $\frac{d\phi}{dt}$ increases. As a result, stronger magnetic fields yield improved signal strength not only because of a larger nuclear polarization but also because of the additional increase in magnetic flux.

² Often, it is more convenient to visualize the effects of rotating magnetization vectors in the **laboratory** (i.e. stationary) frame of reference (x, y, z) using an alternative reference frame representation (Rabi et al., 1954). In this rotating frame of reference (x', y', z'), the reference frame rotates about the *z*-axis at angular frequency ω_1 (cf. Fig. 2.2). As such, a magnetization vector rotating at ω_1 in the laboratory frame will appear stationary in the rotating frame of reference. Both reference frames are connected using the following reference frame coordinate transformation (Tourais et al., 2022):

towards the *xy*-plane on the surface of a sphere with radius $|\mathbf{M}|$. Indeed, $|\mathbf{M}|$ remains constant as the RF field rotates the spin distribution as a whole. In the rotating frame, the magnetization rotates perpendicular to \mathbf{B}_1 at angular frequency ω_1 . Applying the RF pulse during a time interval Δt (order ~ms), flips the magnetization over an angle $\alpha = \gamma B_1 \Delta t$, termed as the *flip angle*. In most MR sequences, RF pulses are applied so that flip angles are 90° or 180°.



Figure 2.2: Evolution of magnetization vector M experiencing a static longitudinal magnetic field B_0 and a transversal time varying magnetic field B_1 , (a) in the laboratory frame of reference (x, y, z), and (b) in the rotating frame of reference (x', y', z').

In addition, the motion of the magnetization vector M in the static field B_0 from the moment the RF field B_1 is switched off and after the vector has been rotated over a certain angle α , can be described using classical electromagnetism. Given that there exist no other interactions than with the static magnetic field B_0 along the *z*-axis, the macroscopic magnetization M will experience a torque:

$$\frac{dM}{dt} = \gamma \left(\boldsymbol{M} \times \boldsymbol{B}_{0} \right)$$

$$= \gamma \begin{vmatrix} \boldsymbol{e}_{x} & \boldsymbol{e}_{y} & \boldsymbol{e}_{z} \\ \boldsymbol{M}_{x} & \boldsymbol{M}_{y} & \boldsymbol{M}_{z} \\ \boldsymbol{B}_{0,x} & \boldsymbol{B}_{0,y} & \boldsymbol{B}_{0,z} \end{vmatrix} = \gamma \begin{vmatrix} \boldsymbol{e}_{x} & \boldsymbol{e}_{y} & \boldsymbol{e}_{z} \\ \boldsymbol{M}_{x} & \boldsymbol{M}_{y} & \boldsymbol{M}_{z} \\ \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{B}_{0} \end{vmatrix}$$

$$= \gamma B_{0} M_{y} \boldsymbol{e}_{x} - \gamma B_{0} M_{x} \boldsymbol{e}_{y}. \qquad (2.10)$$

The components of $\frac{dM}{dt}$ along the x, y, and z-direction are then given by:

$$\frac{dM_x}{dt} = \gamma B_0 M_y, \qquad \frac{dM_y}{dt} = -\gamma B_0 M_x, \qquad \frac{dM_z}{dt} = 0.$$
(2.11)

These differential equations can be solved using the definition of the Larmor frequency $\omega_0 = \gamma B_0$, resulting in an expression for the components of M(t) in the laboratory frame of reference:

$$\begin{bmatrix} M_x(t) \\ M_y(t) \\ M_z(t) \end{bmatrix} = \begin{bmatrix} \cos(\omega_0 t) & \sin(\omega_0 t) & 0 \\ -\sin(\omega_0 t) & \cos(\omega_0 t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} M_x(0) \\ M_y(0) \\ M_z(0) \end{bmatrix}$$
(2.12)

Equation (2.12) denotes that the macroscopic magnetization vector precesses around the direction of the external static magnetic field, as shown in Fig. 2.3, with angular frequency ω_0 , the Larmor frequency.



Figure 2.3: Schematic representation of the torque on the macroscopic magnetization vector M, caused by the static magnetic field B_0 .

2.3.4 Relaxation

In reality, the free precession of the magnetization vector M, as described by Eq. (2.11) and Fig. 2.3, happens only temporary: the magnetization will return back to its equilibrium state parallel to the static magnetic field B_0 due to interactions of the spins with their surroundings. This process is called *relaxation*. Phenomenologically, the relaxation process can be added to Eq. (2.10), which results in a general set of equations that describe the evolution of M in time under the influence of an arbitrary magnetic field B. These equations are called the *Bloch equations* (Bloch, 1946):

$$\frac{d\boldsymbol{M}}{dt} = \gamma \left(\boldsymbol{M} \times \boldsymbol{B}_{0} \right) - \begin{bmatrix} T_{2}^{-1} & 0 & 0\\ 0 & T_{2}^{-1} & 0\\ 0 & 0 & T_{1}^{-1} \end{bmatrix} \boldsymbol{M} + \begin{bmatrix} 0\\ 0\\ M_{0}T_{1}^{-1} \end{bmatrix}, \quad (2.13)$$

where again M_0 describes the net magnetization vector at equilibrium, and where T_2 and T_1 are the transverse and longitudinal relaxation times, respectively. In the reference rotating (x', y', z')-frame, Eq. (2.13) simplifies to:

$$\frac{dM_{x',y'}}{dt} = -\frac{M_{x',y'}}{T_2}
\frac{dM_{z'}}{dt} = -\frac{M_{z'} - M_0}{T_1},$$
(2.14)

The longitudinal relaxation time, T_1 , characterizes the relaxation process of the longitudinal component $M_{z'}$, whereas the transverse relaxation time, T_2 , describes the relaxation curve of the transverse component $M_{x',y'}$. The causes of relaxation are diverse, though typically, one distinguishes two prominent types of relaxation processes: *spin-lattice relaxation* and *spin-spin relaxation*.

Spin-lattice relaxation

Spin-lattice relaxation, also known as longitudinal or T_1 -relaxation, stems from the exchange of energy with other degrees of freedom in the spin system in order to redistribute the population of the nuclear spin states. In the NMR jargon, these degrees of freedom are referred to as *the lattice*. The energy is dissipated to the surrounding lattice by means of molecular vibrations (phonons). This interaction with the lattice results in reorientation of the magnetic moments, causing a redistribution of the spins. When the thermal equilibrium is restored, more spins will again occupy the lower energy state (spin up), thereby satisfying Eq. (2.4). Phenomenologically, the spin-lattice relaxation process results in a change of $M_{z'}$ back to the equilibrium amplitude (Fig. 2.4), which is characterized by T_1 , and whose evolution is fully determined by Eq. (2.14). After a 90° RF-pulse, the evolution of the $M_{z'}$ component as a function of time t can be written as

$$M_{z'}(t) = M_0 \left[1 - \exp\left(-\frac{t}{T_1}\right) \right], \qquad (2.15)$$

with M_0 the longitudinal magnetization at equilibrium.



Figure 2.4: T_1 relaxation. After a 90° excitation pulse, due to energy dissipation with the surrounding lattice, the longitudinal component of the magnetization $M_{z'}$ regrows exponentially, over time, to the equilibrium value M_0 .

Spin-spin relaxation

Spin-spin relaxation, also called transversal or T_2 -relaxation, describes the recovery of the magnetization components $M_{x',y'}$ perpendicular to the static magnetic field, which is caused by a dephasing of spins (Fig. 2.5). The dephasing or loss in phase coherence of the spins is caused by local fluctuations in the magnetic field which are two-fold. First, the movement of electrons and nuclei creates rapidly fluctuating magnetic field inhomogeneities leading to irreversible transverse relaxation termed T_2 -decay. Second, the static magnetic field B_0 is spatially inhomogeneous, both inherently and due to differences in magnetic susceptibility between tissue types. These type of inhomogeneities also lead to transverse relaxation. When both types of field inhomogeneities are considered, a *pseudo transversal relaxation time* is introduced, called T_2^* , which is related to T_2 by the following inverse relation (Chavhan et al., 2009):

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'},\tag{2.16}$$

with $\frac{1}{T_2'} = \gamma \Delta B_{\text{inhomo}}$ the relaxation rate contribution attributable to magnetic field inhomogeneities ($\Delta B_{\text{inhomo}} \ge 0$) across a voxel. Note that since $\gamma > 0$, T_2^* is always shorter or equal to T_2 , resulting in a faster transversal decay. Furthermore, the solution of Eq. (2.14) for $M_{X',Y'}$ has also an exponential form,

$$M_{x',y'}(t) = M_{x',y'}(0) \exp\left(-\frac{t}{T_2^*}\right),$$
(2.17)

with $T_2^* = T_2$ when field inhomogeneities are disregarded, and where $M_{x',y'}(0)$ represents the transversal component of the net magnetization vector in the RF-rotating frame immediately after the $\alpha = 90^\circ$ RF-pulse.

The relaxation times T_1 and T_2 vary among different tissue types and in certain pathological states. A comprehensive review of reported normal *in vivo* relaxation times at 3T is provided by Bojorquez et al. (2017). It is important to note that both T_1 and T_2 relaxation times depend on the magnetic field strength; they are not inherent biomarkers of a certain tissue type (Korb & Bryant, 2002).



Figure 2.5: T_2 relaxation. Exponential decay, over time, of the transverse component of the magnetization $M_{x',y'}$ in the rotating frame of reference. The transverse component recovers to the equilibrium value zero after excitation. The coherence loss of magnetic moments precessing in the transverse plane and the corresponding net transverse magnetization $M_{x',y'}(t)$ is shown right after the excitation pulse and for two subsequent time points.

2.4 MR image formation and readout strategies

As explained in section 2.3.3, by generating rotating transverse magnetization using an RF pulse, the MR signal can be detected. However, such a signal is not yet linked to a specific spatial location. Spatial encoding of the MR signal is an essential step to perform MRI. It can be achieved by superimposing time dependent magnetic field gradients that vary linearly in space onto the main magnetic field B_0 . These gradients can be applied in any direction by combining the gradient coils, allowing for various strategies of spatial encoding. The most common method is Cartesian readout along the three orthogonal spatial directions: frequency encoding (FE) in the *x*-direction, phase encoding (PE) in the *y*-direction, and slice encoding strategies, most prominently two-dimensional (2D) and three-dimensional (3D) readout. In 2D readout, images are acquired slice by slice, each slice being encoded separately. In 3D readout, the entire volume is encoded and imaged simultaneously. Both readout strategies are discussed hereafter.

2.4.1 2D readout

When a slice encoding magnetic gradient G_z is applied along the *z*-axis, the strength of the total magnetic field in a plane at location *z* is equal to $B_0 + G_z z$. This results in the angular frequency of the precessing spins becoming dependent on the location *z*. When such a slice encoding gradient is applied simultaneously with an RF pulse rotating at the Larmor frequency, that RF pulse will be off-resonant for all spins at locations $z \neq 0$, whereas, theoretically, only spins at z = 0 that are precessing at the Larmor frequency would be excited. However, a real RF pulse is characterized with a finite bandwidth:

$$\Delta \omega = \gamma G_Z \Delta z. \tag{2.18}$$

As a result, spins within the frequency range $[\omega_L - \Delta \omega/2, \omega_L + \Delta \omega/2]$ will be excited in a slice with a thickness equal to Δz . The slice position can be changed by using RF pulses with a frequency $\omega_1 = \omega_L + \delta \omega$, with $\delta \omega$ a certain frequency offset.

After slice selection, application of specifically timed frequency encoding gradients G_x and phase encoding gradients G_y along the x- and y-axis allows to spatially encode the MR signals within the excited slice (Liang et al., 2000; Bloembergen, 1957):

$$S(k_x, k_y) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} S(x, y) \exp\left(-i2\pi(k_x x + k_y y)\right) dx dy,$$
(2.19)

with S(x, y) the magnitude of the magnetization vector at the spatial location with coordinate (x, y). The wave numbers k_x and k_y are equal to the temporal integral of the magnetic gradients applied for spatial encoding. More specifically,

$$k_{\rm x} = \frac{\gamma}{2\pi} \int_0^T G_{\rm x}(t) dt, \qquad (2.20)$$

$$k_{y} = \frac{\gamma}{2\pi} \int_{0}^{T} G_{y}(t) dt, \qquad (2.21)$$

where T denotes the total time duration to acquire $S(k_x, k_y)$, where it is assumed that $G_x(t)$ and $G_y(t)$ are only nonzero at specific times when the respective gradients are switched on.



(e) horizontal *k*-space slab = vertical blurring

(f) k-space rotation = image space rotation

Figure 2.6: Six examples that illustrate some basic operations using the Fourier transform (image-to-k space) and inverse Fourier transform (k-to-image space) in MR imaging. For each example, the k-space is shown on the left, while the image space is shown on the right.

The MR signals $S(k_x, k_y)$ are acquired in so-called *k*-space. This space can be sampled at multiple frequencies (k_x, k_y) so to obtain a 2D data set in *k*-space. It follows from Eq. (2.19)

that the MR signal $S(k_x, k_y)$ is the **Fourier transform** of the spin density S(x, y). Therefore, an image of the spin density in the so-called spatial or *image* domain can be obtained by performing a 2D inverse Fourier transform on the 2D data set in *k*-space (Kumar et al., 1975).

Fig. 2.6 illustrates the effect of the (inverse) Fourier transform in MR imaging for a 2D brain image with an additional high intensity dot in the image. Note that both the k-space (shown left) and real space image (shown right) are originally complex valued, yet their respective magnitudes are shown here on a log-scale. Different properties of the Fourier transform are highlighted. In Fig. 2.6a, the original image is computed as the inverse Fourier transform of the fully sampled k-space. In Fig. 2.6b, subsampling is performed along the horizontal dimension by adding zeros every other vertical k-space line, which results in aliasing in the image space along that dimension. In Fig. 2.6c, the effect of low-pass filtering on the original k-space domain is illustrated, which maintains only the contrast of the underlying image in image space. The information of the high spatial frequencies, that contains the details and contours of the brain, has disappeared. In Fig. 2.6d, high-pass filtering, where only the high spatial frequencies have been selected in the k-space, is demonstrated, providing only information about the details and edges of the brain in the image domain. Next, Fig. 2.6e demonstrates the selection of a horizontal slab in k-space, which reduces resolution in the vertical dimension in image space. Finally, Fig. 2.6f illustrates how a rotation in k-space corresponds with a rotation in image space, and vice versa.

Spin echo

The spin echo (SE) sequence is a well-known two-pulse sequence to generate contrastweighted images (Hahn, 1950; Carr & Purcell, 1954). It is a fundamental pulse sequence that forms the basis for many of the more advanced pulse sequences in MRI. As such, it is illustrative to briefly explain some core principles of MRI acquisition using this sequence. The SE sequence is schematically represented in Fig. 2.7. It uses a $\pi/2 - TE/2 - \pi - TE/2 - echo$ pulse sequence, where TE represents the echo time, denoting the time between the excitation pulse and the time of the readout. First, a single 90° excitation pulse rotates the magnetization, within a certain slice, into the (x', y')-plane. Next, the different spin packets start to dephase in the transversal plane due to all effects contributing to the T_2^* relaxation. At t = TE/2, the magnetization is flipped by applying a 180° refocusing pulse. Finally, after another period of TE/2, the spins are rephasing, thus producing a measurable echo signal. The signal decay at TE, compared to the start of the experiment, now solely originates in the T_2 relaxation. The time between two repetitions of the SE sequence, is called the *repetition time* (TR). TEs are typically in the order of tens of milliseconds, while TRs range in the order of seconds. Variation of the TR and TE allows to obtain different contrast-weighted images. For T_1 -weighted contrast, a short TR and short TE are used. For proton density-weighted contrast, a long TR and short TE are applied. For T_2 -weighted contrast, a long TR and long TE are employed. Depending on the external magnetic field strength, different TR and TE values are needed to obtain the same contrast-weighted image.



Figure 2.7: A: Spin Echo (SE) Sequence: The net magnetization vector is tipped into the transverse plane by a 90° pulse. During a time TE/2, the free-induction decay (FID) signal decays due to strong T_2^* dephasing. At time t = TE/2, the magnetization is flipped by a 180° refocusing pulse. The spin packets continue to rotate, and after a time TE the different magnetization vectors are again in phase along the -y' axis, thereby producing a measurable echo. By acquiring multiple MR signals (images) at different TEs, the decay of the transversal magnetization towards zero can be sampled. The time between two repetitions of the spin echo sequence, is called the repetition time TR. **B**: Effect of the spin echo sequence on the net magnetization vector as seen in the RF-rotating frame. B_0 represents the external magnetic field vector. (a) Initial net longitudinal magnetization in alignment with B_0 , (b) a 90° pulse rotates the magnetization 90° on the y' axis, (c) during a time TE/2 a dephasing of the different spin packets occurs, (d) at t = TE/2 a 180° pulse flips the sign of the transversal y'-component of the individual magnetization vectors. (e) The spin packets continue to precess at the same frequency, and after a time TE the different magnetization vectors are again in phase along the -y' axis. At this point the echo is maximal, and detected in the transversal plane.

Multi-echo spin echo

As an extension of a standard SE sequence, multi-echo spin echo (MESE) sequences have been proposed (Feinberg et al., 1985), which stimulate the spin system to repeatedly rephase using a train of 180° pulses (Fig. 2.8). As long as T_2 relaxation is not complete and MR signal is present, this allows to generate extra echoes within a given repetition time. The amplitude of each echo is progressively smaller due to the T_2 decay. Also, the echo time spacing (TE), i.e. the time between consecutive spin echoes, is inherently fixed. MESE sequences offer the advantage of acquiring different echoes, i.e. different T_2 -weighted images, in a single TR, which makes them suitable for quantitative T_2 mapping protocols.



Figure 2.8: Schematic representation of a multi-echo spin echo (MESE) sequence. Compared to a standard spin echo sequence, the MESE sequence stimulates the spin system with additional 180° pulses. As long as T_2 -relaxation is not complete and MR signal is present, this allows to generate extra echoes within a given repetition time. The amplitude of each echo is progressively smaller due to the T_2 decay. Also, the echo time spacing (TE), i.e. the time between consecutive spin echoes, is inherently fixed.

Inversion recovery

The gold standard method to create images with T_1 -weighted contrast is the inversion recovery (IR) sequence (Drain, 1949; Hahn, 1949), as shown in Fig. 2.9. In this pulse sequence, the longitudinal net magnetization vector is initially flipped by a 180° inversion pulse. Next, during an *inversion time* TI, the longitudinal magnetization component $M_{z'}$ will have partly relaxed to equilibrium according to the Bloch equation Eq. (2.14). Tissues with different T_1 relaxation values recover at different rates, creating a T_1 contrast among them. At time t = TI, the differences in the longitudinal magnetization are converted into differences in the transverse magnetization, by applying a 90° excitation pulse. Again, similar as for a spin echo sequence, this gives rise to a free induction decay (FID) signal. The amplitude of this FID signal depends on the recovery of the longitudinal magnetization component $M_{z'}(t)$ during the period TI, which is given by:

$$M_{z'}(t) = M_{z'}(0) \left[1 - 2 \exp\left(-\frac{t}{T_1}\right) \right].$$
 (2.22)

By acquiring multiple MR signals (images) at different TIs, the recovery of the longitudinal magnetization towards its equilibrium value can be sampled.

Inversion recovery fast/turbo spin echo (IR FSE/TSE)

As illustrated in Fig. 2.10, in a fast spin echo (FSE) or turbo spin echo (TSE) sequence an echo train of evenly spaced refocusing pulses is used to acquire multiple phase encoding lines of the data, i.e. a different phase encoding line is acquired for each echo. The time between the successive echoes is called the inter echo spacing (IES). As the refocusing RF pulses are all evenly spaced in time, also the IES remains fixed. Furthermore, the echo train length (ETL) or turbo factor (TF) denotes the number of echoes in the spin echo train. When FSE is combined with an IR module, the maximum number of slices that can be acquired within one TR will not only depend on the multi-slice readout, but also on the ETL and IES.



Figure 2.9: A: Inversion Recovery Sequence: The longitudinal net nuclear magnetization vector is inverted by a 180° pulse. After inversion time TI, the longitudinal component is tipped into the transverse plane by a 90° pulse, after which the (T_1 -weighted) MR signal is measured. By acquiring multiple MR signals (images) at different inversion times, the recovery of the longitudinal magnetization towards its equilibrium value can be sampled. The time between two repetitions of the sequence, i.e. the time between the inversion pulses, is called the repetition time TR. **B**: Effect of the inversion recovery sequence on the net nuclear longitudinal magnetization vector $M_{z'}$ as seen in the RF-rotating frame. B_0 represents the external magnetic field vector. (a) Initial net nuclear longitudinal magnetization in alignment with B_0 , (b) the 180° pulse inverts the longitudinal magnetization $M_{z'}$, (c)-(d) the longitudinal magnetization $M_{z'}$ relaxes and recovers to equilibrium, (e) after an inversion time TI the relaxing longitudinal magnetization $M_{z'}$ is tipped into the transverse plane by a 90° pulse before readout.



Figure 2.10: Schematic representation of an inversion recovery fast spin echo (IR-FSE) sequence. Compared to a standard inversion recovery sequence, the IR-FSE sequence uses an echo train with multiple refocusing 180° pulses to acquire multiple phase encoding lines of the image in a single repetition time (TR). The inter echo spacing (IES) is fixed when the refocusing pulses are evenly spaced in time.

For sequential IR FSE, the number of slices is given by:

$$N_{\rm slice} = \frac{{\sf TR}}{{\sf TI} + ({\sf IES} \times {\sf ETL})}. \tag{2.23}$$

For very high ETL, the number of slices is restricted, indicating that a trade-off exists between the number of slices and the scan time acceleration.

In theory, by using long echo trains, IR FSE provides the opportunity to acquire multiple T_1 -weighted images at a high spatial resolution within a short scan time, making it a go-to sequence for quantitative T_1 mapping. However, in practice, the scan time reduction is somewhat limited by the high specific absorption rate (SAR) that inherently comes with the high RF energy deposition of the repeated RF pulses in the echo train (Weigel et al., 2007). SAR levels can be reduced by decreasing the number of slices acquired within one TR, either by acquiring thicker slices (and decreasing the spatial resolution), or by increasing the number of excitations (thus increasing the scan time). Clearly, there exists a need for methods that can increase spatial resolution without sacrificing scan time, and which are not limited by SAR restrictions.

Echo planar imaging

Single-shot echo planar imaging (ss-EPI) is performed using a pulse sequence in which multiple echoes of different phase steps are acquired using rephasing gradients as opposed to repeated 180° RF pulses following a standard spin echo sequence (Mansfield, 1977). ss-EPI is designed to collect a 2D image as rapidly as possible. After excitation (i.e., slice selection), the entire 2D k-space is traversed by an efficient use of time-varying gradients G_x and G_y . The EPI readout strategy is illustrated in Fig. 2.11. After shortly switching on G_{ν} , a line in the Cartesian coordinate system is sampled during the application of G_x (i.e., different frequencies k_x for a fixed frequency k_y). Subsequent lines for different k_y are sampled by applying the phase encoding gradient for a very short time (a so-called *blip*), in between the positive and negative lobes of the frequency encoding gradient. When data from multiple slices is required, the entire readout procedure is repeated, with slice excitation at different locations along the z-axis. While EPI has the benefit of being extremely fast, it is prone to several imaging artifacts. The most prominent ones are ghosting and potential distortions along the phase-encoding direction (Hu et al., 2020). Nevertheless, ss-EPI remains a method of choice in diffusion and perfusion imaging. This is partly because it is a highly efficient method that can produce whole-brain volumes in 10 seconds or less. This enables the acquisition of tens or even hundreds of image volumes with different diffusion encodings or post-labeling delay times. Moreover, ss-EPI is fairly robust to in-plane motion, since a complete slice can be acquired in the order of 100 ms, which effectively "freezes" typical head motion (Skare et al., 2018).

2.4.2 3D readout

The excitation of a slice, as described in the previous section for 2D readout, can also be performed for a thicker slab. Indeed, following Eq. (2.18), the thickness of the excited volume is equal to $\Delta z = \Delta \omega / (\gamma G_z)$, and can be increased by increasing the bandwidth $\Delta \omega$ of the RF pulse or by reducing the gradient G_z applied during excitation. Once such a thicker slab is prepared, spatial encoding is performed in the three orthogonal directions: phase encoding



Figure 2.11: Single-shot EPI sequence: Schematic representation of a ss-EPI readout (left), and the associated *k*-space traversal during readout (right).

along the *z*-axis and *y*-axis and frequency-encoding along the *x*-axis (Bernstein et al., 2004). Thus, MR signals are encoded in the three spatial directions within the excited slab:

$$S(k_x, k_y, k_z) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} S(x, y, z) \exp\left(-i2\pi(k_x x + k_y y + k_z z)\right) dx dy dz, \quad (2.24)$$

with k_z the wave number to the temporal integral of the magnetic gradient used for encoding along the *z*-axis. In this case, a 3D image of the spin densities is obtained by performing a 3D inverse Fourier transformation on the acquired 3D *k*-space data set.

Compared to 2D readout, 3D readout is more SNR-efficient because a much larger volume is excited within a single excitation. When performing the discrete 3D inverse Fourier transformation instead of the 2D version, more k-space data points contribute to the generation of each data point in the spatial domain. For a formal comparison of the SNR efficiency between 2D and 3D readout, the reader is referred to section 11.6.1 in the work of Bernstein et al. (2004).

2.5 The MRI trade-off

In MRI, a trade-off exists between the scan time, the spatial resolution, and the signal-to-noise ratio (SNR). Fig. 2.12 shows a visual representation of this trade-off.

MRI allows to record 3D images in high resolution, either using 2D readout (section 2.4.1) or 3D readout (section 2.4.2). A problem with high resolution imaging is the prolonged scan time (cf. Fig. 2.12, top left), since multiple RF excitations are required to obtain sufficiently high SNR values. A certain time is needed between repeated excitations of a volume (i.e., the TR) and this time depends on the type of acquisition. The TR is especially long when each new excitation of a specific volume requires T_1 relaxation to obtain the desired contrast, as the T_1 relaxation time is of the order of seconds in most tissue types. With standard 2D readout, every excitation influences only a single slice, so within each TR, a part of the *k*-space of all slices can be recorded. However, when the slices are very thin for high resolution imaging, the total signal power emitted by the slice is low for a



Figure 2.12: Visual representation of the MRI trade-off.

short scan time, resulting in a low SNR (cf. Fig. 2.12, bottom). To increase the SNR, it is possible to average the signal of multiple excitations, requiring scan times that are longer than the minimal TR for T_1 relaxation. This extends the total acquisition time compared to an acquisition with thicker slices. Often, therefore, a compromise between scan time and spatial resolution is found in the acquisition of slices that are substantially thicker than the in-plane resolution, i.e. images are acquired with an anisotropic resolution (cf. Fig. 2.12, top right). In this way, a good in-plane resolution is combined with a (relatively) short scan time. The obvious disadvantage is the reduced spatial resolution in the dimension in which the slices are stacked. In the following sections, each component of the MRI trade-off is discussed more carefully.

2.5.1 Scan time

The scan time or acquisition time for 2D multi-slice imaging depends on several factors, such as the number of signal averages N_{SA} , the repetition time TR, the number of phase encoding steps N_{PE} , and the number of slices N_s when considering a 2D (multi-slice) readout scheme:

scan time
$$\propto N_{\text{SA}} \cdot \text{TR} \cdot \frac{N_{\text{PE}}}{N_{\text{PE/TR}}} \cdot \frac{N_s}{N_{s/\text{TR}}},$$
 (2.25)

with $N_{\text{PE/TR}}$ and $N_{s/\text{TR}}$ the number of phase encoding lines and the number of slices that are acquired within one TR, respectively.

In addition, in-plane acceleration or so-called **parallel imaging (PI)** offers a robust way to provide scan time reduction by acquiring a reduced amount of *k*-space data with an array

of receiver coils (Deshmane et al., 2012). The undersampled data can be acquired faster, but the undersampling results in images corrupted by aliasing. To reconstruct MR images without undersampling artifacts, imaging algorithms such as Sensitivity Encoding (SENSE, (Pruessmann et al., 1999)) or Generalized Auto-Calibrating Partially Parallel Acquisition (GRAPPA, (Griswold et al., 2002)) are commonly used. Although PI leads to scan time reduction, the penalty for acquiring fewer signals using PI is a loss of the SNR in the reconstructed image, typically by a factor of the square root of the acceleration factor, due to reduced signal averaging (Robson et al., 2008). Furthermore, image noise in PI is further amplified by the ill-conditioning of the image reconstruction process. In addition, the noise amplification is spatially variant and depends on the specific geometry of the RF coil array being used, making the development of robust noise models more challenging.

In Chapter 7, another more recent approach is mentioned to accelerate the data acquisition. This technique, referred to as **simultaneous multi-slice (SMS) or multiband**, allows for the simultaneous excitation and imaging of several slices using 2D readout (Barth et al., 2016). The primary benefit is an acceleration in data acquisition that is equal to the number of simultaneously excited slices. Furthermore, unlike in-plane parallel imaging, SMS only has a marginal intrinsic SNR penalty, and the full acceleration is attainable at fixed echo time, as is required for many EPI applications, e.g. for the acquisition of 2D perfusion MRI data.

2.5.2 Spatial resolution

In MRI, the spatial resolution of an image is more accurately defined by the effective width of the system's point spread function (PSF), which describes the system's ability to distinguish between two point sources. This is often referred to as *two-point resolution*: the minimum distance between two point sources that can be distinguished as separate entities. The voxel size, given by $[\Delta x, \Delta y, \Delta z]$, also plays a significant role in determining the spatial resolution, but it should not be confused with the inherent resolution of the imaging system.

A voxel is the 3D volumetric equivalent of a 2D pixel. The through-plane resolution is defined by the slice thickness Δz . The in-plane resolution is defined by:

$$\Delta x = \frac{\text{FOV}_x}{N_{\text{FE}}}, \qquad \Delta y = \frac{\text{FOV}_y}{N_{\text{PE}}}, \qquad (2.26)$$

with N_{FE} the number of frequency encoding steps, N_{PE} the number of phase encoding steps, and FOV is the field of view, which refers to the area over which an MR image is acquired (or displayed). The image matrix size is defined as $N_{\text{FE}} \times N_{\text{PE}}$.

The through-plane spatial resolution can be improved by reducing the slice thickness, either by using a stronger slice-encoding gradient or a narrower RF pulse bandwidth. Thinner slices are less susceptible to partial volume effects, i.e. the effect where a voxel with nominal resolution will consist of a mixture of signals stemming from different anatomical structures at smaller resolution scale. Thinner slices will also contain fewer proton spins and thus will emit less signal. Moreover, decreasing the slice thickness increases the number of slices needed for a full coverage of the subject, which in turn might increase the acquisition time. In practice, the voxel size and resolution are constrained by the gradient strength, acquisition time, and targeted SNR.

Moreover, the through-plane PSF in MRI is significantly influenced by the Fourier relationship between the slice profile and the finite duration of the slice selection pulse (see also Fig. 3.3

in section 3.4.6.4 hereafter) (Noll et al., 1997). This relationship determines the effective width of the PSF, which affects resolution in the through-plane direction. Specifically, the slice profile in k-space, derived from the temporal profile of the RF pulse and gradient characteristics, defines how frequencies are encoded along the slice direction.

Spatial resolution and k-space

A key difference between MRI and other medical imaging modalities is the control the user has over how the data is acquired, manipulated and reconstructed into an image. By adjusting, among other things, the timing of the pulses in an MR sequence, the order of data acquisition, or the strengths and gradients of the auxiliary magnetic fields, the user can change the resolution, the field of view (FOV), the contrast, the acquisition time, and so on. As explained in section 2.4, the agent of this control is k-space, the abstract platform onto which data are acquired, positioned, and then transformed into the desired image (Mezrich, 1995). No such flexible agent exists for X-ray imaging, ultrasound, or positron emission tomography, which prevents these methods from supporting a rich interaction between the user and the image, similar as with MRI. However, the price for this rich interaction in MRI is the need for an intuitive understanding of the concepts and mechanisms of k-space manipulation.

For research that intends to improve the spatial resolution of images or parameter maps, as investigated in this thesis, it is important to understand the relationship between k-space and spatial resolution. Fig. 2.13 illustrates this relationship in more detail.

To avoid loss of information, the sampling interval, i.e. the distance between two *k*-space points (Δk_x along the frequency-encoding direction, and Δk_y along the phase-encoding direction) has to satisfy the Nyquist criterion³. In addition, the *k*-space sampling is finite, i.e. the signal $S(k_x, k_y)$ is not sampled for $|k_x| > k_{\max,x}$ and $|k_y| > k_{\max,y}$, with $k_{\max,x} = (N_{\text{FE}}/2)\Delta k_x$ and $k_{\max,y} = (N_{\text{PE}}/2)\Delta k_y$ the maximum frequency sampled in the frequency and phase encoding direction, respectively. Therefore, according to the Nyquist criterion, the largest acceptable pixel size of the image is given by (Mezrich, 1995):

$$\Delta x = \frac{1}{\mathsf{FOV}_{k,x}}, \qquad \Delta y = \frac{1}{\mathsf{FOV}_{k,y}}, \tag{2.27}$$

with $\text{FOV}_{k,x} = 2k_{\max,x}$ and $\text{FOV}_{k,y} = 2k_{\max,y}$. Since $\text{FOV}_x = N_{\text{FE}}\Delta x$ and $\text{FOV}_y = N_{\text{PE}}\Delta y$, the FOV will thus be determined by the sampling interval:

$$FOV_x = \frac{1}{\Delta k_x}, \qquad FOV_y = \frac{1}{\Delta k_y}.$$
 (2.28)

In Figs. 2.13(c)-(d) the inverse relationship between the spacing of the data samples (Δk_x and Δk_y) and the FOV is shown. When the spacing between the acquired data points is increased, the resulting image will have the same pixel size, but the FOV will be smaller. Since the Nyquist criterion is not fulfilled, the edges of the brain which fall outside the smaller FOV will wrap over the sides of the reconstructed images. This phenomenon is called **aliasing**.

 $^{^{3}}$ The Nyquist criterion, a.k.a. Nyquist-Shannon theorem, defines the minimum sample rate for the highest frequency that you want to measure. The Nyquist rate should be two times (2×) the given frequency to be measured accurately. If the Nyquist theorem is not met, higher frequency information is acquired in too low a sample rate, resulting in aliasing artifacts.



Figure 2.13: Relation between *k*-space sampling and image resolution/FOV. From a fully sampled *k*-space (a) the corresponding MRI image (b) can be computed. Undersampling of the *k*-space (c) results in aliasing in the image space (d). Decreasing the maximum sampled frequency (e), decreases the spatial resolution of the corresponding image (f).

Figs. 2.13(e)-(f) visualize the inverse relationship between the pixel size and the range of sampled frequencies in *k*-space. The sampling rate and spacing (Δk) is kept constant, but the N_{FE} and N_{PE} are reduced, which reduces the maximum acquired frequency k_{max} as well. This manipulation of *k*-space results in an increase of the pixel size (Δx , Δy). Thus, sampling high frequencies in *k*-space is required to achieve a high spatial resolution in MRI.

2.5.3 Signal-to-noise ratio

When an MRI scan is performed, the acquired data is known to be affected by several sources of quality deterioration due to limitations in the hardware, scanning time, or movement of patients. One source of degradation that affects most of the acquisitions is **noise**. The term noise in MRI can have different meanings depending on the context. It has been applied to degradation sources such as physiological and respiratory distortions in some MR applications and acquisitions schemes, or even acoustic sources (the sound produced by the pulse sequences in the magnet). In this section, the term *noise* is strictly limited to the thermal noise introduced during data acquisition, also known as Johnson-Nyquist noise (Aja-Fernández et al., 2016).

The principal source of thermal noise in most MR scans is the subject (or object to be imaged) itself, followed by electronic noise during the acquisition of the signal in the receiver chain (Edelstein et al., 1986; Jezzard et al., 1993; Krüger & Glover, 2001). It is produced by the stochastic motion of free electrons in the RF coil, which is a conductor, and by eddy current losses in the patient, which are inductively coupled to the RF coil (Aja-Fernández et al., 2016). The presence of noise over the acquired MR signal not only affects the visual assessment of an image, but it also may interfere with any post-processing steps such as segmentation, registration, functional MRI analysis, and in particular, with the numerical estimation of quantitative parameters in the context of qMRI applications. To this extent, image-derived metrics are typically used to compare the level of an expected signal to the level of noise corrupting the measurement of that signal. One such metric is the **signal-to-noise ratio (SNR)** of the MR scan. As its name suggests, it takes the ratio of the (power of the) signal and the (power of the) unwanted noise.

In MRI, the *signal* intensity depends on the specific pulse sequence and associated sequence parameters being used during acquisition, as well as on the spatial resolution or voxel dimensions (Edelstein et al., 1986):

signal
$$\propto \Delta x \Delta y \Delta z F_{\text{sequence}} F_{B_0}$$
 (2.29)

with F_{sequence} a sequence-dependent factor incorporating the influence of signal relaxation, i.e. F_{sequence} depends on chosen sequence parameters such as TE and TR, as described in section 2.4. Furthermore, as discussed in section 2.3.2, the magnitude of the bulk magnetization vector is directly proportional to the strength of the magnetic field B_0 , meaning that the observed signal will be larger for increased magnetic field strengths. This proportionality on B_0 is incorporated in the factor F_{B_0} .

The *noise*, on the other hand, is related to the bandwidth (BW) and a set of pulse sequence parameters N_{SA} , N_{PE} , N_{FE} (Dietrich et al., 2007):

noise
$$\propto \frac{\sqrt{BW}}{\sqrt{N_{SA}N_{PE}N_{FE}}}$$
. (2.30)

Briefly put, the bandwidth corresponds to the range of frequencies captured during readout of the *k*-space. Typically, the speed with which the *k*-space is traversed is proportional to the bandwidth. Hence, a larger bandwidth means that more information can be collected in a single readout, speeding up the acquisition. Alas, also the thermal noise power in the coil is proportional to the bandwidth, which means that an increased bandwidth leads to an increase of the noise level (Redpath, 1998). On the other hand, a low bandwidth increases the risk of chemical shift artefacts (Babcock et al., 1985). In addition, the selection of an appropriate RF receiver coil is essential to prevent noise amplification, e.g. nowadays dedicated head coils with 32 or even 64 transmitter/receiver channels are commonly used for brain MRI (Keil et al., 2013).

Given the aforementioned definitions of signal and noise, the SNR of a multi-slice MR image is given by:

$$SNR \propto \frac{\Delta x \Delta y \Delta z F_{sequence} F_{B_0} \sqrt{N_{SA} N_{PE} N_{FE}}}{\sqrt{BW}}.$$
(2.31)

Since Δx , Δy , and Δz define the spatial resolution (see section 2.5.2), and N_{SA} , N_{PE} , and N_{FE} define the scan time (see also Eq. (2.25)), it follows that the SNR is dependent on spatial resolution and scan time, i.e.:

$$SNR \propto (voxel size) \sqrt{scan time}$$
 (2.32)

As such, there exists a **trade-off between SNR**, **spatial resolution**, **and scan time** which complicates acquisition of an MRI image.

References

- Aja-Fernández, S., et al. (2016). *Statistical Analysis of Noise in MRI: Modeling, Filtering and Estimation*. Springer International Publishing, Cham, Switzerland, 1ed. ed.
- Babcock, E., Brateman, L., Weinreb, J., Horner, S., & Nunnally, R. (1985). Edge artifacts in mr images: Chemical shift effect. *Journal of computer assisted tomography*, *9*(2), 252–257.
- Barth, M., Breuer, F., Koopmans, P. J., Norris, D. G., & Poser, B. A. (2016). Simultaneous multislice (SMS) imaging techniques. *Magn. Reson. Med.*, 75(1), 63–81.

Bernstein, M. A., et al. (2004). Handbook of MRI pulse sequences. London: Elsevier Acad. press.

Bloch, F. (1946). Nuclear induction. Phys. Rev., 70, 460-474.

- Bloembergen, N. (1957). Proton relaxation times in paramagnetic solutions. *The Journal of Chemical Physics*, 27(2), 572–573.
- Bohr, N. (1976). Collected works. volume 3: The correspondence principle (1918–1923). J. R. Nielsen (ed.). Amsterdam: North-Holland Publishing.
- Bojorquez, J. Z., Bricq, S., Acquitter, C., Brunotte, F., Walker, P. M., & Lalande, A. (2017). What are normal relaxation times of tissues at 3 T? *Magn. Reson. Imaging*, *35*, 69–80.

Bushberg, J. T., Seibert, J. A., Leidholdt, E. M. J., & Boone, J. M. (2012). The Essential Physics of Medical Imaging, Third Edition. Lippincott Williams & Wilkins, Philadelphia, PA, USA.

Carr, H. Y., & Purcell, E. M. (1954). Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Phys. Rev.*, *94*, 630–638.

Chavhan, G. B., et al. (2009). Principles, Techniques, and Applications of T_2^* -based MR Imaging and Its Special Applications. *RadioGraphics*, 29(5), 1433–1449.

Damadian, R. (1971). Tumor detection by nuclear magnetic resonance. Science, 171(3976), 1151–1153.

- Damadian, R. (1974). Apparatus and method for detecting cancer in tissue. U.S. Patent no. US3789832A.
- Deshmane, A., Gulani, V., Griswold, M. A., & Seiberlich, N. (2012). Parallel MR imaging. J. Magn. Reson. Imaging, 36(1), 55–72.

- Dietrich, O., Raya, J. G., Reeder, S. B., Reiser, M. F., & Schoenberg, S. O. (2007). Measurement of signal-to-noise ratios in MR images: Influence of multichannel coils, parallel imaging, and reconstruction filters. *Journal of Magnetic Resonance Imaging*, 26(2), 375–385.
- Drain, L. E. (1949). A direct method of measuring nuclear spin-lattice relaxation times. Proc. Phys. Soc. London, Sect. A, 62, 301–306.
- Dreizen, P. (2004). The Nobel prize for MRI: a wonderful discovery and a sad controversy. *Lancet*, *363*(9402), 78.
- Edelstein, W. A., Glover, G. H., Hardy, C. J., & Redington, R. W. (1986). The intrinsic signal-to-noise ratio in NMR imaging. *Magnetic Resonance in Medicine*, 3(4), 604–618.
- Ernst, R. R., & Anderson, W. A. (1966). Application of fourier transform spectroscopy to magnetic resonance. *Review of Scientific Instruments*, 37(1), 93–102.
- Feinberg, D. A., et al. (1985). Multiple spin-echo magnetic resonance imaging. *Radiology*, 155(2), 437–442.
- Gorter, C. J. (1936). Negative result of an attempt to detect nuclear magnetic spins. *Physica*, *3*(9), 995–998.
- Griswold, M. A., et al. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.*, 47(6), 1202–1210.
- Hahn, E. L. (1949). An accurate nuclear magnetic resonance method for measuring spin-lattice relaxation times. *Phys. Rev.*, 76(1), 145–146.
- Hahn, E. L. (1950). Spin echoes. Phys. Rev., 80, 580-594.
- Hanson, L. G. (2008). Is quantum mechanics necessary for understanding magnetic resonance? *Concepts* in Magnetic Resonance Part A, 32A(5), 329–340.
- Hu, Z., Wang, Y., Zhang, Z., Zhang, J., Zhang, H., Guo, C., Sun, Y., & Guo, H. (2020). Distortion correction of single-shot EPI enabled by deep-learning. *NeuroImage*, 221, 117170.
- Jeener, J. (1971). Lecture notes from Ampere Summer School in Basko Polje, Yugoslavia, September, 1971.
- Jeener, J., Meier, B. H., Bachmann, P., & Ernst, R. R. (1979). Investigation of exchange processes by two-dimensional NMR spectroscopy. *The Journal of Chemical Physics*, 71(11), 4546–4553.
- Jezzard, P., LeBihan, D., Cuenod, C., Pannier, L., Prinster, A., & Turner, R. (1993). An investigation of the contribution of physiological noise in human functional MRI studies at 1.5 Tesla and 4 Tesla. In Proc. Soc. Mag. Reson. Med., vol. 12, (p. 1392).
- Keil, B., et al. (2013). A 64-channel 3T array coil for accelerated brain MRI. Magnetic Resonance in Medicine, 70(1), 248–258.
- Korb, J.-P., & Bryant, R. G. (2002). Magnetic field dependence of proton spin-lattice relaxation times. Magnetic Resonance in Medicine, 48(1), 21–26.
- Krüger, G., & Glover, G. H. (2001). Physiological noise in oxygenation-sensitive magnetic resonance imaging. Magnetic Resonance in Medicine, 46(4), 631–637.
- Kumar, A., Welti, D., & Ernst, R. R. (1975). NMR Fourier zeugmatography. *Journal of Magnetic Resonance (1969)*, 18(1), 69–83.
- Larmor, J. (1897). LXIII. On the theory of the magnetic influence on spectra; and on the radiation from moving ions. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(271), 503–512.
- Lauterbur, P. C. (1973). Image formation by induced local interactions: Examples employing nuclear magnetic resonance. *Nature*, 242, 190–191.
- Liang, Z., Lauterbur, P., in Medicine, I. E., & Society, B. (2000). *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. IEEE Press Series on Biomedical Engineering. Wiley.
- Liu, Y., et al. (2021). A low-cost and shielding-free ultra-low-field brain MRI scanner. *Nature Communications*, *12*(1).
- Mansfield, P. (1977). Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10(3), L55.
- Mezrich, R. (1995). A perspective on K-space. Radiology, 195(2), 297-315.
- Moser, E., Stahlberg, F., Ladd, M. E., & Trattnig, S. (2012). 7-T MR—from research to clinical applications? *NMR in Biomedicine*, 25(5), 695–716.

- Noll, D. C., Boada, F. E., & Eddy, W. F. (1997). A spectral approach to analyzing slice selection in planar imaging: Optimization for through-plane interpolation. *Magn. Reson. Med.*, 38(1), 151–160.
- OECD (2021). Magnetic resonance imaging (MRI) exams (indicator). OECD, 2021. doi: 10.1787/1d89353f-en.
- Pruessmann, K. P., Weiger, M., Scheidegger, B. M., & Boesiger, P. (1999). SENSE: sensitivity encoding for fast MRI. *Magn. Reson. Med.*, 42(5), 952–962.
- Purcell, E. M., Torrey, H. C., & Pound, R. V. (1946). Resonance absorption by nuclear magnetic moments in a solid. *Phys. Rev.*, 69, 37–38.
- Rabi, I. I., Ramsey, N. F., & Schwinger, J. (1954). Use of Rotating Coordinates in Magnetic Resonance Problems. *Rev. Mod. Phys.*, *26*, 167–171.
- Rabi, I. I., Zacharias, J. R., Millman, S., & Kusch, P. (1938). A new method of measuring nuclear magnetic moment. *Phys. Rev.*, *53*, 318–318.
- Redpath, T. (1998). Signal-to-noise ratio in MRI. British Journal of Radiology, 71(847), 704-707.
- Robson, P. M., Grant, A. K., Madhuranthakam, A. J., Lattanzi, R., Sodickson, D. K., & McKenzie, C. A. (2008). Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn. Reson. Med.*, 60(4), 895–907.
- Roguin, A. (2004). Nikola tesla: The man behind the magnetic field unit. J. Magn. Reson. Imaging, 19(3), 369–374.
- Skare, S., et al. (2018). A 1-minute full brain MR exam using a multicontrast EPI sequence. *Magn. Reson. Med.*, 79(6), 3045–3054.
- Tourais, J., Coletti, C., & Weingärtner, S. (2022). Chapter 1 Brief Introduction to MRI Physics. In M. Akçakaya, M. Doneva, & C. Prieto (Eds.) Magnetic Resonance Image Reconstruction, vol. 7 of Advances in Magnetic Resonance Technology and Applications, (pp. 3–36). Academic Press.
- Vaughan, T., et al. (2006). 9.4T human MRI: Preliminary results. *Magnetic Resonance in Medicine*, 56(6), 1274–1282.
- Weigel, M., Zaitsev, M., & Hennig, J. (2007). Inversion recovery prepared turbo spin echo sequences with reduced SAR using smooth transitions between pseudo steady states. *Magnetic Resonance in Medicine*, 57(3), 631–637.
- Zavoisky, E. K. (1945). Paramagnetic relaxation of liquid solutions for perpendicular fields. *Zhur. Eksperiment. i Theoret. Fiz.*, *15*(7), 344–350.
- Zeeman, P. (1897). On the influence of magnetism on the nature of the light emitted by a substance. *Philosophical Magazine, Series 5*, 43(262), 226–239.

3

The Advent of Quantitative MRI

Contents

3.1	Introduction	43
3.2	Basic concept of quantitative MRI:	44
3.3	Clinical interests of quantitative MRI	46
3.4	MR Relaxometry	48
3.5	Arterial Spin Labeling	60
3.6	Parameter estimation	76
References		82

3.1 Introduction

Since the early 1980s, when magnetic resonance imaging (MRI) first became clinically available, advances in magnetic resonance (MR) scanners and the development of tailored MR acquisition protocols have led to widespread availability of MRI in clinical practice. MRI offers remarkable soft tissue contrast and the benefit of non-ionizing radiation, making it a safe and highly valuable imaging modality for disease diagnosis and preoperative planning for a wide variety of clinical applications. To date, MRI is considered the gold standard imaging technique for diagnosis and monitoring of many neurological diseases.

Conventional (anatomical) brain MRI images are qualitative in nature, showing signal intensities that depend on many different factors, including not only the underlying biophysical tissue properties of interest, such as the longitudinal relaxation time (T_1), the transverse relaxation time (T_2), or the proton density, but also the MRI acquisition technique (pulse sequence) and the chosen acquisition parameters (e.g., flip angle, repetition time, echo time, inversion time, etc.). A trained neuroradiologist can interpret the relative intensity values (i.e., with respect to a control/reference region) in these images to define hyper- or hypo-intensive structures and to decide on any diagnosis. However, this makes conventional qualitative MRI inherently subjective, i.e. there is no absolute and objective quantification (or measurement) of biophysical tissue parameters, which increases the risk of an erroneous diagnosis.

When using MRI images for clinical diagnosis and prognosis of brain disorders, it is imperative that images be quantitative, reproducible (i.e. precise), indicative of tissue parameters, and independent of imaging sites or scanner vendors. In this way, biological changes in the disease states and their response to possible treatments can be carefully detected and progressively monitored. As such, *MRI can evolve from a process of picture-taking, where observations are made on the basis of unusually bright, dark, small or large structures, to a measurement process where a whole range of objective quantities can be tested to see whether they lie in a normal range and whether they have changed from the time of a previous examination (Tofts, 2004). The plethora of MRI techniques aimed at absolute quantification of biophysical tissue parameters is termed quantitative MRI, which will be abbreviated as 'qMRI' in the remainder of this thesis. As will be shown in this chapter, qMRI replaces qualitative images with <i>quantitative parameter maps*. While these quantitative parameter maps appear similar to a contrast-weighted MRI scan, they are conceptually different with voxel values having a biological meaning rather than representing signal intensity on an arbitrary scale.

In what follows, section 3.2 of this chapter briefly repeats the basic principles of quantitative MRI, illustrating its concept for quantitative T_1 mapping. Next, section 3.3 highlights some clinical interests of qMRI, so that the reader gets a clear understanding how qMRI can contribute to new impactful scientific insights. Subsequently, sections 3.4 and 3.5 provide a more in-depth explanation of two important qMRI applications relevant for study of neurological diseases, namely **MR relaxometry** and **Arterial Spin Labeling MRI**, aimed at quantitative measurements of relaxation and perfusion parameters, respectively. Particular attention is given to existing technical needs that prevent routine clinical use of these qMRI applications. Both qMRI applications will be further investigated in the contributions of this thesis. Finally, section 3.6 elaborates on the process of absolute quantification and statistical parameter estimation in the context of qMRI.

3.2 Basic concept of quantitative MRI:

qMRI is a technique for estimating biophysical properties of (brain) tissue, such as the T_1 and T_2 relaxation times, the proton density (PD), the mean diffusivity (MD), the apparent diffusion coefficient (ADC), the magnetic susceptibility, or perfusion measures such as the cerebral blood flow (CBF) and the arterial transit time (ATT). The estimation (or quantitative parameter mapping) of these properties typically consists of **two steps** (Tofts, 2004; Karakuzu et al., 2022):

 Collecting multiple MRI images, where the contributions of effective¹ micrometer-level MRI parameters, such as local T1 and T2 values, are systematically manipulated by adapting specific acquisition parameters.

¹For water protons, the nm-level quantum-mechanical couplings typically average out to produce *effective* μ m-level parameters, such as the local T1 and T2 values and the local diffusion coefficient (Novikov et al., 2018).

2. Voxel-wise fitting a (biophysical) signal model to the resultant voxel intensity variations across the images.

These steps result in one or multiple quantitative maps of the estimated tissue properties across the imaged volume. Because the tissue properties can be indicators of the biological state of the tissue and their change during disease, their precise and accurate estimation is of high importance. The number of acquired images in a qMRI images series is variable, with typically five to over a hundred images being acquired, depending on the tissue properties of interest. Fig. 3.1 illustrates the concept of qMRI for a T_1 relaxation signal model.



Figure 3.1: Schematic illustration of the concept of qMRI using T_1 mapping as a carrying example. On the top two rows a series of eight contrast-varying images of a brain is shown. Each image is characterized by a unique T_1 contrast-weighting as a result of a different acquisition parameter (in this example the inversion time (TI) within an inversion recovery sequence). The yellow dots indicate one specific pixel location \boldsymbol{x} in each of the images, for which the corresponding intensities are plotted in the graph below. The signal model $m(\vartheta)$ is fitted to these pixel intensities, where ϑ contains the tissue properties of interest (in this example the T_1 relaxation time). The model is fitted for each pixel in the image so that a map of the element(s) in ϑ can be created (in this example a T_1 map). Such a map is shown in the bottom right.

By providing quantitative measurement of tissue parameters, gMRI not only offers additional information for radiologists, but also provides an opportunity for improved harmonisation and calibration between scanners and as such it is well-suited to large-scale investigations such as clinical trials and longitudinal studies (Cashmore et al., 2021). Realising these benefits, however, also comes with its own challenges. When measuring a parameter quantitatively, it is crucial that the reliability² and reproducibility² of the technique are well understood. From a scientific point of view, a numerical result of a measurement is meaningless unless it is accompanied by a description of the associated measurement uncertainty. Therefore, parameter estimation and the monitoring of estimation uncertainties are an important part of qMRI research. As the process of determining a physical property from the raw MR signal is complicated and multistep, estimation of uncertainty is challenging and there are many aspects of the MRI process that require validation and research, most notably: signal model selection (i.e., selecting the appropriate model and parameter dependencies), motion and artefact correction, noise model selection (i.e., choosing a valid noise model for the data corresponding with the MR coil setup), and optimal experiment design (i.e., choosing the optimal number of contrast-weighted images and acquisition parameter settings).

3.3 Clinical interests of quantitative MRI

Early attempts at quantitative MRI date back to the late 70's (Gupta, 1977), and were primarily designed for the mapping of a single parameter, e.g. the T_1 or T_2^* relaxation time. Since then, the field has witnessed many waves of impactful developments, driven by technological advances and emerging trends in MRI research (Stikov et al., 2019). In addition to the technological developments associated with the roll-out of qMRI, it is useful to briefly highlight in which areas qMRI can be interesting from a clinical point of view:

- **Improved characterization of tissues and pathologies** The acquisition of PD, T_1 , or T_2 'maps' using qMRI can facilitate improved characterization of tissue, enhance image tissue contrast, and provide a more direct link between the observed signal changes and the micro-anatomical alterations distinguished via histochemistry and histology (Tofts, 2004; Cheng et al., 2012). Particularly, alterations between healthy and affected tissue can be detected locally with high specificity and sensitivity, the extent of tissue injury can be characterized, or the temporal evolution of both individual lesions and the overall disease activity can be monitored. All of the above can provide a greater understanding of the natural history of diseases and associated pathological changes, potentially paving the way towards effective treatment or therapy.
- **Earlier detection of pathophysiological changes** The detection of neurodegenerative diseases and cancer (NDAC) before the manifestation of clinical symptoms is paramount to prevent or delay its progression. Considering that prevention, delay, and treatment is more likely to be successful for patients in the earliest disease phases, it is important

²*Reliability* relates the magnitude of the measurement error in observed measurements to the inherent variability in the true underlying level of the quantity between subjects. High reliability means small measurement errors compared to true differences between subjects, enabling clear distinction between subjects based on error-prone measurements. Low reliability occurs when measurement errors are large compared to true differences, leading to potential confusion between genuine differences in true values and errors in measurements. *Reproducibility* refers to the consistency of measurements. It is the extent to which a measurement tool can produce the same result when used repeatedly under the same circumstances. For a more elaborate discussion of both concepts, the reader is referred to (Bartlett & Frost, 2008).

to discover reliable and accessible biomarkers that can detect NDAC prior to its clinical manifestation. However, current diagnostic biomarkers, including MRI markers, are often invasive and require specialized personal or expensive hardware. These constraints, together with financial and logistical issues limit broad-based implementation of these biomarkers for wide application for screening, in primary care settings. The growing responsibility of primary care physicians and care teams in the screening and early diagnosis of NDAC is imperative in this era of an aging population, combined with a shortage of specialists such as geriatricians (Rowe, 2021). Therefore, to promote the early detection of at-risk individuals, there is a need to identify accessible and scalable biomarkers of brain health that can be obtained regularly in the general population at point-of-care facilities. By establishing a unique relationship between MR parameter maps and physiology to provide a noninvasive surrogate for biopsy and histology, qMRI has attracted increased interest as a method to discover potential biomarkers for the detection of subtle or diffuse pathophysiological changes (Cheng et al., 2012). Particularly, various studies have shown that quantification of relaxation time variation can be important for the early diagnosis and progress monitoring of diseases in the human brain, e.g. in studies concerning autism (Deoni, 2011), dementia (Erkinjuntti et al., 1987; Tang et al., 2018; Knight et al., 2019), Parkinson's disease (Baudrexel et al., 2010; Vymazal et al., 1999), multiple sclerosis (Larsson et al., 1989; Stevenson et al., 2000; Parry et al., 2002; Gracien et al., 2016), epilepsy (Okujava et al., 2002; Townsend et al., 2004), stroke (Bernarding et al., 2000), and tumors (Just & Thelen, 1988; Badve et al., 2017; Chekhonin et al., 2023).

- **Objective quantification as a prerequisite of precision medicine** The increased emphasis on evidence-based and precision medicine requires physicians to integrate data from clinical examinations, laboratory tests, and imaging studies when deciding on patient care, and to assess and alter treatment plans as necessary. Data integration from multiple sources is becoming increasingly automated, and this requires that input data be inter-operable, machine-readable, and, ideally, quantitative. By providing clear numeric differentiation of disease states, qMRI increases the quality of information available to artificial intelligence algorithms for automated decision-making (Keenan et al., 2019). It is important to note that this comes with the requirement that parametric quantification using qMRI is reproducible and standardized, so to limit variability of quantitative values derived from radiological images (Keenan et al., 2019; Hagiwara et al., 2020).
- **Longitudinal follow-up and multi-centric evaluation** The quantitative measurement of tissue parameters without the confounding influence of other (hardware-specific) MR parameters allows to directly compare qMRI maps across subjects and in time, for proper follow-up study. As an example, in addition to providing early diagnosis, qMRI relaxometry is used to follow cartilage repair treatment (Matzat et al., 2013), or to detect muscle changes following acute muscle tear in soccer and rugby players, and predict the return-to-sport time for such injuries (Biglands et al., 2020). Moreover, quantitative results facilitate group comparisons in which data from multiple MRI centers or hospitals is combined and evaluated (Voelker et al., 2021). This is not possible for conventional MRI, which is prone to inter-site and intra-site variability of scans, even when the same sequence is used on different scanners with the same subject (Voelker et al., 2016).

3.4 MR Relaxometry

3.4.1 Introduction

In the context of quantitative MRI, MR relaxometry refers to the plethora of methods that are aimed at the quantitative measurement of the intrinsic T_1 , T_2 , or T_2^* relaxation times of a particular tissue (Deoni, 2010). As explained in the preceding sections of this chapter, relaxation times are unique per tissue type and reflect changes in tissue density or chemical composition. Therefore, MR relaxometry can add sensitivity to conventional MRI scans and detect abnormalities not observed with conventional MRI. In what follows, the biophysical basis of T_1 and T_2 relaxation in neuroimaging is briefly reflected on. Next, some relevant pulse sequences for the acquisition of T_1 or T_2 weighted images are provided, followed by a brief recap of the basic principle of MR relaxometry parameter mapping, as already introduced in section 3.2. Additionally, some of the main clinical applications of T_1 and T_2 relaxometry are highlighted. Finally, at the end of this section on MR relaxometry, some existing bottlenecks and issues in MR relaxometry are discussed.

Parametric signal models for quantitative MR relaxometry mapping, as used in this thesis, are explained in the respective contribution chapters (T_1 mapping: chapters 5-6, and T_2 mapping: chapter 6).

3.4.2 Biophysical basis of relaxation

Biological basis of \mathcal{T}_1 relaxation

Fluctuating magnetic fields, largely arising from the motion of molecules near magnetic moments, play a significant role in the recovery of T_1 . Consequently, the T_1 relaxation process is frequently linked to water mobility and structural density, indicative of water molecule binding. In the brain, T_1 has also been shown to be strongly correlated with myelin or macromolecular volume content, both in gray matter (Stüber et al., 2014), and white matter (Mezer et al., 2013). As a result, researchers often leverage T_1 contrast in brain studies, exploiting the fact that myelin causes white matter to exhibit a shorter T_1 compared to gray matter, thereby creating a distinct contrast.

Moreover, T_1 can change due to pathologies. For instance, edema around tumors or inflammatory acute multiple sclerosis (MS) lesions leads to an increase in T_1 (Brück et al., 1997). Chronic MS lesions also exhibit an increased T_1 , likely attributed to the reduction of myelin and an increase in water content. At the rim of active MS lesions, T_1 is reduced due to the presence of cell debris that forms extra-relaxation centers in the fluid. Other alterations, such as myelination of developing brain (Paus et al., 2001), or decrease in myelination due to aging can benefit from T_1 quantification (Cho et al., 1997). A comprehensive review of T_1 values in normal and pathological tissues over a range of field strengths can be found in (Bottomley et al., 1987; Bojorquez et al., 2017).

Biological basis of T_2 relaxation

In the brain, the degree of binding and water compartmentalization is reflected by T_2 . For example, research in premature human neonates previously revealed a decrease of T_2 as the brain underwent maturation (Ferrie et al., 1999). During brain maturation, tissue water decreases, myelin precursors such as cholesterol and proteins appear, glial cells proliferate and



Figure 3.2: A structural T_1 -weighted and T_2 -weighted MRI of a brain showing the white matter (WM), the gray matter (GM) and the cerebrospinal fluid (CSF).

differentiate, and a number of biochemical cell membrane changes occur. These alterations increase the ratio of bound to free water, consequently shortening the T_2 relaxation time (Dobbing & Sands, 1973). While T_1 relaxation times also decrease during brain maturation, especially with the onset of myelination, the decrease in T_2 occurs at a faster rate. Hence, T_2 has been proposed as a particularly interesting biomarker for assessing the developmental stages of the brain (Tofts, 2004).

A number of studies have demonstrated that gray matter has a longer T_2 relaxation time than white matter (Bojorquez et al., 2017). This distinction is believed to be due to differences in water compartmentalization, vascularity and iron concentration. The paramagnetic nature of iron shortens proton relaxation times. For instance, earlier research reported a shorter T_2 in areas such as substantia negra, in globus pallidus and putamen in Parkinson's patients (Vymazal et al., 1999).

3.4.3 Methods for T_1 and T_2 measurements

There are multiple MR pulse sequences that provide T_1 or T_2 contrast-weighted images (see Fig. 3.2). Some relevant sequences are listed below (Tsialios et al., 2017; Boudreau et al., 2020; Dortch, 2020):

\mathcal{T}_1 -weighted imaging

Inversion recovery based sequences The gold standard method to measure T_1 relies on inverting the longitudinal magnetization (by applying a 180° RF pulse), sampling its recovery at different time points (referred to as TI), and then fitting an exponential model of magnetization recovery to the data (see Fig. 3.1). This basic inversion recovery (IR) approach is known for its high accuracy and precision (Drain, 1949; Hahn, 1949). However, this approach is not used in a clinical setting because of its long total
acquisition time, since it requires a repetition time (TR) approximately five times longer than T_1 to allow complete recovery of the longitudinal magnetization. Nonetheless, IR is continuously used as a reference measurement during the development of new techniques, or when comparing different T_1 mapping techniques.

- **Look-Locker methods** The Look-Locker (LL) method (Look & Locker, 1970) improves the efficiency of the IR technique by sampling multiple TIs per repetition, using a train of low-angle RF readout pulses. The problem with LL is that it is particularly sensitive to RF pulse errors. The effect of the readout pulses is to hasten the recovery of the magnetization. This means that the magnetization recovers at an apparent T_1 , or T_1^* . The conversion between T_1^* and T_1 requires accurate knowledge of the flip angles, making the LL sequence sensitive to inhomogeneities in the applied RF field (B_1) (Kaptein et al., 1976). Furthermore, the small flip angles of the LL method result in a low SNR (Crawley & Henkelman, 1988).
- **Variable flip angle methods** The variable flip angle (VFA) method (Christensen et al., 1974), also known as driven equilibrium single pulse observation T_1 (DESPOT1) (Homer & Beevers, 1985), can generate a T_1 map from spoiled gradient echo (SPGR/FLASH) images at two or more flip angles with constant repetition time (Deoni et al., 2003, 2005). Whole-brain coverage can be achieved at high resolutions and reasonable scan times. Indeed, DESPOT1 enables the acquisition of a T_1 map with 1 mm³ resolution in approximately 7 minutes (Deoni et al., 2003). However, the precision and accuracy are low (~10%). Also, DESPOT1 suffers from a strong dependence on excitation flip angle (Deoni, 2007; Yarnykh, 2007), and is highly sensitive to proper SPGR sequence spoiling (Yarnykh, 2010), which may require the use of large gradients and increase the overall time of the technique.
- Dictionary-based methods Dictionary-based qMRI techniques use numerical dictionaries databases of pre-calculated signal values simulated for a wide range of tissue and protocol combinations - during the image reconstruction or post-processing stages. Notable examples of dictionary-based techniques are MR Fingerprinting (MRF) (Ma et al., 2013) and Magnetization Prepared 2 Rapid Acquisition Gradient Echoes (MP2RAGE) (Marques et al., 2010). MRF typically leverages information redundancy from parametric data to assist in accelerated imaging, while MP2RAGE uses dictionaries to estimate quantitative maps using the MR images after reconstruction. MP2RAGE, an extension of the conventional MPRAGE pulse sequence (Haase et al., 1989; Mugler III & Brookeman, 1990), is increasingly adopted as a standard pulse sequence for T_1 mapping on many MRI systems. It can be seen as a hybrid between the inversion recovery and VFA pulse sequences (Boudreau et al., 2020): a 180° inversion pulse is used to prepare the magnetization with T_1 sensitivity at the beginning of each TR, and then two images are acquired at different TIs using gradient recalled echo (GRE) imaging blocks with low flip angles and short repetition times. Because two images at different TI times are acquired, information about the T_1 values can be inferred, thus making it possible to generate quantitative T_1 maps using this data. Typically, MP2RAGE does not use a conventional minimization algorithms to fit a signal model to the observed data. Instead, to limit post-processing times, it uses pre-calculated signal values for a wide range of T_1 parameter values, and then interpolation is done within this dictionary of values to estimate the T_1 value that matches the observed signal.

Saturation pulse modified sequences The need to wait for a full recovery of the magnetization in inversion recovery (IR) methods may be circumvented by the introduction of a saturation pulse followed by a delay before each inversion pulse, combined with appropriate modifications of the T_1 fitting procedure (Deichmann et al., 1999). The saturation pulse ensures that the longitudinal magnetization is always in a fixed state when the inversion pulse is applied. A very fast and accurate T_1 mapping method, dubbed TAPIR (T_1 -mApping-with-Partial-Inversion-Recovery), is based on this approach (Shah et al., 2001). TAPIR combines the LL sequence with a pre-saturation scheme and an advanced multislice, multi-time point data acquisition protocol. Using TAPIR, the whole brain can be covered by acquiring 13 slices with 8 mm thickness, with 12 time points in 5 min 44 s and 1 mm in-plane resolution. TAPIR has been shown to deliver sub 1% accuracy (Shah et al., 2001). Unfortunately, the through-plane resolution is very low and increasing it while maintaining the in-plane resolution would unavoidably increase the acquisition time. For example, the total acquisition time for a T_1 map with volume coverage of 4 slices with a slice thickness of 4 mm is about 17 minutes (TR=12ms and 40 time points, FA=25°) (Möllenhoff et al., 2010).

T_2 -weighted imaging

- **Spin echo sequences** Gold standard methods to produce T_2 maps rely on spin echo (SE) measurements at different echo times (TEs) and fitting them to an exponential model of signal decay. Typically, a 90° excitation pulse is followed by a 180° refocusing pulse at time TE/2, and signal is measured at time TE where the spin-echo is formed and effects of inhomogeneity in the main magnetic field (B_0) are eliminated (cf. also section 2.4 of Chapter 2). T_2^* measurements are similar to T_2 , except that an 180° refocusing pulse is not used. Unfortunately, single SE acquisitions require extremely long scan times, extending on the order of tens of minutes. As an alternative, multi-echo spin echo (MESE) sequences have been proposed (Feinberg et al., 1985), which stimulate the spin system with repeated 180° pulses. As long as T_2 relaxation is not complete and MR signal is present, this allows to generate extra echoes within a given repetition time. The primary drawback of multi-echo methods is their use of multiple refocusing pulses, which results in signal contributions from non-spin-echo pathways (i.e., stimulated echoes) that can contaminate the signal decay and bias T_2 estimates, even in the presence of relatively minor B_0 and B_1 imperfections (Dortch, 2020).
- Accelerated T_2 mapping More recently, to accelerate T_2 quantification and subsequent generation of synthetic T_2 -weighted image contrast for clinical research and routine, model-based approaches for rapid T_2 and proton density (PD) quantification have been developed. For instance, a technique called GRAPPATINI has been proposed, in which a model-based approach for high-speed T_2 and PD quantification based on *k*-space subsampling (*Model-based Accelerated Relaxometry by Iterative Non-linear Inversion*, MARTINI) (Sumpf et al., 2011), was complemented by parallel imaging (*generalized autocalibrating partially parallel acquisition*, GRAPPA) (Griswold et al., 2002) to provide a high-resolution T_2 mapping of the whole brain within 1:44 min (Hilbert et al., 2018). Moreover, this approach also provides synthetic T_2 -weighted images with different echo times at no additional acquisition time (Hilbert et al., 2018). These synthetic images have already been evaluated for pediatric (Kerleroux et al., 2019) and musculoskeletal (Roux et al., 2019) applications, and recently also in brain

imaging (Gruenebach et al., 2023). A downside of GRAPPATINI is that it introduces additional noise in the resulting T_2 maps, due to the use of GRAPPA undersampling. In particular, it is well known that the SNR of parallel imaging acquisitions scales with $1/g \times \frac{1}{\sqrt{AF}}$ (where AF denotes the acceleration factor, and g being impacted by coil design) when undersampling an MRI dataset (Breuer et al., 2009). When not considering any effect of regularization during the reconstruction, a loss of 30% in SNR was experimentally observed when adding the additional GRAPPA (Hilbert et al., 2018).

Combined T_1 and T_2 -weighted imaging

Apart from the sequences above that are tailored to provide T_1 -weighted or T_2 -weighted contrast separately, some techniques exist that can generate multiple contrast-weighted images simultaneously. For instance, **inversion recovery TrueFISP (Fast Imaging with Steady Precession)** acquires T_1 -weighted, T_2 -weighted, and proton density-weighted images together (Schmitt et al., 2004). It uses an inversion pulse to adjust T_1 contrast, utilizes the steady-state free precession (SSFP) sequence for inherent T_2 contrast sensitivity, and captures proton density-weighted information from steady-state magnetization, all in one MRI sequence. As such, this method enhances efficiency by reducing scan time while also improving consistency across different tissue contrast images in MRI exams.

3.4.4 Quantitative relaxometry mapping

The objective of quantitative MR relaxometry mapping is to estimate a relaxation time at different spatial positions of the scanned object, thereby generating a spatial map of the relaxation parameter of interest.

To achieve this, three crucial ingredients have to be present:

- 1. **Pulse sequence and signal model selection** First, a specific MR pulse sequence has to be selected. The applied pulse sequence disturbs the net magnetization vector out of equilibrium, inducing the spin ensemble system to enter the relaxation phase (see Chapter 2). Here, different pulse sequences result in different ways of magnetization recovery, which in turn obey different (relaxation) signal models (Barral et al., 2010; Bojorquez et al., 2017). These models, which typically involve nonlinear exponential decay functions, depend on the relaxation time and also on user-defined acquisition parameters such as specific timings (echo time, inversion time, repetition time) or flip angles.
- 2. Signal variation over time A set of contrast-weighted images needs to be created. As the relaxation times needs to be probed at different spatial positions, spatial encoding of the images is pivotal. Hence, *k*-space data is acquired as described in Chapter 2, and the corresponding magnitude MR images are reconstructed. Crucial in the acquisition process is that the different images in a set are obtained using different time points or flip angles, so that the contrast-weighting changes per image (i.e., different signal intensities per image), and the signal model of interest can be optimally sampled over time. See also the schematic illustration in Fig. 3.1, demonstrating this concept for T_1 mapping as a carrying example.

3. Voxel-wise model fitting – Finally, the relaxation time parameters, which parameterize the signal change over time, are estimated from the set of contrast-weighted images by voxel-wise fitting the selected signal model to the acquired data points. It is crucial that accuracy and precision of the estimation framework are carefully monitored. Once the quantitative relaxation parameter maps are estimated, they can be used to assess tissue properties and potentially diagnose or characterize diseases, as will be highlighted in the next section.

3.4.5 Clinical applications

In addition to the list of clinical interests of qMRI introduced in section 3.3, the current section briefly highlights some clinical applications of MR relaxometry, focusing on neuroimaging in particular (Deoni, 2010).

- Multiple sclerosis In the study of patients with multiple sclerosis (MS), quantitative MR relaxometry has shown to be an invaluable tool for studying changes in myelin and iron content in the brain and spinal cord. MS is a disabling neurodegenerative and neuroinflammatory disease affecting over 400.000 persons in Europe. MS is hallmarked by characteristic hypo (T_1) and hyper (T_2) intense macroscopic lesions on spin echo MRI images as well as changes in diffusion characteristics caused by destruction of myelin and axons. Disease modifying drugs have become available, but the success of these drugs relies on early diagnosis and adequate tools for therapy monitoring (Noyes & Weinstock-Guttman, 2013). In recent years, MRI-derived biomarkers for diagnosis of MS in clinical practice have been launched (Jain et al., 2015, 2016; Sima et al., 2017). These biomarkers rely on analysis of weighted MRI images to extract for example the (total) volume of T_2 lesions. A limitation of such volumetric biomarkers is the necessity of macroscopic lesions, whereas the underlying pathology first manifests itself at the microscopic level. Volumetric measures intrinsically only detect (irreversible) tissue-damage, typically occurring in a late stage when treatment options are limited. Increasing evidence suggests that MR relaxometry has the ability to detect subtle microscopic tissue damage during early neurodegeneration. For instance, in MS both normally appearing gray and white matter have shown prolonged relaxation times (Vrenken et al., 2006; Roosendaal et al., 2009; Papadopoulos et al., 2010). These observations support the hypothesis that MR relaxometry could lead to sensitive early biomarkers of MS. By facilitating the accessibility to quantitative MR sequences on novel MRI scanners and by improving the robustness of relaxation parameter mapping methods, qMRI will likely play a fundamental role in the upcoming decades as a sensitive tool to quantitatively assess brain damage in patients with MS, with relevant implications for prognostic stratification and treatment-response evaluation (Tranfa et al., 2022).
- **Epilepsy** Epilepsy is one of the most common neurological diseases. Epileptic seizures may occur as a result of intercurrent events such as fever, hypoglycemia (low blood sugar levels), acute central nervous system infections and the like, and are then termed *occasional seizures*. When they recur spontaneously without known cause, they constitute epilepsy (Achten, 2001). It is a chronic condition in which occasional seizures tend to occur repeatedly as a result of either structural brain damage or of an intrinsic functional propensity to have seizures. The most common cause of temporal lobe epilepsy is hippocampal sclerosis (HS), or atrophy of the hippocampus

(Winston et al., 2017). Although HS is typically associated with increased signal intensity in T_2 -weighted images, the ambiguity of conventional T_2 -weighted signal changes hinders definitive diagnosis (Reutens et al., 1996). At this point, quantitative T_2 mapping in normal and pathological hippocampal tissue allows for a more objective and improved effectiveness of the detection and monitoring of hippocampal structure changes (Jackson et al., 1993; Rodionov et al., 2015).

Dementia and Alzheimer's disease Alzheimer's disease (AD) is an irreversible neurodegenerative disorder affecting millions of people each year. AD is characterized by a progressive accumulation of two toxic proteins, the amyloid-beta and hyperphosphorylated tau, inducing neuronal loss, cognitive impairment, and dementia (van den Berg et al., 2022). Due to an aging society, an increasing number of people are expected to suffer from this disorder, which entails enormous economic and social burdens to society. Currently, more than 55 million people are living with AD or related dementia's worldwide, and its prevalence is expected to grow exponentially over the next few decades (World Health Organization et al., 2017). Furthermore, AD is the cause of approximately 70% of all dementia's, and dementia is currently the seventh leading cause of death and one of the major causes of disability and dependency among older people globally³. In low- and middle-income countries, which have a higher population growth rate, dementia will pose an even more severe threat to their development in the near future (World Health Organization et al., 2017). The early detection of AD before the manifestation of clinical symptoms is paramount to prevent or delay the progression of the disease by targeting new interventions, including modification of risk factors, enrolling in clinical trials, or using disease-modifying drug therapies. Unfortunately, AD has an insidious onset that makes early diagnosis challenging (Coupé et al., 2015; Scheltens et al., 2016; Villemagne & Chételat, 2016). Research is therefore increasingly focusing on biomarkers for early diagnosis, disease progression monitoring and potential treatment response predictions. Neuroimaging biomarkers play a crucial role in this field, including regional structural alterations on structural MRI, metabolism alterations on Positron Emission Tomography (PET), detection of amyloid plaque deposits on amyloid PET, and brain function alterations on blood oxygenation level dependent (BOLD) functional MRI (Dustin et al., 2016; Villemagne & Chételat, 2016). However, these methods are invasive, not widely accessible, very expensive, and show high variability that challenges the interpretation of longitudinal studies. These limitations impede the large-scale implementation of such biomarkers. In contrast, quantitative T_1 and T_2 relaxation times have been proposed to serve as non-invasive biomarkers of AD, in which alterations are believed to not only reflect AD-related neuropathology but also cognitive impairment (Tang et al., 2018). Further, accumulating evidence suggests that AD pathology affects biological properties of white matter beginning from pre-symptomatic stages of AD, where AD disrupts integral white matter properties involving axonal transport and packaging, axonal density, axonal tract myelination, and macromolecular lipid composition (Fingerhut et al., 2022). It has been shown that quantitative relaxometry is able to characterize specific biological properties of white matter tissue in vivo (Gozdas et al., 2021), and as such provide biomarkers for AD non-invasively and longitudinally. In addition, therapy monitoring of recent

 $^{^{3}}$ In Belgium, the number of people with dementia was estimated at 192,926 in 2018 (1.69% of total population), with an expected growth to 210,974 in 2025 (1.79% of total population) (Alzheimer Europe, 2019).

FDA approved drugs for AD such as aducanumab (or Aduhelm), an amyloid-beta directed antibody, requires regular examination leveraging MRI, with current guidelines requiring up to four MRI sessions in between drug infusions (Cummings et al., 2022). To get the most information out of such routine MRI scans and to follow the disease development longitudinally, qMRI and MR relaxometry are expected to gain importance in the coming years. Moreover, such regular MRI follow-up examinations also create a clear technological need for regular, fast, and accessible qMRI scans so that patients with AD can be treated in a timely manner in an aging society.

Neurodevelopment Another area of increasing clinical interest is brain development in early infancy. A possible explanation for a variety of psychiatric disorders, including autism, developmental delay, and attention deficit, is disrupted or abnormal connectivity of the complex neurological systems that underlay higher order emotional, social or behavioural functions (Hughes, 2007). Mediating this connectivity are the myelinated white matter pathways of the brain, which develop throughout the first years of life. Quantitatively monitoring the maturation of these pathways using MR relaxometry parameter mapping, in association with behaviour monitoring, may offer new insights into the spatial and temporal origins of these disorders. Particularly, it has been demonstrated that conventional T_1 and T_2 weighted MRI brain scans over the first life year highlight progressive changes in white and grey matter contrast (Ballesteros et al., 1993; Huang et al., 2006). It is believed that the observed decrease in both T_1 and T_2 throughout the first years reflect to the increased presence of lipids, cholesterol and other constituents of the myelin sheath, as well as an increased water compartmentalization of the brain. However, hardware dependent signal profile inhomogeneities inherent to qualitative weighted MRI scans make appreciation and comparison of tissue signal and contrast changes difficult and ambiguous. Quantitative evaluation of T_1 and T_2 relaxation times throughout neurodevelopment can provide a less ambiguous appreciation of age related change and maturation (Deoni et al., 2011).

3.4.6 Existing problems in MR relaxometry

The potential of MR relaxometry is undisputed, however its application remains subject to a number of difficulties, largely attributable to the need to include multiple contrast-weighted scans to provide a voxel-wise fit of the chosen signal model. Some of these existing problems are explained below.

3.4.6.1 Motion

It goes without saying that accurate and precise estimation of a relaxation parameter map via voxel-wise fitting of the relevant signal model to the data is only possible if the different contrast-weighted scans are spatially aligned and assume an **anatomical correspondence between the scans**. However, due to unavoidable patient motion, physiological motion such as cardiac or respiratory motion, and/or geometric distortions caused by the acquisition, small misalignment will occur at the voxel level. This may lead to erroneous parameter estimation, especially at tissue boundaries. Misalignment can be reduced during acquisition by using methods such as gating or breath-holding (van Heeswijk et al., 2012). However, such approaches do not always give the desired effect and can increase the acquisition time. The most common solution is to spatially align the images prior to the voxel-wise signal model fitting. This alignment can be achieved with image registration techniques. However, **image** registration for MR relaxometry (and qMRI) imposes three main challenges:

- **Contrast variation** A typical data set for MR relaxometry consists of scans with different contrast-weightings, complicating image registration based on the image intensities. To deal with these contrast changes a pairwise registration approach is commonly used in which all images are registered to a chosen reference image using a metric based on mutual information⁴ (MI), because this metric is robust against intensity changes in the images (Bron et al., 2013; Mangin et al., 2002). However, it has been shown that the choice of reference image can influence the result of the registration (Huizinga et al., 2016).
- **Multi-image registration** Often more than two images need to be co-registered. To circumvent the need to choose a reference image and perform a pairwise registration, group-wise registration approaches have been proposed that simultaneously register all images to a mean reference space. This offers the advantage that information of all images is taken into account during the registration process, which improves consistency compared over a pairwise registration approach (Bhatia et al., 2007; Metz et al., 2011; Wachinger & Navab, 2013; Hallack et al., 2014; Huizinga et al., 2016).
- **Propagating errors** Many qMRI estimation routines compensate for motion by using a separate processing step in which the motion parameters of the individual images are updated once after registration (Studler et al., 2010; Bron et al., 2013; Guyader et al., 2015; Van Steenkiste et al., 2016, 2017), prior to the voxel-wise fitting of the signal model. A downside to such an approach is the *lack of a feedback mechanism* that connects the motion compensation routine with the final estimation of the biophysical parameters. As a result, potential propagating registration errors will not be corrected for and may lead to inaccurate (i.e., biased) tissue parameter maps. As an example, the unfavorable effects of motion correction interpolation in model blind registration techniques have been discussed for quantitative T_1 mapping, showing substantial errors in the T_1 estimation (Nachmani et al., 2019). A possible retrospective⁵ motion

⁴Mutual information, which is a measure originating from information theory (Maes et al., 1997; Wells et al., 1996), uses entropy as its underlying concept. The entropy of an image can be thought of as a measure of dispersion in the distribution of the image gray values. Given two images A and B, the definition of the mutual information MI(A, B) of these images is MI(A, B) = E(A) + E(B) - E(A, B) with E(A) and E(B) the entropies of the images A and B, respectively, and E(A, B) their joint entropy. The joint entropy E(A, B) measures the dispersion of the joint probability distribution p(a, b): the probability of the occurrence of gray value a in image A and gray value b in image B (at the same position), for all a and b in the overlapping part of A and B. The joint probability distribution should have fewer and sharper peaks when the images are matched than for any case of misalignment. Therefore maximization of mutual information should correspond to the optimal affine transformation describing the motion between both images.

⁵Different motion compensation methods have been proposed in literature, which can be classified into three groups: motion correction based on *k*-space trajectories, prospective motion compensation techniques, and retrospective motion correction. Motion correction based on *k*-space trajectories relies on specially designed and implemented trajectories (Bookwalter et al., 2010; Liu et al., 2004), which limits the flexibility of these techniques and often increases the inherent acquisition time. Prospective motion compensation is achieved by obtaining real-time tracking data of the position and orientation of the object. This tracking data is then passed to the scanner with minimal delay, to allow adaptive adjustments of the MR pulse sequences such that the imaging volume follows the object movements (Maclaren et al., 2013; Callaghan et al., 2015). Retrospective motion correction methods modify the MR image data during reconstruction, which typically requires more complex mathematical considerations and understanding of the MR physics at hand (Loktyushin et al., 2013; Anderson III et al., 2013). A more elaborate overview of existing motion correction methods for MRI of the brain is given in (Godenschweger et al., 2016).

correction strategy to avoid such propagating errors, is to introduce an explicit object motion model in a model-based estimation framework, and estimate the corresponding motion parameters of this model jointly with the tissue relaxation parameters. In particular, it was demonstrated that by combining models of T_1 relaxation, motion, and noise into one unified statistical framework, one is able to obtain substantially more accurate motion and tissue parameter estimates, as compared to a conventional twostep approach in which motion registration precedes the T_1 estimation (Ramos-Llordén et al., 2017). This idea will prove useful later in this thesis, when motion compensation is to be combined with model-based super-resolution reconstruction for qMRI.

3.4.6.2 Specific absorption rate

The specific absorption rate (SAR) is a measure of the absorption of electromagnetic energy in the body (in watts per kilogram [W/kg]). The SAR describes the potential for heating of the patient's tissue due to the application of the RF energy to produce the MR signal. It is proportional to the square of the static magnetic field amplitude (B_0), the square of the flip angle α and the duty cycle D, which corresponds to the fraction of duration of the sequence during which the RF waves are transmitted (Bottomley et al., 1985; Bernstein et al., 2004):

$$SAR \propto B_0^2 \alpha^2 D. \tag{3.1}$$

For some MR pulse sequences, such as the TSE sequences (cf. section 2.4.1 of the previous chapter) which are typically used to accelerate the readout for T1-mapping (IR-TSE), the echo train of the RF pulses deposits a high amount of RF energy, resulting in high SAR values (Oshio & Feinberg, 1991; Weigel et al., 2007). Consequently, to allow practical use, SAR is often reduced by decreasing either the number of slices acquired within one TR, or by acquiring thicker slices (thus reducing spatial resolution), or by increasing the number of excitations (thus increasing the scan time). Once again, this highlights the need for qMRI approaches that can balance imaging trade-offs, in this case between spatial resolution, scan time, and SAR.

3.4.6.3 Acquisition time constraints

An inherent disadvantage of qMRI is the need to record multiple contrast-weighted images. The total scan time of a qMRI acquisition is as such proportional to the number of acquired images and the scan time per image. Extending the scan time in MRI examinations should be avoided, as it increases costs and reduces patient comfort. Moreover, it increases the likelihood of motion artefacts during the scans which can be detrimental for accurate diagnosis.

As an example, consider a single T_1 -weighted image acquired with IR TSE, for which the total scan time T_{scan} is equal to (Bernstein et al., 2004):

$$T_{\rm scan} = \frac{{\rm TR} \times N_{\rm PE} \times N_{\rm EX}}{{\rm ETL}},$$
(3.2)

with N_{PE} the number of phase encoding lines, ETL the echo train length and N_{EX} the number of times the sequence has to be run to have full slice coverage, i.e. the total number of slices divided by the number of slices acquired per TR. The ETL in Eq. (3.2) is equal to one for a standard IR spin echo. Since the IR sequence demands long TR to allow complete recovery of the longitudinal magnetization, it is inherently slow. As already discussed in the previous chapter, one can reduce scan time to some extent by using fast readout approaches such as TSE and single-shot EPI. However, these respective methods come at the cost of either a significant increase in SAR (Weigel et al., 2007) or an increased sensitivity to geometric distortions due to the low bandwidth in the phase-encoding direction and B_0 -field inhomogeneity. The latter can be reduced to some extent by using a segmented simultaneous multi-slice acquisition, combined with slice order shifting across multiple acquisitions (Sanchez Panchuelo et al., 2021). Often, though, the total scan time is reduced either by reducing the number of T_1 -weighted acquisitions or by reducing the number of phase encoding lines N_{PE} and slices in a scan. However, this comes at the expense of a reduction of the precision or spatial resolution of the T_1 parameter mapping, respectively.

3.4.6.4 Slice profile variation

As explained in section 2.4.1 of the previous chapter, a slice encoding gradient is applied in 2D readout to provide spatial encoding along the z-axis. Ideally, 2D imaging would be performed using a series of contiguous thin slices where the corresponding slice-selective RF pulse should be able to excite a narrow slab by using a perfect rectangular slice profile, also known as a boxcar function (see Fig. 3.3). However, the Fourier transform of a boxcar function is a sinc function with infinite support. Since hardware components can only store a finite number of values, actual sinc RF pulses must be truncated (see Fig. 3.3). A truncated RF pulse will cause some parts of the slice, such as the center, to be excited as desired, but other areas, such as the edges of the slice, will be excited either more or less than expected. In addition, there may be excitation outside the slice. Consequently, variations in the flip angles across slices or B_1 field inhomogeneities can occur (Kingsley et al., 1998; Dowell & Tofts, 2007). In addition, imperfect slice profiles may lead to cross-talk between adjacent slices in a 2D multi-slice readout (Bernstein et al., 2004), i.e. there might be a loss of MR signal in one slice due to pre-excitation of an RF pulse meant for an adjacent slice. To prevent this from happening, 2D multi-slice images are often acquired with an inter-slice gap. However, as some parts of the subject are not fully sampled, inter-slice gaps lead to a loss of information.

3.4.6.5 Need for protocol standardization and metrology

Quantitative MR relaxometry necessitates the **standardization of acquisition protocols** to ensure consistent and comparable results across different scanners and institutions (Hagiwara et al., 2020). This standardization is crucial because vendor-specific differences can lead to variations in the resulting images, even when the same sequence is applied. Different MRI manufacturers, such as Siemens, General Electric, and Philips, have unique methods for calculating and setting parameters, which can result in significant discrepancies in the acquired MRI scans, affecting the accuracy and reliability of quantitative measurements. Establishing standardized protocols and calibration procedures is therefore essential to minimize these variations, enabling more accurate comparisons and reproducibility of quantitative MRI data across different platforms and clinical settings. As an example, standardizing image acquisition protocols across scanners before commencing a multicenter study has been shown to be a valuable tool to increase the statistical power and reduce the required sample sizes for detecting disease-related neuroanatomical changes (George et al., 2020).



Figure 3.3: (a) Slice profile of a rectangular excitation RF pulse centered at frequency ω_{sl} , with an excitation bandwidth $\pm \omega_{sl}$, and (b) its corresponding temporal profile, an infinitely long sinc pulse. (d) Example of a time-limited RF pulse, obtained by truncating a sinc pulse with a rectangular window. This causes ripples at the edge of the slice profile (so-called Gibbs ringing) (c). This figure was adapted from the work of Tourais et al. (2022).

On top of that, it is crucial that **standardization of reporting formats and analysis methods** is established to avoid incomplete or inaccurate reporting of parameters. The latter would complicate quantification, analysis, and sharing of data, particularly for studies across multiple sites, platforms, and methods. Standardization can occur in different ways, e.g. by establishing community consensus on (open-source) data standards such as the Brain Imaging Data structure (BIDS), for the organization of data and metadata for particular neuroimaging modalities. As an example, qMRI-BIDS has been proposed as an extension to the BIDS specifically aimed at quantitative magnetic resonance imaging data (Karakuzu et al., 2022), thereby reducing the entrance barrier for qMRI in the field of neuroimaging. Also for other modalities such as Arterial Spin Labeling (cf. Chapter 7), a specific ASL-BIDS data structure extension has been released (Clement et al., 2022b). Such BIDS extensions can act as a catalyst of convergence between qMRI method development and application-driven neuroimaging studies to facilitate the development of quantitative biomarkers.

With the many advantages that MR relaxometry offers and with the emergence of a multitude of new qMRI techniques, it remains important for MRI practitioners to consider the reliability and reproducibility of such techniques (Keenan et al., 2019; Cashmore et al., 2021). Measuring a quantitative parameter is one thing, but estimating the measurement uncertainty this entails is just as important. As the process of determining a physical property from the raw MR signal is complicated and multi-step, estimation of uncertainty is challenging and there are many aspects of the MRI process that require validation. For that reason, there is a clear and urgent need for **metrology** in qMRI and health care in general (Smith et al., 2020). Metrology is the study of measurement processes. A key idea of metrology is *traceability*, i.e. the chain of comparisons which directly relates any given measurement to the primary standard determination of that unit (e.g. metre, second, etc.) (Cashmore et al., 2021). Without understanding the traceability of a quantitative measurement, there is no way of making a meaningful comparison between values. Furthermore, without an evaluation of the

measurement uncertainty, there exists no way of knowing whether a quantitative difference observed with gMRI is significant. This is why, alongside the development of gMRI techniques, a requirement exists to develop new procedures and methods that perform independent validation of qMRI techniques to determine their accuracy, repeatability and reproducibility. For example, specific phantoms from the U.S. National Institute of Standards and Technology (NIST), that model human physiology, are being developed to calibrate (g)MRI devices and techniques (Gunter et al., 2009; Keenan et al., 2016, 2019). International consortia such as the Quantitative Imaging Biomarkers Alliance (QIBA) metrology groups (Sullivan et al., 2015; RSNA, 2020) and the European Imaging Biomarkers Alliance (EIBALL) (ESR, 2020) are being launched to establish standards on which to base quantitative image-based measurements, with collaborations between industry, research institutions and healthcare (Cashmore et al., 2021). Furthermore, inter-site comparisons are critical to determine how accurately image-based biomarkers can be measured. For example, a multi-site study comparing MRI T_1 measurements showed considerable variation and vendor-dependent bias using a gold standard inversion recovery protocol (Keenan et al., 2021), clearly demonstrating the need to define uncertainty intervals in image-based biomarker measurements obtained using MR relaxometry.

3.5 Arterial Spin Labeling

3.5.1 Introduction

Our brains are made up of a collection of neurons and glial cells, with a population of some 100 billions each (Herculano-Houzel, 2009). The cell bodies of the neurons are mostly located at the periphery of the brain or in particular locations, and form the *gray matter*, whereas the so-called *white matter* consists of the neuron's axons, which allow the exchange of information and communication between different areas of the gray matter. (see Fig. 3.2). Interestingly, the brain lacks almost any form of energy storage and therefore all energy needs to be transported into the brain by means of **cerebral perfusion** (Raichle, 2006), i.e. the biological process during which the different brain cells are supplied with oxygen and nutrients through the blood. It occurs largely via the microcirculation blood flow through the capillaries, the smallest vessels within the brain blood supply system.

An important parameter for quantifying the cerebral perfusion process is the **cerebral blood flow (CBF)**, which denotes the volume of blood delivered to a certain brain tissue volume within a certain amount of time, typically expressed in units of millilitres of blood per 100 grams of brain tissue per minute (mL/100g/min). It is a physiological parameter and potential biomarker of high interest in a number of brain disorders, such as stroke, neurodegenerative diseases, epilepsy, and cancer (Alsop et al., 2015; van Osch et al., 2018). Typical CBF values range from 50 to 70 mL/100g/min in cortical gray matter and about 20 mL/100g/min in white matter in young and healthy adults (Parkes et al., 2004).

There exist several methods that try to monitor cerebral perfusion and CBF, yet often with distinct disadvantages. Cerebral perfusion can be measured with ¹⁵O PET (Herscovitch et al., 1983), xenon-computed tomography (Gur et al., 1982) and CT perfusion, but these techniques involve ionizing radiation. Dynamic susceptibility contrast (DSC) MRI is widely used in clinical routine, where perfusion is visualized by injecting a bolus of gadolinium chelate contrast agent and by subsequently imaging it as it passes through the cerebral capillary bed

(Villringer et al., 1988; Belliveau et al., 1990; Rempp et al., 1994; Østergaard et al., 1996b,a). However, the major drawbacks of DSC MRI are the invasiveness of the contrast injection and the difficulty to attain absolute quantification of perfusion (Calamante et al., 2002). Intravoxel incoherent motion (IVIM) is another method that postulates the measurement of tissue perfusion by modeling it as a pseudo-diffusion process by exploiting motion sensitizing gradients (Le Bihan et al., 1986). Yet, the use of IVIM in the brain is challenging, mainly due to low cerebral blood volume fractions and modeling issues (Jezzard et al., 2018). Finally, ¹⁷O-water MRI is capable of absolute quantification of perfusion (Pekar et al., 1991), but the high cost of the label is a stumbling block for large-scale use.

The above list is not exhaustive. An alternative MRI method that allows absolute CBF quantification while being completely non-invasive, is Arterial Spin Labeling (ASL) MRI. ASL is based upon the use of blood as an endogenous tracer by employing spatially selective labeling of the inflowing blood that inverts its longitudinal magnetization. The technique was originally introduced in the early 1990s to assess rat brain perfusion (Williams et al., 1992; Detre et al., 1992). In the next 25 to 30 years, however, thanks to the many developments in MRI hardware, sequence optimization, post-processing and interpretation, ASL expanded into a full-fledged technique for human brain perfusion imaging (Jezzard et al., 2018). Yet, clinical adoption of ASL was hindered by the large amount of implementation options, both in terms of signal generation and perfusion quantification. This changed with the publication in 2014 (early view publication; printed version available early 2015) of the consensus paper by the ISMRM Perfusion Study Group and the EU-COST action 'ASL in dementia' on the recommended clinical implementation of ASL perfusion MRI (Alsop et al., 2015). The publication of this consensus paper was instrumental in the adoption of ASL brain imaging in the clinic and provided a common reference for researchers. Moreover, it provided expert guidelines for ASL sequence implementation for the major MR manufacturers, who since then all offer the same labeling strategy (pseudo-continuous ASL) and similar readouts (3D gradient-and-spin-echo (GRASE)). Consequently, clinical applications of ASL have increased significantly, and a benchmark for comparison of future developments was established.

Since 2015, a range of new technical developments and advances in ASL MRI have been developed, which are regularly reviewed by the ISMRM Perfusion Study Group (Hernandez-Garcia et al., 2022). In line with these new developments, a specific contribution of this thesis was dedicated to the combination of model-based super-resolution reconstruction with an ASL signal model for direct quantitative mapping of CBF. As a background to the respective contribution chapter 7, ASL's basic principle, its current recommended implementation, and a number of important technological concepts are explained in sections 3.5.2-3.5.3. Furthermore, a brief list of some recent clinical applications of ASL is given in section 3.5.4, and some of of the remaining challenges in today's ASL perfusion MRI landscape are highlighted in section 3.5.5.

3.5.2 Basic principle

ASL is a non-invasive magnetic resonance perfusion imaging technique that allows the quantification of cerebral blood flow throughout the vascular system of the brain. ASL uses arterial blood as an endogenous tracer. Before diving into the detailed facets of ASL signal generation, labeling, and readout schemes, it pays to define the core concept of ASL and to identify the various steps involved in the ASL imaging protocol.

Fig. 3.4 provides a schematic overview of the basic principle behind ASL imaging. First, RF pulses are applied in the large neck vessels in a plane at the level of the carotid arteries, to invert the magnetization of the hydrogen nuclei of arterial blood water. At this point, the inflowing blood is magnetically labeled. Subsequently, this label will continue to flow, and if one makes a scan at different times, the labeled blood can be seen to move through the vascular tree of the brain. This angiographic measurement of the supplying arteries provides a great insight into how the brain is supplied with blood, but it does not explain exactly how much blood flows to the brain tissue. Interestingly, the latter can be measured by waiting a little longer before acquisition of a scan after the labeling of the blood. During this waiting time of about 2 seconds, which is referred to as the **post-labeling delay (PLD)**, the label flows to the capillaries of the brain, where the labeled water will flow out of the vascular bed and into the brain tissue to accumulate there. After this PLD, a fast readout imaging module is employed to map the brain magnetization. This results in a so-called label image. Besides a label image, also a **control image** is acquired that is identical to the label condition except for the absence of inversion of inflowing blood. After subtraction of the label image from the control image, an image of the labeled blood that has reached the brain tissue is obtained. Because signal differences due to labeled blood water account for at most 5% of the raw ASL image intensity, the signal-to-noise ratio (SNR) of a single control-label subtraction is usually low (Mehranian et al., 2020; Clement et al., 2022a). In order to increase this SNR, multiple interleaved control-label image pairs are acquired, followed by pair-wise subtraction and subsequent averaging of the resulting subtraction images. The averaged and unitless subtraction image (also referred to as perfusion-weighted image) can be further quantified to obtain a **CBF-map** when the temporal width of the bolus of labeled spins, the labeling efficiency, and \mathcal{T}_1 of tissue are taken into account, and by correcting for the decay of label due to longitudinal T_1 relaxation.

Its non-invasive and quantitative nature makes ASL especially attractive for vulnerable patient populations, such as the elderly, oncological patients with difficult venous access, and patients with renal insufficiency (Grade et al., 2015). ASL is also favourable for pediatric populations, as it avoids the technical difficulties and ethical problems of contrast agents and radiation exposure with CT and nuclear medicine techniques (Wang et al., 2003). Moreover, ASL is repeatable and reproducible, which makes it a suitable quantitative imaging modality for longitudinal evaluation and monitoring of CBF changes and disease progression (Wolf & Detre, 2007; Wang et al., 2011; Mutsaerts et al., 2014). For example, ASL is often used to follow dynamic changes in CBF during functional challenges, such as a CO_2 inhalation test, which is used to assess cerebrovascular reactivity (Tancredi et al., 2012).

The main drawback of ASL is the low SNR. This increases the total necessary scan time, making the technique particularly sensitive to motion artefacts (Petersen et al., 2006). In addition, flow quantification can be complex, as the signal is dependent on a number of physiological parameters (Petersen et al., 2006; Bladt et al., 2020a). Also, post-acquisition processing of ASL images and their preparation for voxel-wise statistical analysis is complex and multi-step, requiring standardization of processing pipelines and adequate training for researchers (Clement et al., 2022a).



Figure 3.4: Conceptual overview of the ASL experiment. The acquisition of a so-called *label* image consists of three subsequent parts: magnetic labeling of arterial blood proximal to the brain, a single post-labeling delay time to allow the labeled blood to flow to and exchange with the brain tissue, and finally acquiring an image of the brain. Besides label images, *control* images without prior labeling are acquired. The difference between control and label images originates from the labeled spins delivered to the brain tissue by perfusion, thus resulting in a *perfusion-weighted* image. Finally, by means of a quantification model, a *CBF map* can be obtained.

3.5.3 Recommended implementation

As mentioned earlier, many new developments and improvements have been proposed to augment the ASL imaging protocol over the years. In what follows, some of the main implementation details recommended by the ASL consensus paper by Alsop et al. (2015) are briefly summarized. A distinction is made between aspects related to the generation of the ASL signal and the subsequent quantification of the perfusion-weighted image to a CBF map.

3.5.3.1 Signal generation

As highlighted in Fig. 3.4, the acquisition of a label image consists of three subsequent steps. First, the hydrogen nuclei of the arterial blood water proximal to the brain are magnetically labeled. Second, a post-labeling delay is used to allow the labeled blood to flow to and exchange with the brain tissue. Third, a brain image is acquired using a fast readout scheme. In what follows, a detailed description of these subsequent steps is provided.

Labeling strategies

Over time, multiple labeling strategies have been proposed, which can be grouped into different categories. The major clinically applied techniques for magnetic labeling of the hydrogen nuclei of arterial blood water are spatially-selective pulsed labeling (PASL) and continuous labeling strategies. Continuous labeling methods are typically further divided into true continuous labeling (CASL) and **pseudo-continuous labeling (pCASL)**. The latter has been put forward as the recommended labeling technique, due to its high labeling efficiency combined with its ease of implementation and hardware specifications for clinical scanners (Alsop et al., 2015). Fig. 3.5 shows a schematic representation of these three spatially-selective labeling strategies.



Figure 3.5: Schematic representation of spatially-selective labeling strategies. Both in pulsed and continuous labeling, arterial blood is labeled proximal to the imaging volume. In PASL, labeling is performed using a single short RF pulse (in the order of milliseconds) that inverts the arterial blood magnetization within a whole slab of tissue below the brain, which includes the supplying large arteries (Kwong et al., 1995; Kim, 1995; Wong et al., 1998). In continuous labeling methods, labeling is performed for a longer period (in the order of seconds) by applying continuous RF energy to a labeling plane which inverts the magnetization of arterial blood as it flows through that plane, a process known as flow-driven adiabatic inversion. The difference between continuous ASL (CASL) and pseudo-continuous ASL (pCASL) lies in the way continuous RF energy is established. In CASL, adiabatic inversion is established by means of a constant gradient and a constant RF pulse (Detre et al., 1992; Williams et al., 1992). In pCASL, the same effect as in CASL is created by a long pulse train of slice-selective RF and gradient pulses (Dai et al., 2008). Note that the figure also indicates the position of two inversion pulses for when background suppression is applied as an option to suppress physiological noise. Furthermore, au denotes the labeling duration and PLD (or TI for PASL) denotes the post-labeling delay, i.e. the time between the end of the labeling pulse (train) and the readout. Note that the total repetition time (TR) to acquire a single label-control image pair is subdivided into two parts (one with and one without labeling) of length TR/2.

More recently, velocity-selective labeling has been proposed as an alternative labeling strategy, which labels blood based on its velocity and creates a magnetic bolus immediately proximal to the microvasculature within the imaging volume (Wong et al., 2006). It provides a significant

innovation over traditional labeling approaches as it eliminates arterial transit time confounds and can provide a significant boost in SNR. As the description of this labeling strategy is quite extensive and not the scope of this thesis, the reader is referred to the original publication and a recent review paper of the ISMRM Perfusion Study Group on the topic of velocity-selective ASL (VSASL) (Qin et al., 2022).

In the remainder of this thesis, if not mentioned otherwise, pCASL labeling is implicitly assumed as the labeling strategy of choice.

Labeling duration

An optimal choice of the labeling duration τ (see Fig. 3.5) in pCASL is important because it determines the final precision of the estimation of the perfusion parameters. Two considerations influence this choice. First, extending the labeling duration is beneficial as it increases the ASL signal and pCASL SNR. Note however that the SNR increase with a longer labeling duration is limited, because for a labeling duration much longer than the T1 relaxation time of blood the signal gain decreases. Typically, for the T1 of blood, one assumes a population average of 1.65 s at 3T (Lu et al., 2004). Second, a longer labeling duration increases the TR per label-control image pair, and as such reduces the number of label-control image pairs (i.e. number of averages) that can be acquired per unit time. Clearly there exists a trade-off: a longer labeling duration increases the SNR, which has a positive effect on the estimation precision, however, it reduces the amount of label-control image pairs per unit time, which reduces the number of averages and thus has a negative effect on estimation precision.

Given this trade-off, a **labeling duration of 1.8 s** was recommended in the consensus paper (Alsop et al., 2015). However, this choice is a pragmatic compromise, the labeling duration remains an acquisition parameter that is in the best case optimized as a function of the estimation precision (Bladt et al., 2020a).

Single-PLD and multi-PLD acquisitions

In pCASL, the time between the end of the labeling pulse train and the start of the readout module is referred to as the post-labeling delay (PLD) (see Fig. 3.5). pCASL acquisitions are typically classified into so-called single-PLD and multi-PLD acquisitions:

Single-PLD In this ASL protocol, each control-label image pair is acquired with the same acquisition settings, thus also using the same PLD. It implies that the entire scan time is used to create repetitions of control-label image pairs with fixed acquisition settings, often referred to as *averaging*. Single-PLD pCASL was advised by the consensus paper (Alsop et al., 2015), using a PLD at least as long as the longest estimated arterial transit time (ATT), i.e. the time it takes for the labeled bolus to travel from the labeling plane through the arterial vascular tree towards a certain part of brain tissue. Like the CBF, the ATT is a biophysical parameter and potential biomarker that varies between different regions of the brain, between individuals and between healthy and pathological tissue (Petersen et al., 2010). By using a recommended PLD>ATT, it is ensured that the labeled bolus reaches the capillaries in the tissue, largely avoiding remaining ASL signal in large supplying arteries which would show up in the perfusion images as bright spots, mimicking hyperperfusion. Moreover, it (theoretically) guarantees the arrival of the entire bolus in the microcirculation of the

target tissue throughout the brain, reducing the dependence of perfusion quantification on the underlying local ATT (Alsop & Detre, 1996). Since the velocity of blood is different in children, adults or clinical patients, the PLD has to be adapted accordingly. Table 3.1 summarizes recommended PLD values (Alsop et al., 2015).

Multi-PLD pCASL In contrast to single-PLD, this protocol acquires ASL images with a different PLD per control-label image pair, which allows to sample the ASL perfusion process dynamically at multiple time points (Gonzalez-At et al., 2000; Wang et al., 2013). On the one hand, this allows for a direct parameter estimation of both CBF and ATT. It has been shown that CBF estimation accuracy depends on the ATT (Alsop et al., 2015; van Osch et al., 2018; van der Thiel et al., 2018), especially when the ATT varies over a large range in a subject or in the studied population (van Osch et al., 2018; van der Thiel et al., 2018; so estimating both parameters simultaneously is advantageous. On the other hand, multi-PLD pCASL has the disadvantage of a reduced SNR for the perfusion-weighted image for a given PLD, as the number of control-label image pairs per PLD is reduced compared to single-PLD, obeying a similar clinically acceptable total scan time.

Table 3.1: Recommended PLD values for single-PLD pCASL (Alsop et al., 2015).

Subject	PLD value [ms]	
neonates	2000	
children	1500	
healthy subjects $<$ 70 years	1800	
healthy subjects $>$ 70 years	2000	
adult clinical patients $<$ 70 years	2000	

While multi-PLD methods provide additional information, they are more complex and require more measurements and processing. At the present time, therefore, **single-PLD methods are recommended** as the default ASL method (Alsop et al., 2015).

Background suppression

As stated in section 3.5.2, image intensities between control and label images only differ by at most 5%. As a consequence, the ASL signal has a relatively low SNR. Therefore, the use of background suppression (BS) is often recommended as a way to significantly increase the overall SNR of the ASL signal (Garcia et al., 2005; Maleki et al., 2012).

To understand the core concept of BS, it is crucial to distinguish the different noise components that occur in MR images. Generally, when speaking about noise in MRI, one considers either thermal noise or scanner induced noise components, which combined constitute the raw noise component with standard deviation σ_0 . This *raw noise* is proportional to the static magnetic field strength B_0 and independent of the MR-signal intensity (Edelstein et al., 1986). However, as highlighted by Krüger & Glover (2001), there exist also noise components that are signal-dependent, denoted as *physiological noise*, described by a standard deviation $\sigma_P = c \cdot S$, with c a constant and S the MR-signal intensity. Such physiological noise stems from different factors, including local motion artifacts caused by cardiac and respiratory function and magnetic field modulations (Krüger & Glover, 2001). Taking into account both signal-independent and signal-dependent noise components, the total image noise can thus be described by the following standard deviation:

$$\sigma = \sqrt{\sigma_0^2 + \sigma_P^2}.$$
(3.3)

For a certain signal intensity S, which is constant over time in the case of repeated acquisition of control-label image pairs when using a fixed PLD, the SNR can be formally defined as (Krüger & Glover, 2001):

$$SNR = \frac{S}{\sigma} = \frac{S}{\sqrt{\sigma_0^2 + \sigma_P^2}} = \frac{SNR_0}{\sqrt{1 + c^2 SNR_0^2}}$$
 (3.4)

with $SNR_0 = S/\sigma_0$ the signal-to-raw-noise ratio in the absence of physiological noise.

Knowing that physiological noise scales with the image signal intensity, it can be significantly reduced by suppressing the signal intensity. In the case of ASL, if the background signal S in the label and control images (which is theoretically equal in both images) can be suppressed without gravely affecting the ASL signal, it could significantly increase the SNR of the ASL signal in the eventual difference image. Indeed, if we assume the signal intensities in the unsubtracted images to be Gaussian distributed, the SNR of the ASL signal can be written as (Bladt, 2020):

$$SNR_{ASL} = \frac{S_{ASL}}{\sqrt{2} \cdot \sqrt{\sigma_0^2 + \sigma_P^2}}$$
(3.5)

Clearly, the SNR_{ASL} will increase as σ_P , which scales with the background signal *S*, reduces.

BS can be achieved using a combination of a saturation pulse and a certain number of inversion pulses applied to the imaging volume (Garcia et al., 2005; Maleki et al., 2012). By timing the inversion pulses correctly, the longitudinal magnetization of the background tissue will pass through zero at the time of readout. Note, however, that while significantly improving the overall SNR of the ASL signal, BS is subject to two main limitations. First, there exists a trade-off in the amount of inversion pulses to be used. Increasing this amount ensures suppression of the static tissue signal over an increasing range of T_1 values. Unfortunately, due to inevitable imperfections in the inversion pulses, the labeling efficiency decreases by approximately 5% for each extra inversion pulse, resulting in unwanted ASL signal loss. In order to balance this trade-off, **background suppression with two inversion pulses is recommended** (Alsop et al., 2015). Second, the longitudinal magnetization of static tissues is only canceled at a given time point, making BS well suited for 3D readout which uses a single excitation pulse for every slice. This second limitation is more extensively discussed in contribution chapter 7.

Readout sequence and spatial resolution

The original ASL consensus paper (Alsop et al., 2015) put forward **segmented 3D sequences, such as 3D GRASE** (Günther et al., 2005; Fernández-Seara et al., 2005), as the recommended readout method, followed by **single-shot 2D multi-slice (EPI) readout as a back-up choice**. Recently, another review/recommendation paper from the ISMRM Perfusion Study Group confirmed this choice, also highlighting a number of anticipated readout improvements (Hernandez-Garcia et al., 2022). The 2D EPI multi-slice readout, which was used for the acquisition of the real data in contribution chapter 7, has already been described in section 2.4.1.

In most ASL works, segmented 3D readout is preferred because of three main advantages over single-shot 2D readout. First, as the entire image is acquired in one excitation, background suppression can be maximal for the entire volume by timing it correctly with the readout excitation (Ye et al., 2000; Krüger & Glover, 2001; Garcia et al., 2005; Maleki et al., 2012; Paschoal et al., 2021). In 2D readout methods, there is an excitation for the acquisition of each slice. In that case, background suppression can only be optimal in the first slice and will become less effective in each subsequent slice. This can be considered the most important reason for choosing 3D over 2D readout in ASL. Second, 3D acquisition is less susceptible to magnetic field inhomogeneities. Third, the total readout time for a volume is generally lower for 3D readout. However, this effect is limited, as the labeling duration and PLD take up most time of the TR. At the same time, it is worth noting that a considerable downside of 3D readout is its sensitivity to subject motion during acquisition. Any movement that occurs during 3D readout, which takes 300-450 ms for whole-brain coverage using the recommended spatial resolution (Vidorreta et al., 2013, 2014), cannot be untangled afterwards in postprocessing. On the contrary, 2D multi-slice readout is less sensitive to motion as the acquisition time per slice is very short (\sim 50 ms) (Vidorreta et al., 2013). Finally, another downside of 3D readout is its use of long echo trains, resulting in through-plane blurring due to T_2 decay along the echo train and in-plane blurring due to T_2^* decay between refocusing pulses. Splitting the readout into more segments can reduce this blurring, but at the cost of a longer acquisition time and increased sensitivity to inter-shot motion and physiological fluctuations (Hernandez-Garcia et al., 2022).

The intrinsically low SNR of ASL difference images can be partially mediated by choosing a low spatial resolution for readout. For this reason, it has been recommended to use **a** spatial resolution of 3-4 mm in-plane and 4-8 mm through-plane (Alsop et al., 2015). An obvious downside of a low spatial resolution is the occurrence of partial volume effects (PVE) which result in a loss of fine anatomical details in the ASL perfusion-weighted images as well as in the quantified perfusion parameter maps. Correcting for these PVEs in ASL MRI has been actively studied (Asllani et al., 2008; Chappell et al., 2011; Liang et al., 2013). Very often, resolution enhancement is based upon the use of additionally acquired high-resolution structural images to assist the reconstruction (Meurée et al., 2019; Mehranian et al., 2020). A downside to the latter is the requirement of accurate co-registration and/or distortion correction to guarantee a spatial correspondence between the high-resolution structural and the low-resolution ASL acquisitions, and to correct for differences in readout methods between both type of images. When image registration and other corrections are performed separately from the final perfusion parameter estimation, this could potentially result in propagating errors. Therefore, in contribution chapter 7 of this thesis, the potential benefit of combining single-PLD ASL with model-based super-resolution reconstruction and joint motion estimation has been extensively explored as a robust idea to provide both resolution enhancement and accurate motion compensation.

3.5.3.2 Perfusion parameter quantification

In MR relaxometry and most other qMRI applications, a given signal that changes over time is sampled at multiple time points and quantitative parameters, which parameterize the signal

change over time, are estimated by fitting a suitable signal model to the acquired data points (cf. section 3.2). In that sense, the concept of a single-PLD pCASL experiment is markedly different, as the dynamic pCASL signal is only sampled at one time point. Based on such data, only one parameter, the CBF, can be quantified in a unique way. It also simplifies the estimation process considerably; data acquired at one time point allows for parameter quantification by means of a closed-form expression between the data and the parameter to be quantified. This closed-form expression is often referred to as the ASL 'quantification formula', which can be derived from a more general single-compartment model. As will be highlighted hereafter, this derivation comes with a number of assumptions that should be well understood as they impact the final accuracy and precision of the CBF quantification.

Quantification formula

The most simple model describing the pCASL signal is the **single-compartment model**. The central assumption in this model is that, when the magnetically labeled water molecules reach the tissue voxel, there is unrestricted transfer of water molecules between the blood compartment and the tissue compartment. In other words, upon arrival in the tissue voxel, there is an immediate equal concentration of labeled water molecules in the blood compartment and the tissue compartment. Therefore, the tissue voxel can be seen as a single compartment. This concept is visualized in Fig. 3.6.

The single-compartment dynamics can be described in two ways: using modified Bloch equations (Detre et al., 1992; Parkes & Tofts, 2002), or by convolution of a labeled spin bolus function with a tissue response function (Buxton et al., 1998). In what follows, we use the derivation from modified Bloch equations.



Figure 3.6: Single-compartment model. Labeled arterial blood enters the tissue voxel with magnetization m_a and perfusion rate f, and leaves with magnetization m_v . The tissue voxel has magnetization M_z and longitudinal relaxation time $T_{1,tissue}$. Figure adapted from Parkes & Tofts (2002).

As already introduced in section 2.3.4 of Chapter 2, a Bloch equation can be used to describe the change in longitudinal magnetization (cf. Eq. (2.14)). For convenience, let us repeat this Bloch equation for a certain longitudinal magnetization M_z in a certain tissue voxel:

$$\frac{dM_z(t)}{dt} = \frac{M_0 - M_z(t)}{T_{1,\text{tissue}}},$$
(3.6)

with $T_{1,tissue}$ the longitudinal relaxation time of tissue. In pCASL imaging, labeled spins will enter and leave the tissue voxel with a perfusion rate f. Under the assumption of a single compartment model, and ignoring magnetic transfer contrast effects, Eq. (3.6) can be modified to include the inflow and outflow of magnetization (Detre et al., 1992; Parkes &

CHAPTER 3

Tofts, 2002):

$$\frac{dM_z(t)}{dt} = \frac{M_0 - M_z(t)}{T_{1,\text{tissue}}} + fm_a(t) - fm_v(t), \qquad (3.7)$$

with $m_a(t)$ the inflowing magnetization from the labeled bolus and $m_v(t)$ the venous outflow of labeled spins. The unrestricted and instantaneous equilibration of the concentration of labeled spins between the vascular and tissue compartments when the labeled bolus enters the tissue voxel means that the blood leaving the tissue carries the same concentration of labeled spins as the water within the tissue voxel (Buxton et al., 1998). Consequently, the outflowing magnetization $m_v(t)$ is proportionate to the tissue magnetization $M_z(t)$, weighted with the ratio of water content between tissue and blood, denoted as λ (Buxton et al., 1998; Parkes & Tofts, 2002):

$$\frac{dM_z(t)}{dt} = \frac{M_0 - M_z(t)}{T_{1,\text{tissue}}} + fm_a(t) - f\frac{M_z(t)}{\lambda}.$$
(3.8)

This ratio λ is often referred to as the blood-brain partition coefficient of water. Since, in (conventional) pCASL a label image is subtracted from a control image, Eq. (3.8) can be adapted to (Parkes & Tofts, 2002):

$$\Delta \frac{dM_z(t)}{dt} = \Delta \frac{M_0 - M_z(t)}{T_{1,\text{tissue}}} + \Delta f m_a(t) - \Delta f \frac{M_z(t)}{\lambda}, \qquad (3.9)$$

where Δ represents the signal difference between the control and label image. Under the reasonable assumption that M_0 , $T_{1,\text{tissue}}$, λ and f do not change between the acquisition of both images, Eq. (3.9) can be further simplified as:

$$\frac{d\Delta M_z(t)}{dt} = -\frac{\Delta M_z(t)}{T_{1,\text{tissue}}} + f\Delta m_a(t) - f\frac{\Delta M_z(t)}{\lambda}$$
$$= -\frac{\Delta M_z(t)}{T_1'} + f\Delta m_a(t), \qquad (3.10)$$

with $1/T'_1 = 1/T_{1,tissue} + f/\lambda$. This differential equation can be solved if $\Delta m_a(t)$, the difference in arterial magnetization flowing into the tissue voxel between the label and control image, is known. If uniform plug flow is assumed for the labeled bolus as it travels from the labeling plane to the tissue voxel, it can be described as (Buxton et al., 1998):

$$\Delta m_a(t) = \begin{cases} 0 & t < \Delta t, \\ 2M_{0b}\alpha \exp\left(-\frac{\Delta t}{T_{1b}}\right) & \Delta t < t < \Delta t + \tau, \\ 0 & t > \Delta t + \tau, \end{cases}$$
(3.11)

with τ the pseudo-continuous labeling duration, Δt the ATT between the labeling plane and the tissue voxel, M_{0b} the equilibrium magnetization of arterial blood in a unit voxel, and α the inversion efficiency of the labeling. The factor $\exp(-\Delta t/T_{1b})$ describes the magnetization loss due to longitudinal relaxation in the arterial blood compartment during the travel time from the labeling plane to the tissue voxel. Note that the factor 2 originates from the fact that Eq. (3.11) describes the difference in magnetization between the label and control image, which is equal to twice the equilibrium magnetization of blood at t = 0due to the 180° inversion of the arterial magnetization at the labeling plane. When Eq. (3.11) is substituted in Eq. (3.10), the following expression is found for the difference magnetization:

$$\Delta M(t) = \begin{cases} 0 & t < \Delta t, \\ 2M_{0b}\alpha f T_1' \exp\left(-\frac{\Delta t}{T_{1b}}\right) \left(1 - \exp\left(-\frac{t - \Delta t}{T_1'}\right)\right) & \Delta t < t < \Delta t + \tau, \\ 2M_{0b}\alpha f T_1' \exp\left(-\frac{\Delta t}{T_{1b}}\right) \exp\left(-\frac{t - \Delta t}{T_1'}\right) \left(\exp\left(-\frac{\tau}{T_1'}\right) - 1\right) & t > \Delta t + \tau, \end{cases}$$

$$(3.12)$$

which is commonly referred to as the *single-compartment* model. As follows from Eq. (3.12), this model consists of three distinct phases: no signal as long as the labeled bolus has not yet reached the tissue voxel ($t < \Delta t$), a build-up of the pCASL signal as labeled spins flow into the tissue voxel ($\Delta t < t < \Delta t + \tau$), followed by a phase where the entire labeled bolus has arrived in the tissue voxel and longitudinal relaxation exponentially decays the pCASL signal ($t > \Delta t + \tau$). An example is shown in Fig. 3.7.

Often, two important additional assumptions are made that further simplify Eq. (3.12) to arrive at a quantification formula for single-PLD pCASL:

- **PLD is longer than the arterial transit time.** This assumes that the entire labeled bolus arrives at the tissue voxel. Since for the PLD it holds that $PLD = t \tau$, the assumption that $PLD > \Delta t$ means that the pCASL signal is described by the third regime of Eq. (3.12).
- The blood vessels are impermeable. This means that the labeled spins remain in the blood compartment within the tissue voxel. In that case, the difference magnetization ΔM decays with the blood longitudinal relaxation time T_{1b} instead of $T_{1,tissue}$.



Figure 3.7: An example of the difference magnetization ΔM as a function of time *t* according to the single-compartment model for a labeling duration $\tau = 1.8$ s. Physiological parameters were chosen to represent a gray matter voxel: f = 50 mL/100 g/min, $\Delta t = 0.6$ s, $T_{1,\text{tissue}} = 1.4$ s, $T_{1b} = 1.65$ s, $\lambda = 0.9$. Values for $T_{1,\text{tissue}}$ and T_{1b} are given assuming a static magnetic field strength $B_0 = 3.0$ T. Also, α and M_{0b} were assumed equal to 0.85 and 1, respectively.

Given these two additional assumptions, the third regime of Eq. (3.12) simplifies to:

$$\Delta M(t) = 2M_{0b}\alpha f T_{1b} \exp\left(-\frac{\Delta t}{T_{1b}}\right) \exp\left(-\frac{t-\Delta t}{T_{1b}}\right) \left(\exp\left(-\frac{\tau}{T_{1b}}\right) - 1\right) \quad (3.13)$$

$$= 2M_{0b}\alpha f T_{1b} \exp\left(-\frac{t-\tau}{T_{1b}}\right) \left(1 - \exp\left(-\frac{\tau}{T_{1b}}\right)\right).$$
(3.14)

It follows that with $PLD = t - \tau$ and with M_{0b} approximated by the equilibrium magnetization in the tissue voxel S_{PD} divided by the blood-brain partition coefficient λ , i.e. $M_{0b} = S_{PD}/\lambda$, rearranging Eq. (3.14) as a function of the CBF f, results in the so-called recommended **quantification formula** (independent of the ATT) (Alsop et al., 2015; Buxton et al., 1998):

Recommended CBF quantification formula

$$\mathsf{CBF} = 6000 \cdot \frac{\lambda \exp\left(\frac{\mathsf{PLD}}{T_{1b}}\right)}{2\alpha T_{1b} \left(1 - \exp\left(-\frac{\tau}{T_{1b}}\right)\right)} \cdot \frac{\Delta S}{S_{\mathsf{PD}}},\tag{3.15}$$

with λ the blood-brain partition coefficient in mL/g, ΔS the averaged difference between the label and control signals, T_{1b} the longitudinal relaxation time of blood, α the labeling efficiency, $S_{\rm PD}$ the proton density signal (which is obtained from a separately acquired calibration image), τ the labeling duration, PLD the postlabeling delay time, and a factor of 6000 to convert the units of the CBF from mL/g/s to mL/100g/min.

Accuracy and precision of the recommended quantification

In essence, the CBF value in Eq. (3.15) can be seen as a multiplication of the relative single-PLD pCASL perfusion signal ΔS by a certain prefactor consisting of a large number of parameters. The latter are typically either assumed to be known or determined through an additional experiment. The PLD and labeling duration τ are parameters that are known and chosen by the practitioner. The proton density signal S_{PD} is acquired as a separate image, which needs to be spatially aligned, i.e., co-registered, to the control-label pCASL image pairs. The remaining parameters are either fixed to population means, in the case of T_{1b} and λ , or to experiment repetition means, in the case of α . Specifically, in the recommendations by Alsop et al. (2015), the following assumptions are made: $\lambda = 0.9 \text{ mL/g}$, $T_{1b} = 1.65 \text{ s at } 3.0 \text{ T}$, and $\alpha = 0.85$. As these parameters can vary significantly between individuals or between repetitions of the experiment, they are a **potential source of bias**. The effects of such a bias are being extensively studied in literature (Bladt et al., 2020b; Bladt, 2020).

The recommended quantification formula in itself is also a source of inaccuracy, as it is based on assumptions that are approximations of reality. Of course, to some point, this is true for any chosen perfusion model, as it is unlikely any model exactly describes the physiology underlying the pCASL perfusion process. Furthermore, in terms of the PLD, efforts have been made in the recommended implementation to minimize bias when quantifying the CBF with Eq. (3.15). However, optimal single-PLD selection is difficult for accurate CBF quantification. In particular, it has been shown that when the ATT varies over a large range in a subject or in a studied population, single-PLD pCASL remains vulnerable to under or overestimation of CBF (van Osch et al., 2018; van der Thiel et al., 2018).

While single-PLD pCASL quantification clearly suffers from potential low accuracy (which could be mitigated by the use of more complex multi-PLD pCASL (cf. section 3.5.3.1)), its strong suit is a **high precision**. Firstly, the use of repeated measurements of control-label image pairs at the same time point greatly increases the SNR of the ASL signal in the eventual averaged ASL difference image ΔS . As all other parameters on the right hand side are known or assumed to be known, error propagation dictates that the precision of CBF quantification scales with the SNR of the averaged ASL difference image ΔS . Secondly, unlike in a nonlinear multi-PLD pCASL model where the number of parameters to estimate is increased to enhance model accuracy, the reduced number of parameters in the single-PLD model may contribute to improved precision.

3.5.4 Clinical applications

As a versatile complement to other medical imaging modalities, ASL imaging can provide insightful information to establish diagnosis, to monitor the evolution of pathologies, or to characterize disease states. In this section, a brief list of some recent applications of ASL is given, paying particular attention to neuroimaging applications. For a more elaborate list, the reader is referred to some excellent reviews in literature (Detre et al., 2012; Grade et al., 2015; Hernandez-Garcia et al., 2019).

- **Cerebrovascular disease** ASL has emerged as a valuable tool in the assessment of cerebrovascular disease, offering non-invasive insights into perfusion alterations (Alsop et al., 2015). As an example, Wang et al. (2012) demonstrated the efficacy of pCASL in detecting regional CBF changes associated with ischemic strokes, aiding in the early identification and characterization of affected brain regions.
- **Dementia** In the realm of dementia research, pCASL has proven instrumental in evaluating cerebral perfusion patterns associated with neurodegenerative disorders. Johnson et al. (2005) utilized pCASL to investigate perfusion abnormalities in Alzheimer's disease, highlighting its potential as a sensitive imaging technique for early detection and monitoring disease progression.
- **Neuro-oncology** The application of pCASL in neuro-oncology has been pivotal in delineating tumor-related perfusion characteristics. Specifically, ASL is able to distinguish between high-grade and low-grade gliomas (i.e, tumor cells that start growing in the brain or spinal cord) based on perfusion patterns, offering valuable information for treatment planning and prognosis assessment (Alsaedi et al., 2019).
- **Psychiatric disease** Preliminary studies exploring the use of pCASL in psychiatric diseases have shown promise in uncovering cerebral perfusion alterations associated with conditions such as schizophrenia and mood disorders. As an example, Oliveira et al. (2018) used pCASL imaging to identify regional perfusion abnormalities in schizophrenia.
- **Epilepsy** In the field of epilepsy, pCASL has been employed to investigate CBF changes associated with seizure activity. Pendse et al. (2010) demonstrated the utility of pCASL in detecting focal hypoperfusion in areas implicated in epileptogenesis (i.e., the gradual process by which a typical brain develops epilepsy), providing valuable insights for presurgical evaluation and treatment planning in epilepsy patients.

3.5.5 Existing problems in ASL

From the previous section, it is evident that pCASL imaging holds significant promise for the assessment of cerebrovascular diseases and neurological disorders in today's clinical routine. However, several technological barriers persist, hindering its full potential.

- **Inherent low SNR** Like with any MRI technique, the trade-off between SNR, scan time, and spatial resolution is present. Notably, as mentioned earlier, pCASL difference images intrinsically have a **low SNR**. Consequently, one typically proceeds to the acquisition of multiple image pairs with **lower spatial resolution**, after which a signal increase is achieved through averaging. Although this increases the SNR, it introduces some other difficulties. First, the required scan time increases. Due to averaging with single-PLD pCASL, the total scan time becomes a multiple of the scan time for one control-label image pair. Preferentially, such extended scan time is to be avoided for clinical routine where scanning protocols (which often include other MRI contrasts in addition to pCASL) are limited in time for the patient's comfort and to maximize patient throughput. In addition, as a larger number of images are recorded, the **risk of artifacts stemming from patient motion** also increases.
- **Motion correction and quantification** A spatial correspondence of the individual pCASL images is crucial for voxel-wise signal averaging. This requires a robust motion correction strategy to align the different image pairs. Additionally, the recommended quantification formula for single-PLD pCASL also demands the acquisition of an extra proton density weighted calibration image, which is required to convert CBF values in arbitrary units to absolute perfusion units of mL/100g/min. To allow a voxel-wise division with this calibration image (cf. Eq. (3.15)), a co-registration of the calibration image with the lower-resolution (LR) pCASL control and label images is required.
- Partial volume effects Due to the inherently low spatial resolution and large voxel size of pCASL images, partial volume effects (PVE) occur, where perfusion of different tissues contribute to the observed perfusion signal in a voxel (Petr et al., 2018; Chappell et al., 2021). However, for quantitative perfusion analysis of the brain, it is crucial that the CBF estimates derived from the pCASL images can be assigned to specific brain tissue types (gray matter, white matter, etc.), ideally at a high-resolution level. Therefore, it is customary to apply partial volume (PV) correction as a post-processing step of the ASL MRI experiment, so that perfusion can be separated from structural effects when computing the mean perfusion for a certain tissue type. Generally, the tissue volume is obtained using high-resolution (HR) PV maps obtained from the segmentation of a HR T1-weighted (T1w) structural image acquired in the same scanning session as the lower-resolution control-label ASL image pairs (Clement et al., 2022a). Typically, the HR PV maps are co-registered and then downsampled to match the LR ASL image space. It is important to be alert to the differences in readout methods between the T1w and ASL images during this resampling process from HR to LR image space (Petr et al., 2018). For example, prior to downsampling, it is common to perform a Gaussian pre-smoothing step on the HR PV maps as an additional post-processing step to take into account differences in acquisition PSFs between both readout types (Cardoso et al., 2015).

- **Multi-step processing and potential error propagation** It is evident that pCASL imaging demands many separate post-processing steps involving image registration, segmentation, quantification, and so on. Although on the one hand these steps are required to arrive at a correct final perfusion estimate, such **multi-step processing** also entails the danger of propagating errors in the subsequent steps, resulting in a biased end result. That is why throughout the pCASL world (and in qMRI more generally) there exists a stringent need for joint estimation frameworks and ideas that provide single-step approaches to minimize potential propagation errors.
- **Intravoxel dephasing**, also known as *phase dispersion*, refers to the phenomenon where the phase coherence of the magnetization within a voxel is disrupted, resulting in a signal decay due to the destructive interference of magnetization signals from different nuclei within that voxel (Chen & Wyrwicz, 1999). E.g., the transverse magnetization of nuclei within a voxel can lose coherence in regions of magnetic field inhomogeneity, at tissue interfaces with differing magnetic susceptibilities, or due to variations in blood velocities within the voxel (Amukotuwa et al., 2016; Ozsarlak et al., 2004). This leads to a decrease of the net signal from that voxel, compromising the accuracy of the acquired ASL data. Some brain structures which are characterized by a more complex vascular network, such as the cerebellum, are more sensitive to intravoxel dephasing since blood velocity changes occur more frequently, like turbulent or accelerating blood flow in small curved arteries. The risk of phase dephasing can be mitigated by using small TE's, to permit the signal to be detected before the magnetization has had time to dephase. Another approach to reduce intravoxel dephasing is the use of smaller voxels, such that the amount of magnetization that is permitted to combine incoherently is reduced, thus limiting potential signal loss. However, as pointed out earlier, ASL inherently needs larger voxels to obtain enough SNR. Clearly, intravoxel dephasing poses another trade-off dimension to be considered alongside SNR, spatial resolution, and scan time.
- Blurring artefacts The utilization of 3D readout methods such as the recommended 3D GRASE sequence comes with an increased risk of blurring artefacts that arise due to T_2 and \mathcal{T}_2^* relaxation effects during the long pulse echo trains, resulting in signal decay and through-plane blurring in the acquired images (Vidorreta et al., 2014; Zhao et al., 2018). Such blurring artefacts can compromise the spatial resolution and interpretability of the acquired data. Splitting the readout into more segments can reduce this blurring, but at the cost of a longer acquisition time and increased sensitivity to inter-shot motion and physiological fluctuations (Hernandez-Garcia et al., 2022). In addition, the long readout time of a 3D imaging sequence holds an increased risk of motion artefacts (Alsop et al., 2015). As a viable alternative to 3D readout, single-shot 2D multi-slice readout methods based on EPI have been suggested (Alsop et al., 2015). Compared to 3D readout, these 2D readout methods are less susceptible to spatial blurring due to T_2 decay (Vidorreta et al., 2013). However, as will be highlighted in contribution chapter 7, the use of a separate excitation pulse for every slice in 2D readout complicates background suppression. In practice, background suppression can be optimal for only one slice and will be progressively less efficient for other slices (Alsop et al., 2015).

Furthermore, it is important that the ASL community continues to clearly communicate guidelines and agreements on the so-called 'best practices' in the use of pCASL imaging, so that results can be interpreted unambiguously and without implementation bias.

3.6 Parameter estimation

The signal models, described in the previous sections on MR relaxometry and ASL, are mathematical relations that describe the contrast-weighted MR images in terms of physical parameters. By using such models, interesting properties of the object under study can be obtained by estimating the quantitative parameters from a series of acquired contrastweighted MR images. However, acquired MR signals are disturbed by noise and are, as such, random (or stochastic) variables. The random variable is best described by its probability distribution function (PDF) over the continuous range of its possible outcomes. Knowledge about PDFs and statistical parameter estimation is indispensable for a contemporary qMRI researcher. In the contribution chapters of this thesis, the PDFs describing the intensity of magnitude MR images will be more thoroughly discussed. In this chapter, a more general introduction to statistical estimators, i.e. methods to extract information about model parameters from noisy measurements, is provided. Some important estimators are the least squares (LS) estimator, the maximum likelihood (ML) estimator and maximum a posteriori (MAP) estimator, which will be used in part III of this thesis, when the different contributions are elaborated. In what follows, the general properties of these estimators, their strengths, and limitations will be briefly discussed. For a more detailed explanation of statistical parameter estimation, the reader is referred to the works of van den Bos (2007) for a general introduction to parameter estimation, and Gelman et al. (1995) for more information on Bayesian data analysis.

3.6.1 Statistical parameter estimators

An **estimator** can be defined as any function of the observed data, thereby providing an estimate of an unknown (physical) quantity of interest. In mathematical terms, this can be expressed as:

$$\hat{\theta}_N = g(y_1, y_2, \dots, y_N).$$
 (3.16)

Here, $\boldsymbol{y} = (y_1, y_2, \dots, y_N)^T$ denotes a vector of N random samples, or *observations*. The probability density function (PDF) that describes the observations is assumed to be parametric in the true parameter vector $\boldsymbol{\theta}_0$. Note that in general $\boldsymbol{\theta}_0$ is unknown, except for simulation experiments. The estimator is then any function g, while $\hat{\boldsymbol{\theta}}_N$ denotes the estimator of $\boldsymbol{\theta}_0$ specifically.

Estimators can be characterized by certain properties that will allow us to distinguish between "good" and "bad" estimators. Since the observations are stochastic variables, so is the estimator. Therefore, just as the observations, the estimator will also have an expectation value $\mathbb{E}[\hat{\theta}_N]$ and a variance var $(\hat{\theta}_N)$. In this thesis, both $\hat{\theta}_N$ and θ_0 are assumed to be real-valued in $\mathbb{R}^{P \times 1}$, however, the definitions in this section can easily be extended to include complex-valued parameters.

3.6.1.1 Accuracy and precision

Important properties of estimators, which might be used to compare different estimators, are the accuracy and precision, see Fig. 3.8. An estimator is said to be *accurate* when the estimates are *on average* close to the true value of the parameter in which we are interested. Another way of saying this, is that the estimator has a small *bias*, which is defined as (van den

Bos, 2007):

bias
$$(\hat{\theta}_N) = \mathbb{E} \left[\hat{\theta}_N \right] - \theta_0.$$
 (3.17)

An estimator is *unbiased*, and thus perfectly accurate, if its bias equals zero, or equivalently:

$$\mathbb{E}\left[\hat{\boldsymbol{\theta}}_{N}\right] = \boldsymbol{\theta}_{0}. \tag{3.18}$$

Moreover, an estimator can also be *asymptotically unbiased* when its bias goes to zero in the limit of $N \rightarrow \infty$:

$$\lim_{N \to \infty} \mathbb{E}\left[\hat{\boldsymbol{\theta}}_{N}\right] = \boldsymbol{\theta}_{0}. \tag{3.19}$$

The bias defines and quantifies the accuracy or, equivalently, the *systematic error* of an estimator. The larger the absolute value of $\mathbb{E}\left[\hat{\theta}_{N}\right] - \theta_{0}$, the larger the bias, the lower the accuracy of the estimator. Bias may have different sources, including a systematic deviation of the estimator caused by fluctuations of the observations (i.e., the MRI data), a mismatch between the estimation model and the underlying true process, or an insufficient number of observations (van den Bos, 2007).



Figure 3.8: Visual representation of the concepts of accuracy and precision of an estimator. The black ×'s denote the estimates $\hat{\theta}_N$ for different noise realizations. The underlying ground truth parameters θ_0 are assumed to coincide with the bullseye. Precision is associated with random errors, accuracy is associated with systematic errors.

Precision, on the other hand, is a desirable property of an estimator that relates to the *statistical variability* of the estimates, i.e. how much the estimates will vary when the experiment is repeated. The precision of an estimator is generally quantified by its *variance*, i.e. the diagonal elements of the corresponding covariance matrix of the estimator:

$$\operatorname{cov}\left(\widehat{\boldsymbol{\theta}}_{N}\right) = \mathbb{E}\left[\left(\widehat{\boldsymbol{\theta}}_{N} - \mathbb{E}\left[\widehat{\boldsymbol{\theta}}_{N}\right]\right)\left(\widehat{\boldsymbol{\theta}}_{N} - \mathbb{E}\left[\widehat{\boldsymbol{\theta}}_{N}\right]\right)^{T}\right].$$
(3.20)

Note that the precision of the estimator for a certain parameter is inversely related to the variance of the estimator for that same parameter; the higher the variance, the lower the precision. Variance is related to *nonsystematic errors* of an estimator, which are caused by unpredictable fluctuations in the observations (i.e., noise in the data) (van den Bos, 2007). Additionally, the precision of an estimator is directly related to the amount of observations and the noise in the data.

3.6.1.2 (Root) mean squared error

Obviously, a *good* estimator is characterized by a high precision (or low variance), and high accuracy (or low bias). An overarching measure to compare the performance, comprising both accuracy and precision, of different estimators is the *mean squared error* (MSE). First,

we define the *error* ϵ of the parameter estimator as the difference between the estimator $\hat{\theta}_N$ and the true value θ_0 :

$$\epsilon = \hat{\theta}_N - \theta_0. \tag{3.21}$$

Then, the MSE, as its name implies, is equal to the expected value of the squared error. Assuming a **scalar-valued** parameter θ , the MSE can be expressed as:

$$MSE (\hat{\theta}_{N}) = \mathbb{E} \left[\left(\hat{\theta}_{N} - \theta_{0} \right)^{2} \right] \\= \mathbb{E} \left[\left(\hat{\theta}_{N} - \mathbb{E} \left[\hat{\theta}_{N} \right] + \mathbb{E} \left[\hat{\theta}_{N} \right] - \theta_{0} \right)^{2} \right] \\= \mathbb{E} \left[\left(\hat{\theta}_{N} - \mathbb{E} \left[\hat{\theta}_{N} \right] \right)^{2} \right] + \mathbb{E} \left[\left(\mathbb{E} \left[\hat{\theta}_{N} \right] - \theta_{0} \right)^{2} \right] + 2\mathbb{E} \left[\left(\hat{\theta}_{N} - \mathbb{E} \left[\hat{\theta}_{N} \right] \right) \left(\mathbb{E} \left[\hat{\theta}_{N} \right] - \theta_{0} \right) \right] \\= \mathbb{E} \left[\left(\hat{\theta}_{N} - \mathbb{E} \left[\hat{\theta}_{N} \right] \right)^{2} \right] + \left(\mathbb{E} \left[\hat{\theta}_{N} \right] - \theta_{0} \right)^{2} + 2 \left(\mathbb{E} \left[\hat{\theta}_{N} \right] - \mathbb{E} \left[\hat{\theta}_{N} \right] \right) \left(\mathbb{E} \left[\hat{\theta}_{N} \right] - \theta_{0} \right) \\= \operatorname{var} \left(\hat{\theta}_{N} \right) + \left[\operatorname{bias} \left(\hat{\theta}_{N} \right) \right]^{2},$$
(3.22)

where we have used that θ_0 and $\mathbb{E}[\hat{\theta}_N]$ are constants, and that the expected value of a constant equals the constant itself. Clearly, it follows from (3.22) that the MSE of an estimator is equal to its variance plus the square of its bias. When **vector-valued** parameters are considered, the MSE can be further defined as:

$$MSE(\hat{\boldsymbol{\theta}}_{N}) = trace\left[\mathbb{E}\left[\left(\hat{\boldsymbol{\theta}}_{N} - \boldsymbol{\theta}_{0}\right)\left(\hat{\boldsymbol{\theta}}_{N} - \boldsymbol{\theta}_{0}\right)^{T}\right]\right]$$
$$= \sum_{i=1}^{P} var(\hat{\theta}_{i}) + \sum_{i=1}^{P} bias(\hat{\theta}_{i})^{2}.$$
(3.23)

Here, $\hat{\theta}_i$ denotes the *i*th component of the vector $\hat{\theta}_N$. Consequently, the MSE of a vectorvalued parameter is equal to the sum of the MSEs of each of its components. The MSE is always non-negative, where values closer to zero are better. As the MSE is measured in units that are the square of the target parameter, it is often more explanatory to take the root of the MSE and obtain the so-called *root mean squared error* or RMSE, which is measured in the same units as the target parameter:

$$\mathsf{RMSE}\left(\hat{\boldsymbol{\theta}}_{N}\right) = \sqrt{\sum_{i=1}^{P} \mathsf{var}\left(\hat{\theta}_{i}\right) + \sum_{i=1}^{P} \mathsf{bias}\left(\hat{\theta}_{i}\right)^{2}}.$$
(3.24)

3.6.2 Maximum likelihood estimators

Assuming the probability distribution function of vector \boldsymbol{y} is known and given by $p_{\boldsymbol{y}}(\boldsymbol{y}|\boldsymbol{\theta})$ with $\boldsymbol{\theta}$ as set of parameters, then a general estimation method with optimal (asymptotical) statistical properties, both in terms of accuracy and precision can be developed. This method is known as the *maximum likelihood* (ML) estimator. As its name suggests, the ML estimator maximizes the *likelihood function*. The likelihood function is typically denoted as $\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{y})$ and closely related to the probability distribution function (PDF). While $p_{\boldsymbol{y}}(\boldsymbol{y}|\boldsymbol{\theta})$ describes the probability of finding a certain set of observations \boldsymbol{y} given a known set of model parameters $\boldsymbol{\theta}$, the likelihood function considers the inverse. It describes how "likely" a certain set of parameters $\boldsymbol{\theta}$ is to produce the set of observations \boldsymbol{y} . In other words, the likelihood function is a function of the model parameters, whereas the probability distribution function is a

function of the observations. However, note that mathematically, both the likelihood function and the probability distribution function are equal. For a given set of observations y, the ML estimator is defined as:

$$\hat{\boldsymbol{\theta}}_{ML} = \arg\max_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}|\boldsymbol{y}).$$
 (3.25)

When the observations $\boldsymbol{y} = (y_1, y_2, \dots, y_N)^T$ are *statistically independent*, the joint PDF is the product of the PDFs of the set of observations, and Eq. (3.25) can be rewritten as

$$\hat{\boldsymbol{\theta}}_{ML} = \arg \max_{\boldsymbol{\theta}} \prod_{i=1}^{N} \mathcal{L}(\boldsymbol{\theta}|y_i)$$

$$= \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^{N} L_i(\boldsymbol{\theta}|y_i),$$
(3.26)

with $L_i(\boldsymbol{\theta}|y_i) = -\log \mathcal{L}(\boldsymbol{\theta}|y_i)$ the negative log-likelihood function. Both expressions in Eq. (3.26) lead to the same outcome since the logarithmic function is monotonically increasing. However, the minimization of the negative log-likelihood function is more convenient because most nonlinear optimization software tools include optimizers that minimize a cost function criterion.

The ML estimator is known to be asymptotically efficient unbiased, which implies that it is an unbiased estimator that reaches the Cramér-Rao lower bound (i.e., a lower bound on the maximal attainable precision of an unbiased estimator) as the number of data points increases (van den Bos, 2007). In some cases, asymptotically efficient unbiased estimators may already behave asymptotically for unexpectedly small numbers of observations (van den Bos, 2007). In addition, the ML estimator is consistent; as the number of data points increases, the set of estimates of repeats of the experiment converges in probability to the underlying ground truth θ_0 (van den Bos, 2007). These properties are related to the fact that **the distribution of the data is exploited in the ML estimator**. Knowledge about the distribution of MRI data is therefore important to be able to derive the ML estimator and benefit from its statistical properties.

3.6.3 Least squares estimators

Another class of estimators which is often used are so-called *least squares* (LS) estimators. The least squares method determines the optimal set of parameters $\hat{\theta}_{LS}$ of the model of interest f, by minimizing the sum of the element-wise squared residuals ρ of the data:

$$\hat{\theta}_{LS} = \arg\min_{\theta_N} \sum_{i=1}^{N} \varrho_i^2.$$
(3.27)

In the above equation, it is assumed that the data are disturbed by noise. A *residual* is defined as the difference between a measurement y_i , and the predicted value of this measurement by the model $f(x, \theta_N)$, given a set of parameters θ_N and corresponding independent variables x:

$$\boldsymbol{\varrho} = \boldsymbol{y} - f(\boldsymbol{x}, \boldsymbol{\theta}_N). \tag{3.28}$$

Note the subtle, though important, difference between *residuals* (ϱ) and *errors* (ϵ) of the data, the latter can be written as:

$$\boldsymbol{\epsilon} = \boldsymbol{y} - f(\boldsymbol{x}, \boldsymbol{\theta}_0), \qquad (3.29)$$

79

with θ_0 the vector of ground truth model parameters. In general, it is assumed that the independent variables x are known exactly and are thus noise-free.

When the function f is a linear function of θ , Eq. (3.27) is termed the *Linear LS (LLS)* estimator, while if the function f is a non-linear function of θ , Eq. (3.27) is termed the *Non-Linear LS (NLLS)* estimator. Since practically all of the models considered in this PhD thesis are non-linear, we will focus on the NLLS case. For a more elaborate introduction to LLS estimators, the reader is referred to the work of van den Bos (2007).

In addition, the squared differences in Eq. (3.27) are sometimes multiplied by deterministic values w_i in order to weight the contribution of every residual differently. In that case, the term weighted (N)LLS is used. Unweighted (or equivalently, uniformly weighted) (N)LLS estimators are sometimes called ordinary (N)LLS estimators to make the distinction more clear.

3.6.4 Bayesian estimators

The ML and LS estimators introduced above are typically referred to as *frequentist approaches*, since these type of estimators consider an event's probability as the limit of its relative frequency in a large number of trials. In the frequentist approach, unknown parameters are treated as having a *fixed but unknown* value. Consequently, this obviates parameters to be treated as random variables in any sense, i.e. there is no probability about the parameters.

The *Bayesian approach*, on the other hand, interprets probability as a reasonable expectation representing a total state of knowledge. Although a bit philosophical, in essence this means the Bayesian approach treats the **unknown parameters as random variables** and allows to associate probabilities with these parameters, thereby representing the experimenter's belief that a given value of the parameter is true.

3.6.4.1 Bayes theorem

Central to the entire statistical branch of Bayesian inference is Bayes' theorem, named after the 18th-century English Reverend Thomas Bayes. Bayes' theorem describes the conditional probability of an event by accounting for prior knowledge of conditions that might affect this event, thereby providing a way to revise existing predictions or theories given new or additional evidence. In its most general form, Bayes' theorem states that:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)},$$
 (3.30)

where P(A|B) represents the conditional probability of the likelihood of event A occurring given that B is true (and vice versa for P(B|A)), while P(A) and P(B) represent the marginal probabilities of, respectively, observing A or B independent of each other. When interpreted specifically in terms of parameter estimation, Eq. (3.30) can be written as:

$$p(\boldsymbol{\theta}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\boldsymbol{y})},$$
(3.31)

with $p(\theta|y)$ the **posterior distribution** of the parameters, $p(y|\theta) = \mathcal{L}(y|\theta)$ the **likelihood** function, as introduced in section 3.6.2, and $p(\theta)$ and p(y) the **prior** distributions of θ and

y, respectively. The posterior distribution function is arguably the most interesting from a parameter estimation point of view, as it describes everything known about the parameter θ after the experimental outcome y has been observed. The prior $p(\theta)$ encompasses everything known about the parameter vector θ before the actual observations are made. The prior distribution of the data p(y), on the other hand, is typically just a scalar positive constant, and can consequently often be ignored for the purpose of parameter estimation. So, from a Bayesian point of view, the likelihood function and prior distribution of the parameters define the statistical model for the estimation problem, whereas the posterior distribution contains its solution. Given that the parameters are regarded as random variables, the posterior density can be used to deduce any characteristic of the PDF of the parameters, given the data. Hence, the posterior density should always be regarded as the most general solution to the estimation problem.

3.6.4.2 Maximum a posteriori estimator

The Maximum a Posteriori (MAP) estimator is a popular Bayesian estimator that maximizes the posterior distribution $p(\theta|y)$ w.r.t. the parameters θ :

$$\hat{\boldsymbol{\theta}}_{MAP} = \arg \max_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\boldsymbol{y})$$

= $\arg \max_{\boldsymbol{\theta}} \left[p(\boldsymbol{y}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) \right],$ (3.32)

By maximizing $p(\theta|y)$ w.r.t. θ , the MAP estimator equals the *mode* of the posterior distribution $p(\theta|y)$ (see Fig. 3.9). The mode represents the parameter set θ that is most likely to be sampled. For symmetric unimodal distributions, such as the normal distribution, the mean, median, and mode all coincide. However, for asymmetric unimodal distributions, such as the Rician distribution which describes noise in magnitude MR images (Gudbjartsson & Patz, 1995; den Dekker & Sijbers, 2014), the mode may differ from the mean or median.



Figure 3.9: Geometric visualization of the concept of the mode of a distribution, as compared to the median and mean of an arbitrary probability density function.

The MAP estimator is an appropriate estimator for problems where **prior knowledge about the parameters is available**, e.g. some prior belief distribution over the possible values that the parameters could take on. In contrast, the ML en LS estimators, i.e. frequentist approaches, estimate parameters based on data alone.

Similar as for the ML estimator in section 3.6.2, one typically optimizes Eq. (3.32) in the (negative) log domain to facilitate the use of modern-day optimization software tools that include optimizers that minimize a cost function criterion. More specifically, it follows from Eq. (3.32) that:

$$\hat{\theta}_{MAP} = \arg\min_{\boldsymbol{\theta}} - \left[\log(p(\boldsymbol{y}|\boldsymbol{\theta})) + \log(p(\boldsymbol{\theta}))\right], \qquad (3.33)$$

where we have used that the logarithmic function is monotonically increasing and that the logarithm of a product of distributions is equal to the sum of the logarithms.

References

- Achten, E. (2001). MR Imaging Strategies in Epilepsy, (pp. 201–214). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Alsaedi, A., et al. (2019). The value of arterial spin labelling in adults glioma grading: systematic review and meta-analysis. *Oncotarget*, 10(16), 1589–1601.
- Alsop, D. C., & Detre, J. A. (1996). Reduced Transit-Time Sensitivity in Noninvasive Magnetic Resonance Imaging of Human Cerebral Blood Flow. J. Cereb. Blood Flow Metab., 16(6), 1236–1249.
- Alsop, D. C., Detre, J. A., Golay, X., Guenther, M., Hendrikse, J., Hernandez-Garcia, L., Lu, H., MacIntosh, B. J., Parkes, L. M., Smits, M., van Osch, M. J. P., Wang, D. J. J., Wong, E. C., & Zaharchuk, G. (2015). Recommended Implementation of Arterial Spin-Labeled Perfusion MRI for Clinical Applications: A Consensus of the ISMRM Perfusion Study Group and the European Consortium for ASL in Dementia. *Magn. Reson. Med.*, 73(1), 102–116.
- Alzheimer Europe (2019). Dementia in Europe Yearbook 2019: Estimating the prevalence of dementia in Europe. *Luxembourg, Alzheimer Europe*.
- Amukotuwa, S. A., Yu, C., & Zaharchuk, G. (2016). 3D Pseudocontinuous arterial spin labeling in routine clinical practice: A review of clinically significant artifacts. J. Magn. Reson. Imaging, 43(1), 11–27.
- Anderson III, A. G., et al. (2013). Adaptive retrospective correction of motion artifacts in cranial MRI with multicoil three-dimensional radial acquisitions. *Magn. Reson. Med.*, 69(4), 1094–1103.
- Asllani, I., Borogovac, A., & Brown, T. R. (2008). Regression algorithm correcting for partial volume effects in arterial spin labeling MRI. *Magn. Reson. Med.*, 60(6), 1362–1371.
- Badve, C., et al. (2017). MR Fingerprinting of Adult Brain Tumors: Initial Experience. *American Journal of Neuroradiology*, *38*(3), 492–499.
- Ballesteros, M. C., Hansen, P. E., & Soila, K. (1993). MR imaging of the developing human brain. Part 2. Postnatal development. *RadioGraphics*, *13*(3), 611–622.
- Barral, J. K., et al. (2010). A robust methodology for in vivo T1 mapping. *Magn. Reson. Med.*, 64(4), 1057–1067.
- Bartlett, J. W., & Frost, C. (2008). Reliability, repeatability and reproducibility: analysis of measurement errors in continuous variables. Ultrasound in Obstetrics & Gynecology, 31(4), 466–475.
- Baudrexel, S., Nürnberger, L., Rüb, U., Seifried, C., Klein, J. C., Deller, T., Steinmetz, H., Deichmann, R., & Hilker, R. (2010). Quantitative mapping of T1 and T2* discloses nigral and brainstem pathology in early Parkinson's disease. *NeuroImage*, 51(2), 512–520.
- Belliveau, J. W., et al. (1990). Functional cerebral imaging by susceptibility-contrast NMR. *Magn. Reson. Med.*, *14*(3), 538–546.
- Bernarding, J., et al. (2000). Histogram-based characterization of healthy and ischemic brain tissues using multiparametric MR imaging including apparent diffusion coefficient maps and relaxometry. *Magn. Reson. Med.*, 43(1), 52–61.
- Bernstein, M. A., King, K. F., & Zou, X. J. (2004). Handbook of MRI pulse sequences. London: Elsevier Academic press.
- Bhatia, K. K., Hajnal, J., Hammers, A., & Rueckert, D. (2007). Similarity metrics for groupwise non-rigid registration. In N. Ayache, S. Ourselin, & A. Maeder (Eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*, (pp. 544–552). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Biglands, J. D., et al. (2020). MRI in acute muscle tears in athletes: can quantitative T2 and DTI predict return to play better than visual assessment? *European Radiology*, *30*(12), 6603–6613.
- Bladt, P. (2020). Accurate and precise perfusion parameter estimation in pseudo-continuous arterial spin labeling MRI. Phd thesis, University of Antwerp.
- Bladt, P., den Dekker, A. J., Clement, P., Achten, E., & Sijbers, J. (2020a). The costs and benefits of estimating T1 of tissue alongside cerebral blood flow and arterial transit time in pseudo-continuous arterial spin labeling. *NMR Biomed.*, *33*(12), e4182.

- Bladt, P., van Osch, M. J. P., Clement, P., Achten, E., Sijbers, J., & den Dekker, A. J. (2020b). Supporting measurements or more averages? How to quantify cerebral blood flow most reliably in 5 minutes by arterial spin labeling. *Magn. Reson. Med.*, 84(5), 2523–2536.
- Bojorquez, J. Z., Bricq, S., Acquitter, C., Brunotte, F., Walker, P. M., & Lalande, A. (2017). What are normal relaxation times of tissues at 3T? *Magn. Reson. Imaging*, *35*, 69–80.
- Bookwalter, C. A., Griswold, M. A., & Duerk, J. L. (2010). Multiple Overlapping k-Space Junctions for Investigating Translating Objects (MOJITO). *IEEE Trans. Med. Imag.*, 29(2), 339–349.
- Bottomley, P. A., Redington, R. W., Edelstein, W. A., & Schenck, J. F. (1985). Estimating radiofrequency power deposition in body NMR imaging. *Magn. Reson. Med.*, 2(4), 336–349.
- Bottomley, P. A., et al. (1987). A review of 1H nuclear magnetic resonance relaxation in pathology: Are T1 and T2 diagnostic? *Medical Physics*, 14(1), 1–37.
- Boudreau, M., Keenan, K. E., & Stikov, N. (2020). Chapter 2 Quantitative T1 and T1 ρ Mapping. In *Quantitative Magnetic Resonance Imaging*, vol. 1 of *Advances in Magnetic Resonance Technology* and *Applications*, (pp. 19–45). Academic Press.
- Breuer, F. A., Kannengiesser, S. A., Blaimer, M., Seiberlich, N., Jakob, P. M., & Griswold, M. A. (2009). General formulation for quantitative G-factor calculation in GRAPPA reconstructions. *Magn. Reson. Med.*, 62(3), 739–746.
- Bron, E. E., et al. (2013). Image registration improves human knee cartilage T1 mapping with delayed gadolinium-enhanced MRI of cartilage (dGEMRIC). *European Radiology*, 23(1), 246–252.
- Brück, W., Bitsch, A., Kolenda, H., Brück, Y., Stiefel, M., & Lassmann, H. (1997). Inflammatory central nervous system demyelination: Correlation of magnetic resonance imaging findings with lesion pathology. *Annals of Neurology*, 42(5), 783–793.
- Buxton, R. B., Frank, L. R., Wong, E. C., Siewert, B., Warach, S., & Edelman, R. R. (1998). A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magn. Reson. Med.*, 40(3), 383–396.
- Calamante, F., Gadian, D., & Connelly, A. (2002). Quantification of perfusion using bolus tracking magnetic resonance imaging in stroke. *Stroke*, *33*(4), 1146–1151.
- Callaghan, M. F., et al. (2015). An evaluation of prospective motion correction (PMC) for high resolution quantitative MRI. *Front. Neurosci.*, *9*(97), 1–9.
- Cardoso, M. J., Modat, M., Vercauteren, T., & Ourselin, S. (2015). Scale factor point spread function matching: Beyond aliasing in image resampling. In N. Navab, J. Hornegger, W. M. Wells, & A. Frangi (Eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Lecture Notes in Computer Science, (pp. 675–683). Cham: Springer International Publishing.
- Cashmore, M. T., McCann, A. J., Wastling, S. J., McGrath, C., Thornton, J., & Hall, M. G. (2021). Clinical quantitative MRI and the need for metrology. *The British Journal of Radiology*, *94*(1120), 20201215.
- Chappell, M. A., Groves, A. R., MacIntosh, B. J., Donahue, M. J., Jezzard, P., & Woolrich, M. W. (2011). Partial volume correction of multiple inversion time arterial spin labeling MRI data. *Magn. Reson. Med.*, 65(4), 1173–1183.
- Chappell, M. A., McConnell, F. A. K., Golay, X., Günther, M., Hernandez-Tamames, J. A., van Osch, M. J., & Asllani, I. (2021). Partial volume correction in arterial spin labeling perfusion MRI: A method to disentangle anatomy from physiology or an analysis step too far? *NeuroImage*, 238, 118236.
- Chekhonin, I., et al. (2023). Magnetic resonance relaxometry in quantitative imaging of brain gliomas: A literature review. *Neuroradiol. J.*
- Chen, N.-k., & Wyrwicz, A. M. (1999). Removal of intravoxel dephasing artifact in gradient-echo images using a field-map based RF refocusing technique. *Magnetic Resonance in Medicine*, 42(4), 807–812.
- Cheng, H. M., Stikov, N., Ghugre, N. R., & Wright, G. A. (2012). Practical medical applications of quantitative MR relaxometry. *Journal of Magnetic Resonance Imaging*, *36*(4), 805–824.
- Cho, S., Jones, D., Reddick, W. E., Ogg, R. J., & Steen, R. G. (1997). Establishing norms for age-related changes in proton T1 of human brain tissue in vivo. *Magnetic Resonance Imaging*, 15(10), 1133–1143.
- Christensen, K. A., Grant, D. M., Schulman, E. M., & Walling, C. (1974). Optimal determination of relaxation times of fourier transform nuclear magnetic resonance. Determination of spin-lattice relaxation times in chemically polarized species. *The Journal of Physical Chemistry*, 78, 1971–1977.

Clement, P., et al. (2022a). A Beginner's Guide to Arterial Spin Labeling (ASL) Image Processing.

Front. Radio., 2.

- Clement, P., et al. (2022b). ASL-BIDS, the brain imaging data structure extension for arterial spin labeling. *Scientific Data*, 9(1), 543.
- Coupé, P., et al. (2015). Detection of Alzheimer's disease signature in MR images seven years before conversion to dementia: Toward an early individual prognosis. *Hum. Brain. Mapp.*, 36(12), 4758–4770.
- Crawley, A. P., & Henkelman, R. M. (1988). A comparison of one-shot and recovery methods in T1 imaging. *Magn. Reson. Med.*, 7(1), 23–34.
- Cummings, J., et al. (2022). Aducanumab: Appropriate Use Recommendations Update. J. Prev. Alzheimers Dis., 9(2), 221–230.
- Dai, W., Garcia, D., de Bazelaire, C., & Alsop, D. C. (2008). Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magn. Reson. Med.*, 60(6), 1488–1497.
- Deichmann, R., Hahn, D., & Haase, A. (1999). Fast T1 mapping on a whole-body scanner. *Magn. Reson. Med.*, 42(1), 206–209.
- den Dekker, A. J., & Sijbers, J. (2014). Data distributions in magnetic resonance images: A review. *Physica Med.*, 30(7), 725–741.
- Deoni, S. C. (2007). High-resolution T1 mapping of the brain at 3T with driven equilibrium single pulse observation of T1 with high-speed incorporation of RF field inhomogeneities (DESPOT1-HIFI). *Journal of Magnetic Resonance Imaging*, *26*(4), 1106–1111.
- Deoni, S. C. (2011). Magnetic Resonance Relaxation and Quantitative Measurement in the Brain, (pp. 65–108). Totowa, NJ: Humana Press.
- Deoni, S. C., et al. (2011). Mapping infant brain myelination with magnetic resonance imaging. J. Neurosci., 31(2), 784–791.
- Deoni, S. C. L. (2010). Quantitative relaxometry of the brain. *Top. Magn. Reson. Imaging*, 21(2), 101–113.
- Deoni, S. C. L., Peters, T. M., & Rutt, B. K. (2005). High-resolution T1 and T2 mapping of the brain in a clinically acceptable time with DESPOT1 and DESPOT2. *Magn. Reson. Med.*, 53(1), 237–241.
- Deoni, S. C. L., Rutt, B. K., & Peters, T. M. (2003). Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Magn. Reson. Med.*, 49(3), 515–526.
- Detre, J. A., Leigh, J. S., Williams, D. S., & Koretsky, A. P. (1992). Perfusion imaging. Magn. Reson. Med., 23(1), 37–45.
- Detre, J. A., Rao, H., Wang, D. J., Chen, Y. F., & Wang, Z. (2012). Applications of arterial spin labeled MRI in the brain. *J. Magn. Reson. Imaging*, *35*(5), 1026–1037.
- Dobbing, J., & Sands, J. (1973). Quantitative growth and development of human brain. Arch. Dis. Child., 48(10), 757–767.
- Dortch, R. D. (2020). Chapter 3 Quantitative T₂ and T₂^{*} Mapping. In Quantitative Magnetic Resonance Imaging, vol. 1 of Advances in Magnetic Resonance Technology and Applications, (pp. 47–64). Academic Press.
- Dowell, N. G., & Tofts, P. S. (2007). Fast, accurate, and precise mapping of the RF field in vivo using the 180° signal null. *Magnetic Resonance in Medicine*, 58(3), 622–630.
- Drain, L. E. (1949). A direct method of measuring nuclear spin-lattice relaxation times. Proc. Phys. Soc. London, Sect. A, 62, 301–306.
- Dustin, D., Hall, B. M., Annapragada, A., & Pautler, R. G. (2016). Neuroimaging in Alzheimer's disease: preclinical challenges toward clinical efficacy. *Translational Research*, 175, 37–53.
- Edelstein, W. A., Glover, G. H., Hardy, C. J., & Redington, R. W. (1986). The intrinsic signal-to-noise ratio in NMR imaging. *Magn. Reson. Med.*, 3(4), 604–618.
- Erkinjuntti, T., Ketonen, L., Sulkava, R., Vuorialho, M., & Livanainen, M. (1987). Do white matter changes on MRI and CT differentiate vascular dementia from Alzheimer's disease? J. Neurol. Neurosurg. Psychiatry, 50(1), 37–42.
- ESR (2020). European Imaging Biomarkers Alliance. Available from: https://www.myesr.org/ research/eiball/.
- Feinberg, D. A., Mills, C. M., Posin, J. P., Ortendahl, D. A., Hylton, N. M., Crooks, L. E., Watts, J. C., Kaufman, L., Arakawa, M., Hoenninger, J. C., & Brant-Zawadzski, M. (1985). Multiple spin-echo magnetic resonance imaging. *Radiology*, 155(2), 437–442.

- Fernández-Seara, M. A., et al. (2005). Continuous arterial spin labeling perfusion measurements using single shot 3D GRASE at 3T. *Magnetic Resonance in Medicine*, *54*(5), 1241–1247.
- Ferrie, J. C., et al. (1999). MR assessment of the brain maturation during the perinatal period: quantitative T2 MR study in premature newborns. *Magnetic Resonance Imaging*, *17*(9), 1275–1288.
- Fingerhut, H., Gozdas, E., & Hosseini, S. M. H. (2022). Quantitative MRI Evidence for Cognitive Reserve in Healthy Elders and Prodromal Alzheimer's Disease. J. Alzheimers. Dis., 89(3), 849–863.
- Garcia, D. M., Duhamel, G., & Alsop, D. C. (2005). Efficiency of inversion pulses for background suppressed arterial spin labeling. *Magnetic Resonance in Medicine*, 54(2), 366–372.
- Gelman, A., Carlin, J., Stern, H., & Rubin, D. (1995). Bayesian Data Analysis. Chapman and Hall/CRC.
- George, A., Kuzniecky, R., Rusinek, H., Pardoe, H. R., & for the Human Epilepsy Project Investigators (2020). Standardized Brain MRI Acquisition Protocols Improve Statistical Power in Multicenter Quantitative Morphometry Studies. *J. Neuroimaging*, *30*(1), 126–133.
- Godenschweger, F., et al. (2016). Motion correction in MRI of the brain. *Phys. Med. Biol.*, 61(5), R32–R56.
- Gonzalez-At, J. B., Alsop, D. C., & Detre, J. A. (2000). Cerebral perfusion and arterial transit time changes during task activation determined with continuous arterial spin labeling. *Magn. Reson. Med.*, 43(5), 739–746.
- Gozdas, E., Fingerhut, H., Wu, H., Bruno, J. L., Dacorro, L., Jo, B., O'Hara, R., Reiss, A. L., & Hosseini, S. M. H. (2021). Quantitative measurement of macromolecular tissue properties in white and gray matter in healthy aging and amnestic MCI. *NeuroImage*, *237*, 118161.
- Gracien, R.-M., et al. (2016). Changes and variability of proton density and T1 relaxation times in early multiple sclerosis: MRI markers of neuronal damage in the cerebral cortex. *European Radiology*, 26(8), 2578–2586.
- Grade, M., et al. (2015). A neuroradiologist's guide to arterial spin labeling MRI in clinical practice. *Neuroradiology*, 57(12), 1181–1202.
- Griswold, M. A., et al. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.*, 47(6), 1202–1210.
- Gruenebach, N., et al. (2023). Clinical feasibility and validation of the accelerated T2 mapping sequence GRAPPATINI in brain imaging. *Heliyon*, 9(4), e15064.
- Gudbjartsson, H., & Patz, S. (1995). The Rician distribution of noisy MRI data. *Magn. Reson. Med.*, 34(6), 910–914.
- Gunter, J. L., et al. (2009). Measurement of MRI scanner performance with the ADNI phantom. *Medical Physics*, *36*(6, Part1), 2193–2205.
- Günther, M., Oshio, K., & Feinberg, D. A. (2005). Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magn. Reson. Med.*, *54*(2), 491–498.
- Gupta, R. K. (1977). A new look at the method of variable nutation angle for the measurement of spin-lattice relaxation times using fourier transform NMR. *Journal of Magnetic Resonance (1969)*, 25(1), 231–235.
- Gur, D., Good, W. F., Wolfson, S. K., Yonas, H., & Shabason, L. (1982). In Vivo Mapping of Local Cerebral Blood Flow by Xenon-Enhanced Computed Tomography. *Science*, 215(4537), 1267–1268.
- Guyader, J.-M., Bernardin, L., Douglas, N. H., Poot, D. H., Niessen, W. J., & Klein, S. (2015). Influence of image registration on apparent diffusion coefficient images computed from free-breathing diffusion mr images of the abdomen. *J. Magn. Reson. Imaging*, *42*(2), 315–330.
- Haase, A., Matthaei, D., Bartkowski, R., Dühmke, E., & Leibfritz, D. (1989). Inversion recovery snapshot FLASH MR imaging. *J. Comput. Assist. Tomogr.*, *13*(6), 1036–1040.
- Hagiwara, A., Fujita, S., Ohno, Y., & Aoki, S. (2020). Variability and Standardization of Quantitative Imaging: Monoparametric to Multiparametric Quantification, Radiomics, and Artificial Intelligence. *Investigative Radiology*, *55*(9), 601–616.
- Hahn, E. L. (1949). An accurate nuclear magnetic resonance method for measuring spin-lattice relaxation times. *Phys. Rev.*, *76*(1), 145–146.
- Hallack, A., Chappell, M. A., Gooding, M. J., & Schnabel, J. A. (2014). A New Similarity Metric for Groupwise Registration of Variable Flip Angle Sequences for Improved T10 Estimation in DCE-MRI. In S. Ourselin, & M. Modat (Eds.) *Biomedical Image Registration*, (pp. 154–163). Springer International Publishing.
- Herculano-Houzel, S. (2009). The human brain in numbers: a linearly scaled-up primate brain. Frontiers in Human Neuroscience, 3.
- Hernandez-Garcia, L., Lahiri, A., & Schollenberger, J. (2019). Recent progress in ASL. *NeuroImage*, 187, 3–16.
- Hernandez-Garcia, L., et al. (2022). Recent Technical Developments in ASL: A Review of the State of the Art. *Magn. Reson. Med.*, *88*(5), 2021–2042.
- Herscovitch, P., Markham, J., & Raichle, M. E. (1983). Brain Blood Flow Measured with Intravenous H2(15)O. Journal of Nuclear Medicine, 24(9), 782–789.
- Hilbert, T., et al. (2018). Accelerated T2 mapping combining parallel MRI and model-based reconstruction: GRAPPATINI. J. Magn. Reson. Imaging, 48(2), 359–368.
- Homer, J., & Beevers, M. S. (1985). Driven-equilibrium single-pulse observation of T1 relaxation. A reevaluation of a rapid "new" method for determining NMR spin-lattice relaxation times. *Journal of Magnetic Resonance* (1969), 63(2), 287–297.
- Huang, H., Zhang, J., Wakana, S., Zhang, W., Ren, T., Richards, L. J., Yarowsky, P., Donohue, P., Graham, E., van Zijl, P. C. M., & Mori, S. (2006). White and gray matter development in human fetal, newborn and pediatric brains. *NeuroImage*, *33*(1), 27–38.
- Hughes, J. R. (2007). Autism: The first firm finding = underconnectivity? *Epilepsy & Behavior*, 11(1), 20–24.
- Huizinga, W., et al. (2016). PCA-based groupwise image registration for quantitative MRI. *Medical Image Analysis*, *29*, 65–78.
- Jackson, G. D., Connelly, A., Duncan, J. S., Grünewald, R. A., & Gadian, D. G. (1993). Detection of hippocampal pathology in intractable partial epilepsy: increased sensitivity with quantitative magnetic resonance T2 relaxometry. *Neurology*, 43(9), 1793–1799.
- Jain, S., et al. (2015). Automatic segmentation and volumetry of multiple sclerosis brain lesions from MR images. *NeuroImage: Clinical*, *8*, 367–375.
- Jain, S., et al. (2016). Two Time Point MS Lesion Segmentation in Brain MRI: An Expectation-Maximization Framework. *Frontiers in Neuroscience*, 10.
- Jezzard, P., Chappell, M. A., & Okell, T. W. (2018). Arterial spin labeling for the measurement of cerebral perfusion and angiography. *Journal of Cerebral Blood Flow & Metabolism*, 38(4), 603–626.
- Johnson, N. A., et al. (2005). Pattern of Cerebral Hypoperfusion in Alzheimer Disease and Mild Cognitive Impairment Measured with Arterial Spin-labeling MR Imaging: Initial Experience. *Radiology*, 234(3), 851–859.
- Just, M., & Thelen, M. (1988). Tissue characterization with T1, T2, and proton density values: results in 160 patients with brain tumors. *Radiology*, *169*(3), 779–785.
- Kaptein, R., Dijkstra, K., & Tarr, C. E. (1976). A single-scan fourier transform method for measuring spin-lattice relaxation times. J. Magn. Reson. Imaging, 24(2), 295–300.
- Karakuzu, A., et al. (2022). qMRI-BIDS: An extension to the brain imaging data structure for quantitative magnetic resonance imaging data. Scientific Data, 9(1), 517.
- Keenan, K., Stupic, K., Boss, M., Russek, S., Chenevert, T., Prasad, P., Reddick, W., Zheng, J., Hu, P., & Jackson, E. (2016). Comparison of T1 measurement using ISMRM/NIST system phantom. In *Proc. Intl. Soc. Magn. Reson. Med.*, 3290.
- Keenan, K. E., et al. (2019). Recommendations towards standards for quantitative MRI (qMRI) and outstanding needs. *Journal of Magnetic Resonance Imaging*, 49(7), e26–e39.
- Keenan, K. E., et al. (2021). Multi-site, multi-platform comparison of MRI T1 measurement using the system phantom. PLOS ONE, 16(6), 1–19.
- Kerleroux, B., Kober, T., Hilbert, T., Serru, M., Philippe, J., Sirinelli, D., & Morel, B. (2019). Clinical equivalence assessment of T2 synthesized pediatric brain magnetic resonance imaging. *J. Neuroradiol.*, 46(2), 130–135.
- Kim, S.-G. (1995). Quantification of relative cerebral blood flow change by flow-sensitive alternating inversion recovery (FAIR) technique: Application to functional mapping. *Magn. Reson. Med.*, 34(3), 293–301.
- Kingsley, P. B., Ogg, R. J., Reddick, W. E., & Steen, R. G. (1998). Correction of errors caused by imperfect inversion pulses in MR imaging measurement of T1 relaxation times. *Magnetic Resonance Imaging*, 16(9), 1049–1055.

- Knight, M. J., Wearn, A., Coulthard, E., & Kauppinen, R. A. (2019). T2 Relaxometry and Diffusion Tensor Indices of the Hippocampus and Entorhinal Cortex Improve Sensitivity and Specificity of MRI to Detect Amnestic Mild Cognitive Impairment and Alzheimer's Disease Dementia. J. Magn. Reson. Imaging, 49(2), 445–455.
- Krüger, G., & Glover, G. H. (2001). Physiological noise in oxygenation-sensitive magnetic resonance imaging. Magn. Reson. Med., 46(4), 631–637.
- Kwong, K. K., et al. (1995). MR perfusion studies with T1-weighted echo planar imaging. Magn. Reson. Med., 34(6), 878–887.
- Larsson, H. B. W., et al. (1989). Assessment of demyelination, edema, and gliosis by in vivo determination of T1 and T2 in the brain of patients with acute attack of multiple sclerosis. *Magn. Reson. Med.*, *11*(3), 337–348.
- Le Bihan, D., et al. (1986). Mr imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, *161*(2), 401–407.
- Liang, X., Connelly, A., & Calamante, F. (2013). Improved partial volume correction for single inversion time arterial spin labeling data. *Magn. Reson. Med.*, 69(2), 531–537.
- Liu, C., Bammer, R., Kim, D.-h., & Moseley, M. E. (2004). Self-navigated interleaved spiral (SNAILS): Application to high-resolution diffusion tensor imaging. *Magn. Reson. Med.*, *52*(6), 1388–1396.
- Loktyushin, A., Nickisch, H., Pohmann, R., & Schölkopf, B. (2013). Blind retrospective motion correction of MR images. *Magn. Reson. Med.*, 70(6), 1608–1618.
- Look, D. C., & Locker, R. R. (1970). Time saving in measurement of NMR and EPR relaxation times. *Rev. Sci. Instrum.*, 41(2), 250–251.
- Lu, H., Clingman, C., Golay, X., & van Zijl, P. C. (2004). Determining the longitudinal relaxation time (T1) of blood at 3.0 Tesla. *Magn. Reson. Med.*, *52*(3), 679–682.
- Ma, D., Gulani, V., Seiberlich, N., Liu, K., Sunshine, J. L., Duerk, J. L., & Griswold, M. A. (2013). Magnetic resonance fingerprinting. *Nature*, 495(7440), 187–192.
- Maclaren, J., Herbst, M., Speck, O., & Zaitsev, M. (2013). Prospective motion correction in brain imaging: A review. *Magn. Reson. Med.*, 69(3), 621–636.
- Maes, F., et al. (1997). Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, *16*(2), 187–198.
- Maleki, N., Dai, W., & Alsop, D. C. (2012). Optimization of background suppression for arterial spin labeling perfusion imaging. *MAGMA*, 25(2), 127–133.
- Mangin, J.-F., Poupon, C., Clark, C., Le Bihan, D., & Bloch, I. (2002). Distortion correction and robust tensor estimation for MR diffusion imaging. *Medical Image Analysis*, 6(3), 191–198. Special Issue on Medical Image Computing and Computer-Assisted Intervention - MICCAI 2001.
- Marques, J. P., et al. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *NeuroImage*, *49*(2), 1271–1281.
- Matzat, S. J., van Tiel, J., Gold, G. E., & Oei, E. H. G. (2013). Quantitative MRI techniques of cartilage composition. *Quantitative Imaging in Medicine and Surgery*, 3(3), 162–174.
- Mehranian, A., McGinnity, C. J., Neji, R., Prieto, C., Hammers, A., De Vita, E., & Reader, A. J. (2020). Motion-corrected and high-resolution anatomically assisted (MOCHA) reconstruction of arterial spin labeling MRI. *Magn. Reson. Med.*, 84(3), 1306–1320.
- Metz, C., Klein, S., Schaap, M., van Walsum, T., & Niessen, W. (2011). Nonrigid registration of dynamic medical imaging data using nD+t B-splines and a groupwise optimization approach. *Medical Image Analysis*, 15(2), 238–249.
- Meurée, C., Maurel, P., Ferré, J.-C., & Barillot, C. (2019). Patch-based super-resolution of arterial spin labeling magnetic resonance images. *NeuroImage*, *189*, 85–94.
- Mezer, A., et al. (2013). Quantifying the local tissue volume and composition in individual brains with magnetic resonance imaging. *Nature Medicine*, 19(12), 1667–1672.
- Möllenhoff, K., Shah, N. J., Pracht, E. D., & Stöcker, T. (2010). Accelerated mapping of T1 relaxation times using TAPIR. In Proc. Intl. Soc. Magn. Reson. Med., vol. 18, (p. 2951).
- Mugler III, J. P., & Brookeman, J. R. (1990). Three-dimensional magnetization-prepared rapid gradientecho imaging (3D MP RAGE). *Magn. Reson. Med.*, 15(1), 152–157.
- Mutsaerts, H. J. M. M., Steketee, R. M. E., Heijtel, D. F. R., Kuijer, J. P. A., van Osch, M. J. P., Majoie, C. B. L. M., Smits, M., & Nederveen, A. J. (2014). Inter-Vendor Reproducibility of

Pseudo-Continuous Arterial Spin Labeling at 3 Tesla. PLOS ONE, 9(8), 1-10.

- Nachmani, A., Schurr, R., Joskowicz, L., & Mezer, A. A. (2019). The effect of motion correction interpolation on quantitative T1 mapping with MRI. *Med. Image Anal.*, *52*, 119–127.
- Novikov, D. S., Kiselev, V. G., & Jespersen, S. N. (2018). On modeling. *Magn. Reson. Med.*, 79(6), 3172–3193.
- Noyes, K., & Weinstock-Guttman, B. (2013). Impact of diagnosis and early treatment on the course of multiple sclerosis. *Am. J. Manag. Care*, *19*((17 Suppl)), s321–s331.
- Okujava, M., Schulz, R., Ebner, A., & Woermann, F. G. (2002). Measurement of temporal lobe T2 relaxation times using a routine diagnostic MR imaging protocol in epilepsy. *Epilepsy Research*, 48(1), 131–142.
- Oliveira, I. A., Guimaraes, T. M., Souza, R. M., dos Santos, A. C., Machado-de Sousa, J. P., Hallak, J. E., & Leoni, R. F. (2018). Brain functional and perfusional alterations in schizophrenia: an arterial spin labeling study. *Psychiatry Research: Neuroimaging*, 272, 71–78.
- Oshio, K., & Feinberg, D. A. (1991). GRASE (Gradient-and Spin-Echo) imaging: A novel fast MRI technique. *Magn. Reson. Med.*, 20(2), 344–349.
- Østergaard, L., Sorensen, A. G., Kwong, K. K., Weisskoff, R. M., Gyldensted, C., & Rosen, B. R. (1996a). High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part II: Experimental comparison and preliminary results. *Magn. Reson. Med.*, *36*(5), 726–736.
- Østergaard, L., Weisskoff, R. M., Chesler, D. A., Gyldensted, C., & Rosen, B. R. (1996b). High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part I: Mathematical approach and statistical analysis. *Magn. Reson. Med.*, *36*(5), 715–725.
- Özsarlak, O., Van Goethem, J. W., Maes, M., & Parizel, P. M. (2004). MR angiography of the intracranial vessels: technical aspects and clinical applications. *Neuroradiology*, *46*(12), 955–972.
- Papadopoulos, K., Tozer, D. J., Fisniku, L., Altmann, D. R., Davies, G., Rashid, W., Thompson, A. J., Miller, D. H., & Chard, D. T. (2010). T1-relaxation time changes over five years in relapsing-remitting multiple sclerosis. *Multiple Sclerosis Journal*, 16(4), 427–433.
- Parkes, L. M., Rashid, W., Chard, D. T., & Tofts, P. S. (2004). Normal cerebral perfusion measurements using arterial spin labeling: Reproducibility, stability, and age and gender effects. *Magn. Reson. Med.*, 51(4), 736–743.
- Parkes, L. M., & Tofts, P. S. (2002). Improved accuracy of human cerebral blood perfusion measurements using arterial spin labeling: Accounting for capillary water permeability. *Magn. Reson. Med.*, 48(1), 27–41.
- Parry, A., Clare, S., Jenkinson, M., Smith, S., Palace, J., & Matthews, P. M. (2002). White matter and lesion T1 relaxation times increase in parallel and correlate with disability in multiple sclerosis. *Journal of Neurology*, 249(9), 1279–1286.
- Paschoal, A. M., Leoni, R. F., Pastorello, B. F., & van Osch, M. J. P. (2021). Three-dimensional gradient and spin-echo readout for time-encoded pseudo-continuous arterial spin labeling: Influence of segmentation factor and flow compensation. *Magn. Reson. Med.*, 86(3), 1454–1462.
- Paus, T., Collins, D. L., Evans, A. C., Leonard, G., Pike, B., & Zijdenbos, A. (2001). Maturation of white matter in the human brain: a review of magnetic resonance studies. *Brain Research Bulletin*, 54(3), 255–266.
- Pekar, J., Ligeti, L., Ruttner, Z., Lyon, R. C., Sinnwell, T. M., van Gerderen, P., Fiat, D., Moonen, C. T. W., & McLaughlin, A. C. (1991). In Vivo measurement of cerebral oxygen consumption and blood flow using ¹⁷O magnetic resonance imaging. *Magn. Reson. Med.*, 21(2), 313–319.
- Pendse, N., Wissmeyer, M., Altrichter, S., Vargas, M., Delavelle, J., Viallon, M., Federspiel, A., Seeck, M., Schaller, K., & Lövblad, K. O. (2010). Interictal arterial spin-labeling MRI perfusion in intractable epilepsy. *Journal of Neuroradiology*, 37(1), 60–63.
- Petersen, E. T., Mouridsen, K., & Golay, X. (2010). The QUASAR reproducibility study, Part II: Results from a multi-center Arterial Spin Labeling test–retest study. *NeuroImage*, 49(1), 104–113.
- Petersen, E. T., Zimine, I., Ho, Y.-C. L., & Golay, X. (2006). Non-invasive measurement of perfusion: a critical review of arterial spin labelling techniques. *The British Journal of Radiology*, 79(944), 688–701.
- Petr, J., Mutsaerts, H. J. M. M., De Vita, E., Steketee, R. M. E., Smits, M., Nederveen, A. J., Hofheinz, F., van den Hoff, J., & Asllani, I. (2018). Effects of systematic partial volume errors on the estimation of gray matter cerebral blood flow with arterial spin labeling MRI. *Magn. Reson. Mater. Phys. Biol. Med.*, 31(6), 725–734.

Qin, Q., et al. (2022). Velocity-selective arterial spin labeling perfusion MRI: A review of the state of the art and recommendations for clinical implementation. *Magn. Reson. Med.*, 88(4), 1528–1547.

Raichle, M. E. (2006). The Brain's Dark Energy. Science, 314(5803), 1249-1250.

- Ramos-Llordén, G., et al. (2017). A unified maximum likelihood framework for simultaneous motion and T_1 estimation in quantitative MR T_1 mapping. *IEEE Trans. Med. Imag.*, 36(2), 433–446.
- Rempp, K. A., Brix, G., Wenz, F., Becker, C. R., Gückel, F., & Lorenz, W. J. (1994). Quantification of regional cerebral blood flow and volume with dynamic susceptibility contrast-enhanced MR imaging. *Radiology*, 193(3), 637–641.
- Reutens, D. C., Stevens, J. M., Kingsley, D., Kendall, B., Moseley, I., Cook, M. J., Free, S., Fish, D. R., & Shorvon, S. D. (1996). Reliability of visual inspection for detection of volumetric hippocampal asymmetry. *Neuroradiology*, 38(3), 221–225.
- Rodionov, R., Bartlett, P. A., He, C., Vos, S. B., Focke, N. K., Ourselin, S. G., & Duncan, J. S. (2015). T2 mapping outperforms normalised FLAIR in identifying hippocampal sclerosis. *NeuroImage: Clinical*, *7*, 788–791.
- Roosendaal, S. D., Geurts, J. J. G., Vrenken, H., Hulst, H., Cover, K. S., Castelijns, J. A., Pouwels, P. J. W., & Barkhof, F. (2009). Regional DTI differences in multiple sclerosis patients. *NeuroImage*, 44(4), 1397–1403.
- Roux, M., Hilbert, T., Hussami, M., Becce, F., Kober, T., & Omoumi, P. (2019). MRI T2 Mapping of the Knee Providing Synthetic Morphologic Images: Comparison to Conventional Turbo Spin-Echo MRI. *Radiology*, 293(3), 620–630.
- Rowe, J. W. (2021). The US eldercare workforce is falling further behind. Nature Aging, 1(4), 327-329.
- RSNA (2020). Quantitative Imaging Biomarkers Alliance. Available from: https://www.rsna.org/en/ research/quantitative-imaging-biomarkers-alliance.
- Sanchez Panchuelo, R. M., Mougin, O., Turner, R., & Francis, S. T. (2021). Quantitative T1 mapping using multi-slice multi-shot inversion recovery EPI. *NeuroImage*, 234, 117976.
- Scheltens, P., Blennow, K., Breteler, M. M. B., de Strooper, B., Frisoni, G. B., Salloway, S., & Van der Flier, W. M. (2016). Alzheimer's disease. *The Lancet*, 388(10043), 505–517.
- Schmitt, P., Griswold, M. A., Jakob, P. M., Kotas, M., Gulani, V., Flentje, M., & Haase, A. (2004). Inversion recovery TrueFISP: Quantification of T1, T2, and spin density. *Magn. Reson. Med.*, 51(4), 661–667.
- Shah, N. J., Zaitsev, M., Steinhoff, S., & Zilles, K. (2001). A New Method for Fast Multislice T1 Mapping. NeuroImage, 14(5), 1175–1185.
- Sima, D. M., Loeckx, D., Smeets, D., Jain, S., Parizel, P. M., & Van Hecke, W. (2017). Use Case I: Imaging Biomarkers in Neurological Disease. Focus on Multiple Sclerosis, (pp. 169–180). Springer International Publishing.
- Smith, N. A. S., Sinden, D., Thomas, S. A., Romanchikova, M., Talbott, J. E., & Adeogun, M. (2020). Building confidence in digital health through metrology. *The British Journal of Radiology*, 93(1109), 20190574.
- Stevenson, V. L., et al. (2000). Variations in T1 and T2 relaxation times of normal appearing white matter and lesions in multiple sclerosis. *Journal of the Neurological Sciences*, *178*(2), 81–87.
- Stikov, N., Trzasko, J. D., & Bernstein, M. A. (2019). Reproducibility and the future of MRI research. *Magn. Reson. Med.*, 82(6), 1981–1983.
- Stüber, C., Morawski, M., Schäfer, A., Labadie, C., Wähnert, M., Leuze, C., Streicher, M., Barapatre, N., Reimann, K., Geyer, S., Spemann, D., & Turner, R. (2014). Myelin and iron concentration in the human brain: A quantitative study of MRI contrast. *NeuroImage*, *93*, 95–106.
- Studler, U., White, L. M., Andreisek, G., Luu, S., Cheng, H.-L. M., & Sussman, M. S. (2010). Impact of motion on T1 mapping acquired with inversion recovery fast spin echo and rapid spoiled gradient recalled-echo pulse sequences for delayed gadolinium-enhanced MRI of cartilage (dGEMRIC) in volunteers. J. Magn. Reson. Imaging, 32(2), 394–398.
- Sullivan, D. C., et al. (2015). Metrology Standards for Quantitative Imaging Biomarkers. *Radiology*, 277(3), 813–825.
- Sumpf, T. J., Uecker, M., Boretius, S., & Frahm, J. (2011). Model-based nonlinear inverse reconstruction for T2 mapping using highly undersampled spin-echo MRI. J. Magn. Reson. Imaging, 34(2), 420–428.
- Tancredi, F. B., Gauthier, C. J., Madjar, C., Bolar, D. S., Fisher, J. A., Wang, D. J. J., & Hoge, R. D. (2012). Comparison of pulsed and pseudocontinuous arterial spin-labeling for measuring CO2-induced

cerebrovascular reactivity. J. Magn. Reson. Imaging, 36(2), 312-321.

- Tang, X., Cai, F., Ding, D.-X., Zhang, L.-L., Cai, X.-Y., & Fang, Q. (2018). Magnetic resonance imaging relaxation time in Alzheimer's disease. *Brain Research Bulletin*, 140, 176–189.
- Tofts, P. (2004). *Quantitative MRI of the brain: measuring changes caused by disease*. Chichester, England: John Wiley & Sons.
- Tourais, J., Coletti, C., & Weingärtner, S. (2022). Chapter 1 Brief Introduction to MRI Physics. In M. Akçakaya, M. Doneva, & C. Prieto (Eds.) Magnetic Resonance Image Reconstruction, vol. 7 of Advances in Magnetic Resonance Technology and Applications, (pp. 3–36). Academic Press.
- Townsend, T. N., Bernasconi, N., Pike, G. B., & Bernasconi, A. (2004). Quantitative analysis of temporal lobe white matter T2 relaxation time in temporal lobe epilepsy. *NeuroImage*, 23(1), 318–324.
- Tranfa, M., Pontillo, G., Petracca, M., Brunetti, A., Tedeschi, E., Palma, G., & Cocozza, S. (2022). Quantitative MRI in Multiple Sclerosis: From Theory to Application. Am. J. Neuroradiol., 43(12), 1688–1695.
- Tsialios, P., Thrippleton, M., Glatz, A., & Pernet, C. (2017). Evaluation of MRI sequences for quantitative T1 brain mapping. *Journal of Physics: Conference Series*, 931(1), 012038.
- van den Berg, M., et al. (2022). Altered basal forebrain function during whole-brain network activity at pre- and early-plaque stages of Alzheimer's disease in TgF344-AD rats. *Alzheimer's Research & Therapy*, 14(1), 148.
- van den Bos, A. (2007). *Parameter Estimation for Scientists and Engineers*. Hoboken, New Yersey, USA: John Wiley & Sons.
- van der Thiel, M., et al. (2018). Brain Perfusion Measurements Using Multidelay Arterial Spin-Labeling Are Systematically Biased by the Number of Delays. *American Journal of Neuroradiology*, *39*(8), 1432–1438.
- van Heeswijk, R. B., et al. (2012). Motion compensation strategies in magnetic resonance imaging. *Critical Reviews[™] in Biomedical Engineering*, 40(2), 99–119.
- van Osch, M. J. P., et al. (2018). Advances in arterial spin labelling MRI methods for measuring perfusion and collateral flow. *J. Cereb. Blood. Flow. Metab.*, 38(9), 1461–1480.
- Van Steenkiste, G., et al. (2016). Super-resolution reconstruction of diffusion parameters from diffusionweighted images with different slice orientations. *Magn. Reson. Med.*, 75(1), 181–195.
- Van Steenkiste, G., et al. (2017). Super-resolution T1 estimation: Quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. *Magn. Reson. Med.*, 77(5), 1818–1830.
- Vidorreta, M., Balteau, E., Wang, Z., De Vita, E., Pastor, M. A., Thomas, D. L., Detre, J. A., & Fernández-Seara, M. A. (2014). Evaluation of segmented 3D acquisition schemes for whole-brain high-resolution arterial spin labeling at 3T. *NMR Biomed.*, *27*(11), 1387–1396.
- Vidorreta, M., et al. (2013). Comparison of 2D and 3D single-shot ASL perfusion fMRI sequences. *NeuroImage*, *66*, 662–671.
- Villemagne, V. L., & Chételat, G. (2016). Neuroimaging biomarkers in Alzheimer's disease and other dementias. *Ageing Research Reviews*, 30, 4–16. Brain Imaging and Aging.
- Villringer, A., et al. (1988). Dynamic imaging with lanthanide chelates in normal brain: Contrast due to magnetic susceptibility effects. *Magn. Reson. Med.*, 6(2), 164–174.
- Voelker, M. N., et al. (2016). The traveling heads: multicenter brain imaging at 7 Tesla. Magnetic Resonance Materials in Physics, Biology and Medicine (MAGMA), 29, 399–415.
- Voelker, M. N., et al. (2021). The traveling heads 2.0: Multicenter reproducibility of quantitative imaging methods at 7 Tesla. *NeuroImage*, 232, 117910.
- Vrenken, H., et al. (2006). Whole-Brain T1 Mapping in Multiple Sclerosis: Global Changes of Normalappearing Gray and White Matter. *Radiology*, 240(3), 811–820.
- Vymazal, J., Righini, A., Brooks, R. A., Canesi, M., Mariani, C., Leonardi, M., & Pezzoli, G. (1999). T1 and T2 in the Brain of Healthy Subjects, Patients with Parkinson Disease, and Patients with Multiple System Atrophy: Relation to Iron Content. *Radiology*, 211(2), 489–495.
- Wachinger, C., & Navab, N. (2013). Simultaneous registration of multiple images: Similarity metrics and efficient optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(5), 1221–1233.

- Wang, D. J. J., et al. (2012). The Value of Arterial Spin-Labeled Perfusion Imaging in Acute Ischemic Stroke. Stroke, 43(4), 1018–1024.
- Wang, D. J. J., et al. (2013). Multi-delay multi-parametric arterial spin-labeled perfusion MRI in acute ischemic stroke — Comparison with dynamic susceptibility contrast enhanced perfusion imaging. *NeuroImage: Clinical*, *3*, 1–7.
- Wang, J., et al. (2003). Pediatric perfusion imaging using pulsed arterial spin labeling. *Journal of Magnetic Resonance Imaging*, 18(4), 404–413.
- Wang, Y., et al. (2011). Regional reproducibility of pulsed arterial spin labeling perfusion imaging at 3T. NeuroImage, 54(2), 1188–1195.
- Weigel, M., et al. (2007). Inversion recovery prepared turbo spin echo sequences with reduced SAR using smooth transitions between pseudo steady states. *Magn. Reson. Med.*, *57*(3), 631–637.
- Wells, W. M., Viola, P., Atsumi, H., Nakajima, S., & Kikinis, R. (1996). Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1), 35–51.
- Williams, D. S., Detre, J. A., Leigh, J. S., & Koretsky, A. P. (1992). Magnetic resonance imaging of perfusion using spin inversion of arterial water. *Proc. Natl. Acad. Sci. USA*, 89(1), 212–216.
- Winston, G. P., et al. (2017). Automated T2 relaxometry of the hippocampus for temporal lobe epilepsy. *Epilepsia*, *58*(9), 1645–1652.
- Wolf, R. L., & Detre, J. A. (2007). Clinical neuroimaging using arterial spin-labeled perfusion magnetic resonance imaging. *Neurotherapeutics*, 4(3), 346–359.
- Wong, E. C., Cronin, M., Wu, W.-C., Inglis, B., Frank, L. R., & Liu, T. T. (2006). Velocity-selective arterial spin labeling. *Magn. Reson. Med.*, 55(6), 1334–1341.
- Wong, E. C., et al. (1998). A theoretical and experimental comparison of continuous and pulsed arterial spin labeling techniques for quantitative perfusion imaging. *Magn. Reson. Med.*, 40(3), 348–355.
- World Health Organization, et al. (2017). Global action plan on the public health response to dementia 2017–2025. *World Health Organization*.
- Yarnykh, V. L. (2007). Actual flip-angle imaging in the pulsed steady state: A method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magn. Reson. Med.*, 57(1), 192–200.
- Yarnykh, V. L. (2010). Optimal radiofrequency and gradient spoiling for improved accuracy of T1 and B1 measurements using fast steady-state techniques. *Magn. Reson. Med.*, *63*(6), 1610–1626.
- Ye, F. Q., Frank, J. A., Weinberger, D. R., & McLaughlin, A. C. (2000). Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). *Magn. Reson. Med.*, 44(1), 92–100.
- Zhao, L., Chang, C.-D., & Alsop, D. C. (2018). Controlling T2 blurring in 3D RARE arterial spin labeling acquisition through optimal combination of variable flip angles and k-space filtering. *Magn. Reson. Med.*, *80*(4), 1391–1401.

4

Super-resolution reconstruction as prime protagonist for accelerated (q)MRI

Contents

4.1	Introd	uction	
	4.1.1	Resolution challenges in MRI	
	4.1.2	Resolution enhancement techniques	
	4.1.3	Super-resolution reconstruction	
4.2	Super-	resolution reconstruction applied to MRI	
	4.2.1	Concept and definition	
	4.2.2	Historical misconceptions	
	4.2.3	Key components	
		4.2.3.1 Acquisition strategies	
		4.2.3.2 Imaging model	
		4.2.3.3 Reconstruction	
4.3	The ne	eed for robust motion compensation	
Refer	rences		

Up to this point of this thesis, we have introduced the imperative for absolute quantification of biomarkers for neurodegenerative diseases, and coupled with that the need for technological advancements that can facilitate such quantification. In the preceding chapter, we introduced the strengths of quantitative MRI, and highlighted two pivotal qMRI applications: MR relaxometry and Arterial Spin Labeling. At the same time, it was discussed how the dissemination of quantitative MR imaging methods into clinical routine is complicated by the

fundamental consideration of **how to optimally balance spatial resolution, signal-to-noise ratio (SNR), and scan time** in these methods. In particular since those three imaging parameters are highly interdependent (Plenge et al., 2012): higher spatial resolution allows one to observe smaller details, but typically reduces SNR, and/or increases scan time. Yet, a certain minimum level of SNR is required to distinguish the signal of interest from system noise. In addition, minimizing scan time is paramount, as MRI resources are limited and costly, and long scan times are uncomfortable for the patient. Furthermore, long scan times increase the risk of motion artifacts and distortions in the images.

In this chapter, the concept of **super-resolution reconstruction (SRR)** is introduced, which offers the potential to balance the trade-off between spatial resolution, SNR, and scan time in MRI. As will be elucidated, this technique and its fundamental idea have been around since the early 1980's, finding applications in various imaging domains. However, as is often the case with pioneering scientific efforts, the initial proof-of-concept serves as a stepping stone to a multitude of new innovations that can surpass the original idea's potential, particularly when translated into practical societal applications. In light of this notion, the contributions in this thesis focus on advancing SRR in combination with qMRI and joint patient motion correction.

4.1 Introduction

4.1.1 Resolution challenges in MRI

The fundamental goal of MRI is to create detailed images of the internal structures of the body based on the signals emitted by hydrogen nuclei in an external magnetic field (cf. Chapter 2). While there is often a preference for obtaining a high-resolution threedimensional (3D) volumetric image for accurate medical diagnosis, **the limitations imposed by acquisition time constraints and hardware constraints can render direct full 3D acquisition infeasible or ineffective**. In such cases, it is common practice to acquire a set of two-dimensional (2D) slices, and combine these in a so called **2D multi-slice image**. The advantage of 2D multi-slice images, compared to full 3D acquisitions, is that it is possible to interleave the acquisition of slices. That is, while waiting for the relaxation of the magnetization of a slice, (a part of) the *k*-space of the other slices can be excited and recorded. In general, when the repetition time (TR) of a sequence is limited by the T_1 decay, it is possible to acquire 2D multi-slice images significantly faster than full 3D images with the same resolution (Zimmerman et al., 2000). Moreover, 2D multi-slice images might be less influenced by object motion.

Still, MRI acquisitions are bound to the trade-off between spatial resolution, SNR, and scan time. Acquiring a 2D multi-slice image at high resolution (i.e. smaller voxel size) might allow observation of finer details, but typically results in a reduction of the SNR, as the power of the signal scales approximately linearly with the imaged volume. Yet, a minimum level of SNR is required to distinguish the signal of interest from the noise. Improving SNR can be achieved through signal averaging across multiple acquisitions. However, this extends the acquisition time, incurring additional costs, causing discomfort for the patient, and inducing motion artifacts in the images. Additionally, the slice thickness in a 2D multi-slice image is constrained by the slice-selection pulse, which is determined by hardware limitations coupled

with pulse sequence timing considerations, making the acquisition of thin slices not always feasible. Consequently, **2D multi-slice images are often acquired with high in-plane resolution but lower resolution in the slice-selection (through-plane) direction**, resulting in anisotropic voxels (see Fig. 4.1). On the one hand, the acquisition of thick slices is beneficial as it increases the SNR, again since the signal scales approximately linearly with the imaged volume. On the other hand, however, thick slices give rise to **partial volume effects** (PVEs) (see Fig. 4.1), meaning that multiple tissues within an imaged voxel contribute to the observed signal of that respective voxel. In such instances, the voxel intensity not only depends on the pulse sequence and tissue characteristics but also on the proportions of each tissue type present in the voxel. Hence, these PVEs can introduce significant errors in quantitative brain measurements, such as estimates of brain (tissue) volumes (González Ballester et al., 2002).



Figure 4.1: Illustration of in-plane resolution and through-plane resolution in a 2D multi-slice acquisition. Column (a): A high-resolution image with isotropic resolution, column (b): a low-resolution multi-slice image acquired with slice selection along the *z*-axis, and column (c): a low-resolution multi-slice image acquired with slice selection along the *y*-axis. Notice how the axial slice view of column (a) appears more noisy compared to the axial view of the low-resolution image in column (b), due to the lower SNR of the high-resolution scan, as signal intensity scales linearly with the voxel size. As a consequence of partial volume effects, also a blurring of intensities at tissue boundaries occurs in the low-resolution images (b) and (c).

4.1.2 Resolution enhancement techniques

In an attempt to address the challenge of achieving high-resolution isotropic 3D MRI images, several approaches have been introduced that try to enhance spatial resolution. **Hardware improvements**, such as increasing the number of receiver coils or increasing the main magnetic field strength, directly increase the observed MR signal and intrinsic SNR (cf. section 2.5.3, Eqs. (2.29)-(2.31)). According to Eq. (2.31) in section 2.5.3, it follows that for a similar SNR value, scanners with a high B_0 value or a high number of coil receiver channels can reduce the Δx , Δy , and Δz voxel dimensions, thereby producing images with higher spatial resolution and contrast (Regatte & Schweitzer, 2007). As an example, ultrahigh field (UHF) MRI scanners for human MRI are being developed with $B_0 = 7$ T (Feinberg et al., 2023), $B_0 = 9.4$ T (Ivanov et al., 2023), or even $B_0 = 11.7$ T (Boulant et al., 2023) reporting resolutions in the μ m³ range, as opposed to common MRI resolutions in the mm³ range at $B_0 = 3$ T. However, such scanners are *expensive* and their strong magnetic fields pose an *increased risk of physiological effects*, including transient sensory side effects such as nausea, dizziness, metallic taste, and light flashes (Ladd et al., 2018), impeding their clinical use.

In addition, various image **post-processing techniques** have been introduced that try to enhance spatial resolution after the image acquisition (Kang & Chaudhuri, 2003; Park et al., 2003; Farsiu et al., 2004; Greenspan, 2009; Van Reeth et al., 2012). In most MRI machines, standard interpolation (usually zero-padding) of the k-space data is available to reduce the voxel size of the images. Applying this interpolation facilitates the visualization but several artifacts such as blur and contrast loss are added, while no new information is introduced into the image (Van Reeth et al., 2012). Furthermore, there is a growing focus on resolution enhancement through artificial intelligence (AI)-based methods, including the development of deep learning techniques to enhance MRI image resolution (Ding et al., 2020; Roy et al., 2023; Chen et al., 2023). The effectiveness of such learning-based methods hinges on the quality of the training datasets and the machine learning models. Often these algorithms require large datasets for training. The diversity and size of these datasets are crucial in developing robust AI models that can generalize well to new, unseen images. For learning-based methods targeting resolution enhancement it is crucial that high-resolution isotropic data is available. However, obtaining high-resolution data for training is often challenging, as most existing acquisition protocols prioritize data acquisition with non-isotropic resolution to minimize total scan time.

Efforts have also been made to employ techniques that reduce scan time, thus enabling higher resolutions within the same time frame. **Accelerated acquisition techniques** include parallel MRI (Griswold et al., 2002; Heidemann et al., 2003; Pruessmann, 2006), PROPELLER (Pipe, 1999), compressed sensing (Lustig et al., 2007), and simultaneous multi-slice imaging (Setsompop et al., 2012; Feinberg & Setsompop, 2013; Feinberg et al., 2013). However, the drawback of *SNR loss persists*, partly due to scanning with smaller voxel sizes and partly due to the more complex reconstruction algorithms associated with these techniques, which inherently involve potential SNR loss. For instance, it is known that parallel MRI results in a loss of SNR in the reconstructed images by a factor equal to the square root of the parallel acceleration factor, owing to reduced signal averaging (Robson et al., 2008).

4.1.3 Super-resolution reconstruction

Another very promising approach to enhance spatial resolution, is to use signal processing techniques to obtain a high-resolution (HR) image from multiple observed low-resolution (LR) images. One particular method that has been actively studied, and which was first introduced around the early 1980s as an idea to improve the resolution of image sequences in video applications (Tsai & Huang, 1980; Kim et al., 1990), is called **super-resolution reconstruction (SRR)** (Van Reeth et al., 2012):

The general idea of (multi-frame) super-resolution¹ reconstruction states that multiple low-resolution images of the same object, acquired with slightly different imaging conditions, can be combined to reconstruct a high-resolution image that contains additional frequency content. Each of the low-resolution images transforms and samples the high-resolution scene in a distinct fashion, such that aliased frequency content between the images can be retrieved. As such, the spatial resolution can be enhanced beyond the inherent capabilities of the imaging system (Park et al., 2003; Kang & Chaudhuri, 2003; Tian & Ma, 2011).

Over time, SRR has developed in a research field of its own, with applications in many real-world problems in different fields, from satellite and aerial imaging to medical imaging, remote sensing, image or video forensics, and many other fields. A comprehensive survey of super-resolution (SR) methods, including an algorithm taxonomy classification, has been given by Nasrollahi & Moeslund (2014), and some noteworthy review papers have been published over the years (Park et al., 2003; Van Reeth et al., 2012; Plenge et al., 2012; Yue et al., 2016)

Generally, SRR methods can be categorized based on the number of LR images involved, the imaging domain employed and the corresponding reconstruction method. In terms of the number of the LR images involved, **single-frame** (or single-image) SR and **multi-frame** SR methods, can be distinguished. In single-frame SR, an HR image is reconstructed from a single LR image. Mostly some learning algorithms are employed that try to hallucinate the missing information of the super-resolved images using the relationship between the LR image and the HR image from a training database (Dong et al., 2015). However, also other single-frame SR approaches exist, such as interpolation-based methods (Zhang & Wu, 2006) or reconstruction-based methods (Candocia & Principe, 1999). While single-frame SR methods present the benefit of a reduced computational complexity and smaller data storage needs, the ill-conditioned² nature of the SR problem makes the recovery of the HR image from a single LR image challenging.

To overcome the drawbacks of single-image SR, multi-frame SR methods have been proposed. Such methods estimate an HR image by exploiting complementary information from multiple

¹Note that the geometric SRR methods discussed in this thesis, should not be confused with *super-resolution restoration* or *super-resolution fluorescence microscopy*. The former referring to the use of algorithms that operate on a single image in an attempt to recover information beyond the diffraction cut-off frequency by extrapolation (without changing the amount of pixels/voxels of the original image) (Andrews & Hunt, 1979), while the latter refers to the range of techniques in optical microscopy that enable imaging beyond resolutions imposed by the diffraction limit, which is due to the diffraction of light, e.g., the Nobel Prize winning research of E. Betzig, W.E. Moerner and S. Hell (Betzig et al., 2006; Moerner & Kador, 1989; Hell & Wichmann, 1994).

²In layman's terms, an *ill-conditioned* problem is one where, for a small change in the inputs (the independent variables) there is a large change in the answer (or dependent variable). This implies that pinpointing the correct solution or answer to the equation is challenging.

LR images (Farsiu et al., 2004; Shilling et al., 2009; Poot et al., 2010). Usually these multi-frame SR methods assume that there is a targeted HR image and the observed LR images have some relative geometric and/or photometric displacements from the targeted HR image. These algorithms then exploit the differences between the LR observations to reconstruct the targeted HR image, and hence are referred to as **reconstruction-based SR algorithms** (Nasrollahi & Moeslund, 2014). Reconstruction-based SR algorithms treat the SR problem as an **inverse problem** and therefore, like any other inverse problem, need to construct a **forward model**. In the case of MRI, this forward model should describe the imaging process or acquisition of an LR image, as will be elaborated in section 4.2.3.2 hereafter. In contrast to single-frame SR, where the inverse problem is ill-posed³ because the HR image obtained from the LR image is non-unique or unstable, multi-frame SR techniques generate superior results (i.e. the recovery of true high frequency content) provided that the **inter-frame motion** between the LR images is estimated or known with high accuracy.

Typically, two basic types of multi-frame methods can be distinguished: **frequency domain** (Kim et al., 1990; Kim & Su, 1993) and **spatial domain** (Van Reeth et al., 2012) methods. Although frequency domain methods can be very efficient, they are not able to incorporate prior knowledge about the spatial domain in their formulation. Consequently, many spatial domain methods have been developed (Nasrollahi & Moeslund, 2014), including approaches based on non-uniform interpolation (Nguyen et al., 2001), iterative back projection (IBP) (Irani & Peleg, 1993; Greenspan et al., 2002), projection onto convex sets (POCS) (Shilling et al., 2006, 2009), maximum likelihood (ML) (Elad & Feuer, 1996, 1997; Beirinckx et al., 2019, 2020), and maximum a posteriori (MAP) estimation (Elad & Feuer, 1996, 1997; Beirinckx et al., 2022, 2024). The last two methods are of particular interest in this work.

In the contribution chapters of this thesis, super-resolution reconstruction as a **multi-frame spatial domain approach** is studied and applied to medical MRI imaging. The next section will delve into some specific aspects and considerations essential for the application of SRR to MRI.

4.2 Super-resolution reconstruction applied to MRI

The first example of SRR applied to MRI was described in a 2001 patent (filed in 1997) (Fiat, 1997). While SRR has been studied in MRI for different applications, most work has concentrated on brain MRI (Peled & Yeshurun, 2001; Greenspan et al., 2002; Peeters et al., 2004; Zhang et al., 2008; Rousseau et al., 2006, 2010; Gholipour et al., 2010). The reason for this being the fact that SRR is highly dependent on accurate registration of the different LR images (Robinson & Milanfar, 2006; Lin & Shum, 2004), and because relatively simple global motion models can be applied to brain MRI, as opposed to MRI of objects that exhibit more complex often non-rigid motion, e.g. muscle contraction and relaxation in some organs.

Problems that are not well-posed in the sense above are termed *ill-posed*.

³In mathematics, a *well-posed* problem is one for which the following properties hold (Hadamard, 1902):

⁻ The problem has a solution.

⁻ The solution is unique.

⁻ The solution's behavior changes continuously with the initial conditions.

4.2.1 Concept and definition

In the context of MRI, the goal of super-resolution reconstruction can be defined as:

The goal of super-resolution reconstruction in MRI is the estimation of a highresolution MR image (or high-resolution parameter map(s)) with isotropic resolution from a set of 2D multi-slice MR images, with a low through-plane resolution and with varying slice-encoding directions.

Note that for the application to MRI, SRR demands the use of a **2D multi-slice** pulse sequence to provide aliasing via the slice-selection profile along the slice-encoding (i.e., through-plane) direction. More specifically, **the rationale of SRR is that spatial aliasing occurring along the slice-encoding direction of a 2D multi-slice image can be exploited to achieve an increase in resolution along that same dimension (Greenspan et al., 2002; Van Reeth et al., 2012). Here, aliasing refers to high spatial frequency information that is being disguised as low frequency information in the 2D multi-slice readout sampling process. A more detailed description of the theory behind sampling and aliasing in the slice-encoding direction can be found in the works of NoII et al. (1997) and Pipe (1998). Retrieving aliased content is the major advantage of SRR over standard interpolation techniques.**

4.2.2 Historical misconceptions

Historically, efforts were initially made to employ SRR in both the in-plane (frequency-encoded) and through-plane (slice-encoding) dimensions of 2D multi-slice MRI acquisitions (Peled & Yeshurun, 2001; Carmi et al., 2006; Tieng et al., 2011). However, subsequent findings revealed that while SRR techniques could enhance through-plane resolution, **achieving resolution enhancement in the in-plane dimensions within 2D multi-slice or 3D readouts was not feasible** (Scheffler, 2002; Peled & Yeshurun, 2002; Greenspan et al., 2002; Poot et al., 2010; Uecker et al., 2011; Plenge et al., 2012). This limitation arises from the **Fourier encoding scheme's band-limited nature in the frequency and phase encoding directions, inherently excluding aliasing**. Therefore, when defining SRR as the recovery of high-frequency components corrupted by aliasing (Kang & Chaudhuri, 2003), true resolution enhancement using SRR is not possible in-plane in 2D readout, nor in 3D readout. The only viable method for enhancing resolution in the in-plane directions involves the acquisition of data beyond the *k*-space span (Luong, 2009), i.e. the *k*-space coverage as defined by $[-k_{max}, k_{max}]$, with k_{max} the highest measured frequency.

Whereas the individual slices in a 2D slice stack are Fourier encoded without aliasing, **in the slice-encoding direction there are no inherent limitations on the frequency spectrum, and aliased frequencies may potentially be recovered**. The amount of aliasing depends on the slice profile which ideally is a rect function for non-overlapping slices (cf. Chapter 3, section 3.4.6.4). However, since the slice profile is determined by the Fourier transform of the finite length slice selection pulse, it will be only an approximation of the rect function. Because of the aliasing present when the object is convoluted with this slice profile, SRR is possible in the slice-encoding direction.

4.2.3 Key components

Generally, a number of essential aspects are key when studying SRR. Namely, the choice of SR **acquisition strategy** to obtain the images with low through-plane resolution (section 4.2.3.1), providing an accurate **imaging model** of the acquisition process (section 4.2.3.2), and the implementation of an appropriate **model-based reconstruction** technique to estimate the HR image (or HR tissue parameter maps) from the acquired LR images (section 4.2.3.3).

4.2.3.1 Acquisition strategies

Since the spatial resolution can only be improved in the through-plane direction, the in-plane resolution (i.e. in-plane voxel size) is usually chosen higher than the through-plane resolution (i.e. through-plane voxel size) when acquiring a set of images for SRR. Consequently, the spatial resolution of the acquired LR images for SRR is anisotropic (see Fig. 4.1). It is customary to define the latter property in terms of an associated **anisotropy factor (AF)**, which can be defined as *the ratio of the through-plane resolution and the in-plane resolution* (see Fig. 4.4).

Over the years, a plethora of SRR acquisition strategies for MRI have been developed, some strategies more effective than others. In what follows, the main multi-frame super-resolution acquisition strategies for MRI are highlighted.

Sub-voxel shift in the through-plan direction In this approach, the LR acquisition matrix is shifted by sub-pixel distances along the slice-encoding direction for subsequent images (Greenspan et al., 2002; Ben-Ezra et al., 2009), see also Fig. 4.2. In this sampling scheme, all LR images sample the same part of *k*-space, causing the SRR to rely exclusively on recovering the aliased frequencies in the slice-encoding direction. A drawback of this method is that the highest frequency regions of *k*-space are not sampled in all dimensions (Plenge et al., 2012).



Figure 4.2: Schematic representation of the sampling strategy using sub-pixel shifts along the slice encoding direction. To reach isotropic resolution, a minimum of N low-resolution images is required, whereby N is equal to the ratio of the through-plane and in-plane resolution.

Orthogonal slice orientations A second approach consists of acquiring a set of multi-slice images for which the slice-encoding direction is chosen along the three orthogonal directions (Souza & Senn, 2008; Gholipour et al., 2010), see also Fig. 4.3. In essence, this approach improves upon the sub-voxel shift method in terms of coverage of high-frequencies in 3D *k*-space. However, for large anisotropy factors, some high-frequency regions are still not uniformly sampled.



Figure 4.3: Schematic representation of the sampling strategy using three orthogonal multi-slice image scans for SRR. The coloured box represents the isotropic resolution of the to be reconstructed HR image.

Slice orientations around a common frequency-encoding axis To guarantee a more uniform sampling of higher *k*-space frequencies, LR images can be acquired with rotational increments of the slice-encoding direction around a common phase- or frequency-encoding direction (Shilling et al., 2009; Poot et al., 2010; Plenge et al., 2012). This is illustrated in Fig. 4.4. As highlighted in Chapter 2, rotation in image space results in a rotation in frequency domain. As such, acquiring the LR images with different slice orientations ensures that each LR image covers a different part of *k*-space. When the rotation is only performed about one fixed axis (frequency or phase encoding direction), the *k*-space can only be sampled in a cylinder with radius $\frac{1}{a}$, with *a* the voxel size in phase or frequency encoding direction (see Fig. 4.4). To limit scan time, the minimal number of slice orientations that maximally covers the *k*-space by rotating about the center, is chosen. Preferably, the cylinder is completely sampled while the overlap between the different *k*-spaces is as small as possible. Hence the number of different slice orientations, *N*, needed to fill the *k*-space of the HR imaged object with a minimal overlap is given by (Plenge et al., 2012):

$$N = \lceil \frac{\pi}{2} \times \mathsf{AF} \rceil,\tag{4.1}$$

where $\lceil x \rceil$ denotes the ceiling function that maps x to the smallest integer greater than or equal to x. The N images are then acquired rotated about the chosen axis with rotational increments of $180^{\circ}/N$. For example, for AF = 4, this means that $N = \lceil 6.2831 \dots \rceil = 7$ different slice orientations are required.



Figure 4.4: Schematic comparison between image space and k-space (2D and 3D view) for a multi-orientation low resolution acquisition. The anisotropy of the voxels in image space is defined by the anisotropy factor, $AF = \frac{b}{a}$ with b, the slice thickness, and a, the voxel size in the frequency encoding direction (and phase encoding direction). Since we choose to rotate only around the phase encoding axis, the k-space can only be sampled in a cylinder with diameter $\frac{1}{a}$.

Which of these acquisition strategies is most optimal?

Recent work by Nicastro et al. (2022), which evaluated the estimation performance of shift-based and rotated-based multi-slice SR acquisition strategies in a Bayesian framework, has demonstrated the superiority of the rotated acquisition scheme in terms of accuracy, precision, and Bayesian mean squared error. Moreover, the rotation-based acquisition scheme proved to be more resilient to motion. Therefore, in the remainder of this thesis, the acquisition strategy with rotation of the slice-encoding direction around a common frequency encoding (or phase encoding) axis will consistently be used. For EPI imaging,

which is susceptible to geometric distortions along the phase-encoding direction due to its long readout and low-frequency bandwidth in that dimension, rotation around a common phase encoding axis is preferred. This approach ensures that geometric distortions and associated acquisition artifacts occur consistently across all images, minimizing additional blurring in the super-resolution reconstructed (SRR) result. It also reduces the need for corrective steps to address varying distortion directions.

4.2.3.2 Imaging model

Since SRR is an inverse problem, it requires a proper **imaging model** or observation model that relates the underlying HR image (or HR parameter maps) to the observed LR images. This imaging model consists of two main parts: a forward model of the physical acquisition process, and a statistical model that appropriately describes the noise in the observed magnitude LR data.

Forward model of physical acquisition

Let $s = \{s_n\}_{n=1}^N$ be the set of the *N* vectorized noiseless LR 2D multi-slice contrastweighted magnetide images, where $s_n = \{s_n\}_{l=1}^{N_s} \in \mathbb{R}^{N_s \times 1}$ is sampled at the LR grid points $y_n = \{y_{nl}\}_{l=1}^{N_s} \in \mathbb{R}^{N_r \times 1}$ employed at the lR grid points $y_n = \{r_n\}_{l=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ represent the virtual, noise-free HR image assumed to be acquired with the same contrast-weighting settings as s_n and defined at the targeted HR grid points $x = \{x_j\}_{j=1}^{N_r} \in \mathbb{R}^{3 \times N_r}$, with N_r the number of isotropic voxels of the HR image. Finally, let $A_n \in \mathbb{R}^{N_{s_n} \times N_{r_n}}$ be a linear operator that defines the transformation of the high resolution image r_n to the low resolution image s_n . Then, in its most general form, each s_n can be modelled using the following **SRR forward model**:

$$\boldsymbol{s}_n = \boldsymbol{A}_n \boldsymbol{r}_n, \tag{4.2}$$

or for equal contrast-weighting settings per LR image s_n , i.e. when the underlying virtual, noise free HR image r is assumed identical and anatomical (i.e., scalar) SRR is considered:

$$\boldsymbol{s}_n = \boldsymbol{A}_n \boldsymbol{r},\tag{4.3}$$

or when the transformation operator A_n is further decomposed:

$$s_n = DBG_n M_{\theta_n} r, \qquad (4.4)$$

where $M_{\theta_n} \in \mathbb{R}^{N_r \times N_r}$, $G_n \in \mathbb{R}^{N_r \times N_r}$, $B \in \mathbb{R}^{N_r \times N_r}$, and $D \in \mathbb{R}^{N_s \times N_r}$ are linear operators that describe unwanted motion, a known geometric transformation, spatially invariant blurring, and down-sampling, respectively (see also Fig. 4.5). Each of these operators needs to be further specified in the context of MRI.



Figure 4.5: The SRR acquisition forward model.

- **Unwanted motion** A crucial requirement for multi-frame SRR to succeed is that inter-frame motion between the LR images is known with high accuracy. In addition to the known geometric transformation resulting from the chosen acquisition strategy (see below), small unforeseen spatial motion in the LR scans, e.g., due to patient motion, must also be integrated into the forward model. This unforeseen (or unwanted) motion is modelled by operator M_{θ_n} . The unknown and thus to be estimated motion parameters are represented by $\theta = \{\theta_n\}_{n=1}^N$. For the latter, various 3D rigid-body parametrizations are possible, e.g., using Euler angles (3 translations, 3 rotation angles) or quaternions. More details about motion compensation are provided in section 4.3 of this chapter and in the contribution chapters where motion correction is combined with SRR.
- Geometric transformation The operator G_n models a known affine⁴ transformation representing the geometric deformation or mapping of the points in the HR space to the points in the LR space of image s_n . Geometric deformation, often called **image** warping, is crucial in SRR, since it provides the different views of the same object, bringing in additional information. Luckily, for MRI, the affine spatial transformations that serve as input for G_n are known in advance from the chosen acquisition strategy (cf. section 4.2.3.1). Particularly, the orientation and position of the field-of-view of the individual LR images is registered during acquisition in the 'world' (or scanner) coordinate system. After reconstruction of the k-space data to image space, the voxel intensities of the LR image are saved in a standard MRI file format (e.g., DICOM or NIfTI) and the associated low-resolution-voxel-to-world spatial coordinate transformation $(T_{n,lrv2w})$ is stored as a 4 × 4 affine transformation matrix together with the metadata information of the scan (patient identification, study time, acquisition info, etc.). Furthermore, as indicated in Eqs. (4.2)-(4.4), the MRI images for SRR are parameterized as vectors of (anisotropic) voxels, and operators acting on these image vectors operate through matrix-vector multiplication in so-called voxel space. For

$$y = Hx + t, \tag{4.5}$$

$$\begin{bmatrix} \mathbf{y} \\ 1 \end{bmatrix} = \mathbf{T} \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix}, \quad \text{with } \mathbf{T} = \begin{bmatrix} \mathbf{H} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix}.$$
(4.6)

⁴An affine coordinate transformation between two coordinate systems $x \in \mathbb{R}^{n \times 1}$ and $y \in \mathbb{R}^{n \times 1}$ is a geometric transformation that preserves collinearity and parallelism (i.e., sets of parallel lines remain parallel after transformation). It is composed of a linear transformation (rotation, scaling/reflection or shears) and a translation:

where $H \in \mathbb{R}^{n \times n}$ specifies the linear transformation matrix in *n* dimensions, and $t \in \mathbb{R}^{n \times 1}$ specifies the translation part of the transformation. Alternatively, the augmented affine transformation can be specified using a single matrix multiplication following an augmented matrix-vector representation:

operator G_n , the associated transformation should therefore be defined as a voxel-tovoxel affine transformation. In brief, to accomplish this, it typically suffices to choose a virtual HR template reconstruction grid⁵ with isotropic resolution and define its associated *high-resolution-voxel-to-world* transformation (T_{hrv2w}). Next, the required *high-resolution-voxel-to-low-resolution-voxel* affine transformation can be created from concatenating both voxel-to-world affine transformations, i.e. (operators acting from right to left): $T_{n,hrv2lrv} = T_{n,lrv2w}^{-1} T_{hrv2w}$. Finally, since operator $G_n \in \mathbb{R}^{N_r \times N_r}$ inherently operates in HR space, an additional affine transformation is necessary to create the corresponding *high-resolution-voxel-to-high-resolution-voxel* transformation. For this, both the scaling of the voxel sizes and a small additional translation resulting from the voxel center shift between a low and high-resolution voxel needs to be compensated for: $T_{n,hrv2hrv} = T_{n,scale-centershift} T_{n,hrv2lrv}$, where $T_{n,scale-centershift}$ is further defined as

$$\boldsymbol{T}_{n,\text{scale-centershift}} = \begin{bmatrix} v_{n,1}/w_1 & 0 & 0 & (v_{n,1}/w_1 - 1)/2 \\ 0 & v_{n,2}/w_2 & 0 & (v_{n,2}/w_2 - 1)/2 \\ 0 & 0 & v_{n,3}/w_3 & (v_{n,3}/w_3 - 1)/2 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
(4.7)

with $v_n = \{v_{n,j}\} \in \mathbb{R}^{3 \times 1}$ the 3D voxel size of LR image s_n (in mm), and $w = \{w_j\} \in \mathbb{R}^{3 \times 1}$ the 3D voxel size of the virtual HR image r (in mm).

- **Blur and PSF modeling** The blurring operator B represents the amount of blur added during the overall acquisition process, and is therefore often assimilated to the blur introduced by the imaging system. For 2D multi-slice acquisitions that sample a rectangular part of k-space, the point spread function (PSF) can be assumed separable and therefore modeled as the product of three separate one-dimensional (1D) PSFs that are applied in the three orthogonal directions aligned with the MR image coordinate axes (Poot et al., 2010). The PSFs in the frequency and phase encoding direction are defined by the rectangular part of the k-space that is regularly sampled, and can thus be modeled by Dirichlet or periodic sinc functions. The PSF in the slice encoding direction depends on the slice selection profile. Such slice selection is often performed by applying either a (windowed) sinc or a Gaussian shaped RF pulse. Therefore, the sampling in the through-plane direction can be modeled by a (smoothed) box or a Gaussian function, respectively (Poot et al., 2010).
- **Down-sampling operator** This operator D is required to resample the HR grid to LR grid with the aim of generating the aliased LR images from the warped and blurred HR image. The down-sampling only proceeds in the slice-encoding direction, not in the in-plane directions.

Statistical noise model

Next, the acquired LR images are subject to noise:

$$\tilde{\boldsymbol{s}}_n = \boldsymbol{s}_n + \boldsymbol{e}_n, \tag{4.8}$$

⁵An easy approach to create the virtual isotropic HR reconstruction grid is to resample one of the input LR images (e.g., the LR image with 0° slice orientation), using MRtrix3 (Tournier et al., 2019). Particularly, using the following command: $\mbox{mrgrid lrimg_0.nii regrid -scale 1,1,AF hrgrid.nii}$, which assumes a prior image conversion to NIfTI format and resampling of the (3rd) slice-encoding dimension with anisotropy factor AF. Conveniently, this command also updates the affine transformation in the metadata accordingly, creating the required *high-resolution-voxel-to-world* (T_{hrv2w}) transformation.

with $e_n \in \mathbb{R}^{N_s \times 1}$ a vector representing the noise. Depending on the acquisition method, different noise models need to be applied. When magnitude images are reconstructed from single-coil k-space data, the noisy voxel intensities can be modeled as **Rician**⁶ distributed random variables (Gudbjartsson & Patz, 1995; den Dekker & Sijbers, 2014). For a multi-coil acquisition, the magnitude data are governed by a non-central chi distribution (Constantinides et al., 1997; den Dekker & Sijbers, 2014), or they are again Rician distributed for multi-coil data acquired with SENSE (Aja-Fernández et al., 2014), or with GRAPPA jointly with a spatial-matched-filter or the Adaptive Combine method (Walsh et al., 2000). Generally, when the SNR is high enough (> 3), which is a valid assumption when low-resolution voxels with larger voxel size are acquired, the aforementioned distributions can be well approximated by a Gaussian distribution (Gudbjartsson & Patz, 1995; Andersen & Kirsch, 1996; Constantinides et al., 1997). In contribution chapter 6, a Rician noise model will be assumed, while in chapters 5 and 7, a Gaussian noise model will be assumed. The reader is referred to these respective chapters for a more in dept description of both data distributions. Generally, the use of a Rician noise model is more complicated, since it involves the evaluation of modified Bessel functions of the first kind (Gudbjartsson & Patz, 1995), for which typically some polynomial approximations have to be used (Abramowitz, 1974; Press et al., 1992).

Finally, the full imaging model and the sampling of all N multi-slice low-resolution MR images can be combined into a single matrix multiplication,

$$\tilde{s} = Ar + e, \tag{4.10}$$

with

$$\tilde{\boldsymbol{s}} = \begin{bmatrix} \tilde{\boldsymbol{s}}_1 \\ \tilde{\boldsymbol{s}}_2 \\ \vdots \\ \tilde{\boldsymbol{s}}_N \end{bmatrix}, \quad \boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_1 \\ \boldsymbol{A}_2 \\ \vdots \\ \boldsymbol{A}_N \end{bmatrix}, \quad \boldsymbol{e} = \begin{bmatrix} \boldsymbol{e}_1 \\ \boldsymbol{e}_2 \\ \vdots \\ \boldsymbol{e}_N \end{bmatrix}.$$
(4.11)

4.2.3.3 Reconstruction

Once the anisotropic LR images are acquired and the imaging model is specified, the actual reconstruction part of SRR can be performed. The goal is to recover an HR image (or HR parameter maps) with an isotropic resolution from the set of noisy LR images.

Anatomical SRR

As highlighted in section 3.6 of the previous chapter, different choices of estimators exist. When the distribution of the data is to be taken into account, the maximum likelihood estimator (MLE) is the go-to estimator. Assuming zero mean, Gaussian distributed noise, which is a reasonable assumption when the SNR of the magnitude MRI data is sufficiently

$$M = |(x + N_1(0, \sigma)) + i(y + N_2(0, \sigma))| = \sqrt{(x + N_1(0, \sigma))^2 + (y + N_2(0, \sigma))^2}$$
(4.9)

⁶Note that under the assumption of Rician distributed magnitude data the noise is not additive. Instead, magnitude data M is considered as the result of a nonlinear transformation of two independent Gaussian distributed variables corresponding with the real and imaginary components of the raw MR images, as recorded in k-space and both corrupted with zero mean Gaussian distributed noise (here denoted as $N_1(0, \sigma)$ and $N_2(0, \sigma)$, with σ the standard deviation of the noise):

high (> 3) (Gudbjartsson & Patz, 1995; Andersen & Kirsch, 1996), the MLE of r is given by (Elad & Feuer, 1997):

$$\hat{\boldsymbol{r}}_{\mathsf{ML}} = \arg\max_{\boldsymbol{r}} \mathcal{L}\left(\boldsymbol{r}|\tilde{\boldsymbol{s}}\right) = \arg\min_{\boldsymbol{r}} \left(\tilde{\boldsymbol{s}} - \boldsymbol{A}\boldsymbol{r}\right)^{\mathsf{T}} \boldsymbol{R}_{ee}^{-1} \left(\tilde{\boldsymbol{s}} - \boldsymbol{A}\boldsymbol{r}\right), \qquad (4.12)$$

with $\mathcal{L}(\mathbf{r}|\tilde{\mathbf{s}})$ the likelihood function of \mathbf{r} given the observed input data $\tilde{\mathbf{s}}$, and $\mathbf{R}_{ee} \in \mathbb{R}^{NN_s \times NN_s}$ the autocorrelation matrix of the Gaussian noise. If all voxels in \mathbf{s} are statistically independent and the noise variance σ^2 is assumed to be spatially invariant, $\mathbf{R}_{ee} = \sigma^2 \mathbf{I}$ with \mathbf{I} the identity matrix. Correspondingly, Eq. (4.12) simplifies to the unweighted least-squares solution:

$$\hat{\boldsymbol{r}}_{\mathsf{ML}} = \arg\min_{\boldsymbol{r}} \left(\tilde{\boldsymbol{s}} - \boldsymbol{A} \boldsymbol{r} \right)^{\mathsf{T}} \left(\tilde{\boldsymbol{s}} - \boldsymbol{A} \boldsymbol{r} \right).$$
 (4.13)

As this solution satisfies the normal equation, the MLE can be written as a closed-form expression:

$$\hat{\boldsymbol{r}}_{\mathsf{ML}} = \left(\boldsymbol{A}^{\mathsf{T}}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^{\mathsf{T}}\tilde{\boldsymbol{s}}.$$
(4.14)

Unfortunately, for realistic MR image dimensions, the matrices in Eq. (4.14) will be too large to store explicitly, even as sparse matrices. Therefore, Eq. (4.13) is typically solved by using **iterative optimization algorithms**. In time, different iterative optimization algorithms for SRR have been studied (Plenge et al., 2012), yet no particular method outperformed the others. For anatomical, linear SRR, the conjugate gradient method⁷ is a widely used option that generally reaches convergence quickly (Poot et al., 2010), but also gradient descent methods using Barzilai-Borwein step size selection are known to exhibit fast convergence rates (Barzilai & Borwein, 1988). As will also be highlighted in the contribution chapters, to reduce memory and to increase computational efficiency, the matrix vector multiplications that are required in these methods are typically evaluated by a function, without storing the operator matrices explicitly.

Regularized SRR

While the use of iterative reconstruction techniques along with efficient strategies for implementing forward operator $\mathbf{A} \in \mathbb{R}^{NN_s \times N_r}$ make SRR feasible from a computational point-of-view, Eq. (4.13) remains a badly conditioned or even under-determined problem due to the high resolution at which the image is set to be reconstructed. Certain high spatial frequencies in the HR grid will not be present in any of the acquired LR images, which is the reason for Eq. (4.13) being potentially ill-conditioned. In order to remedy this, a regularization term can be added to Eq. (4.13), which reduces the variance of the solution (but also introduces a small bias to the solution):

$$\hat{\boldsymbol{r}} = \arg\min_{\boldsymbol{r}} \left[\left(\tilde{\boldsymbol{s}} - \boldsymbol{A} \boldsymbol{r} \right)^{\mathsf{T}} \left(\tilde{\boldsymbol{s}} - \boldsymbol{A} \boldsymbol{r} \right) + \boldsymbol{r}^{\mathsf{T}} \boldsymbol{K}^{\mathsf{T}} \boldsymbol{K} \boldsymbol{r} \right],$$
 (4.15)

with $K \in \mathbb{R}^{N_r \times N_r}$ specifying the regularization term. As some high frequencies will not be present in the acquired LR images, a suitable regularization is one that constrains these high frequencies. In some previous SRR works (Poot et al., 2013; Van Steenkiste et al., 2016, 2017), the squared Laplacian operator was used as a regularization term:

$$\hat{\boldsymbol{r}} = \arg\min\left[\|\tilde{\boldsymbol{s}} - \boldsymbol{A}\boldsymbol{r}\|_2^2 + \lambda \|\Delta \boldsymbol{r}\|_2^2\right], \qquad (4.16)$$

 $^{^7 \}rm An$ efficient MATLAB implementation of the conjugate gradient least squares (CGLS) method is that of M. Saunders *et al.*, adapted from (Paige & Saunders, 1982), which can be found at: https://web.stanford.edu/group/SOL/software/cgls/

with Δ the Laplace operator, and λ a hyperparameter that determines the weight of the regularization. The net effect of this type of regularization is a spatial smoothing of the reconstructed image.

The addition of a regularization term can also be understood within a Bayesian framework, as will be highlighted in contribution chapters 6 and 7. In Bayesian statistics, regularization corresponds to incorporating prior knowledge about the solution. The regularization term $\mathbf{r}^{\mathsf{T}}\mathbf{K}^{\mathsf{T}}\mathbf{K}\mathbf{r}$ of Eq. (4.15) then acts as a prior, penalizing unlikely solutions and thus guiding the reconstruction towards more plausible outcomes. This approach reduces the solution's variance by imposing prior beliefs about the image's characteristics, such as smoothness or sparsity, while introducing a small bias consistent with these beliefs. By solving the problem in this Bayesian context, the regularization not only addresses the ill-conditioned nature of Eq. (4.13) but also integrates prior knowledge to produce more reliable and robust HR reconstructions.

Apart from a squared Laplacian prior in Eq. (4.16), which was implemented for contribution chapters 5 and 7 of this thesis, other **regularization types** have been proposed, e.g. Total Variation (TV) based regularization (Rudin et al., 1992; Chambolle et al., 2011), which assumes that MR images consist of areas which are piecewise constant. Such TV regularization was used in contribution chapter 6. The main benefit of TV models is that they are very well suited to remove random noise or incoherent noise-like artifacts from random sub-sampling, while preserving the edges in the image (Knoll et al., 2011). In some practical MRI situations, however, the assumption of piecewise constancy is not necessary valid due to the inhomogeneities of the exciting B_1 field and the receive coils. Moreover, the use of TV may lead to staircasing artifacts with unnatural appearance (Knoll et al., 2011). In an attempt to eliminate the aforementioned restrictions, higher order total generalized variation (TGV) has been proposed by Bredies et al. (2010), which has been applied in MRI reconstruction problems to improve the image quality over conventional TV based regularization (Knoll et al., 2011). Although this type of regularization was not explicitly used in this work, it is worth noting to the reader.

Other regularization types exist for ill-conditioned image restoration problems, such as L1-norm wavelet transform regularization. The underlying idea of wavelet regularization is that natural images such as MRI images tend to be sparse in the wavelet domain. By minimizing the L1-norm of the wavelet transform, the regularization enforces sparsity in the reconstructed images, favoring results with only a few significant wavelet coefficients. The connection between L1 regularization and sparsity has garnered attention in MRI due to its use in Compressed Sensing (Donoho, 2006; Lustig et al., 2007; Gamper et al., 2008). Additionally, L1 wavelet regularization has proven effective in other rapid MRI reconstruction problems (Liu et al., 2008; Guerquin-Kern et al., 2009, 2011).

In addition to the type of regularization model, a reliable method is required to determine the optimal regularization hyperparameter λ , which plays an important role in balancing the so-called data-fidelity term and regularization term in Eq. (4.16). When the parameter λ is too small, the reconstructed HR image will fit the observed LR images properly but retain noise in homogeneous regions. When the parameter λ is too large, the reconstructed HR image will be over-smoothed, not only suppressing the noise but also eliminating details in the HR image. By adjusting the parameter λ , a compromise is achieved to suppress the noise and preserve the original HR image. Different methods exist for regularization hyperparameter selection, including the L-curve criterion (Hansen, 1992), Generalized Cross Validation (GCV) (Galatsanos & Katsaggelos, 1992), the unbiased predictive risk estimator (UPRE) (Vogel, 2002), or Morozov's discrepancy principle (Morozov, 1966). As an example, Fig. 4.6 demonstrates the hyperparameter selection using the discrepancy principle, which can be used when some extra information is available about the variance of the noise vector e. More specifically, Morozov's discrepancy principle suggests to choose the regularization parameter λ such that the norm of the residual corresponding to the regularized solution r_{λ} is approximately equal to the assumed bound for the noise level δ in the data:

$$\|\tilde{\boldsymbol{s}} - \boldsymbol{A}\hat{\boldsymbol{r}}_{\lambda}\|_{2} \approx \delta, \tag{4.17}$$

where the left side of the equation refers to the so-called 'discrepancy'. The discrepancy principle requires extra information about the noise, which poses a restriction on its usability. Under the (earlier) assumption of Gaussian distributed data with spatially invariant noise variance σ^2 , Eq. (4.16) can be extended to the following unconstrained regularized problem:

$$\hat{\boldsymbol{r}}_{\lambda} = \arg\min_{\boldsymbol{r}} \left[\frac{\|\tilde{\boldsymbol{s}} - \boldsymbol{A}\boldsymbol{r}\|_{2}^{2}}{\sigma^{2}} + \lambda \|\Delta \boldsymbol{r}\|_{2}^{2} \right].$$
(4.18)

It can be shown that the expected value of the data fidelity evaluated in the true underlying HR image follows a chi-squared distribution with M degrees of freedom (Mead, 2020): $\mathbb{E}\left[1/\sigma^2 \|\tilde{s} - Ar\|_2^2\right] \sim \chi_M^2$. Since the expected value of any χ_M^2 distributed value x is given by $\mathbb{E}[x] = M$, this allows to construct the following constrained optimization problem:

$$\hat{\boldsymbol{r}} = \arg\min_{\boldsymbol{r}} \|\Delta \boldsymbol{r}\|_2^2$$
, subject to $\frac{\|\tilde{\boldsymbol{s}} - \boldsymbol{A}\boldsymbol{r}\|_2^2}{\sigma^2} = M.$ (4.19)

Consequently, an appropriate selection of λ is obtained by solving:

$$\hat{\lambda} = \arg\min_{\lambda} \left\| \|\tilde{\boldsymbol{s}} - \boldsymbol{A}\boldsymbol{r}_{\lambda}\|_{2}^{2} - M\sigma^{2} \right|, \qquad (4.20)$$

where r_{λ} follows from Eq. (4.18) and where the degrees of freedom *M* correspond with the number of LR voxels in the SRR problem. A drawback of this example is that the reconstruction in Eq. (4.18) needs to be solved for different values of λ , which can be computationally demanding and inefficient for large scale reconstruction problems. Consequently, in many research publications, λ is heuristically set and scaled to the actual SNR of the acquired data.

Model-based SRR

Up to this point, the assumption was made that the only real difference between the LR images and the to be reconstructed HR image is spatial resolution and grid orientation. However, **SRR can also be applied for quantitative MRI and be combined with a certain parametric model for direct quantification of HR parameter maps from a set of contrast-weighted LR images**. The potential of SRR for HR isotropic quantitative parameter mapping from LR images has been shown for relaxometry (Van Steenkiste et al., 2017; Bano et al., 2020; Beirinckx et al., 2019, 2020, 2022), diffusion MRI (Van Steenkiste et al., 2016), and perfusion MRI (Bladt et al., 2017, 2020; Beirinckx et al., 2024). Conceptually, the application of SRR for qMRI can be seen as an extension of the anatomical acquisition forward model as described in Eq. (4.4). In particular, a **parametric signal model** is introduced describing the

CHAPTER 4



Figure 4.6: Regularization hyperparameter selection using **Morozov's discrepancy principle**. The green annotated reconstruction corresponds with the most optimal selection for λ , as predicted by Eq. (4.20). Note that for this anatomical SRR example 7 noisy LR images (dimensions $128 \times 128 \times 32$, AF = 4) were generated from a numerical brain phantom (Brainweb, (Cocosco et al., 1997)), with zero-mean, Gaussian distributed noise, where *M* corresponded to the number of LR voxels and $\sigma = \overline{s}_1/\text{SNR}$, with \overline{s}_1 the mean signal intensity of the first image and SNR = 4.5.

relationship between the to be estimated parameters and the HR contrast-weighted image. As highlighted in the previous chapter on qMRI, often such parametric signal models are **nonlinear**, so they cannot be represented by a linear matrix operator. Therefore, the MLE for the SR reconstructed HR parameter maps is given by:

$$\hat{\boldsymbol{\vartheta}} = \arg\min_{\boldsymbol{\vartheta}} \left(\tilde{\boldsymbol{s}} - \tilde{\boldsymbol{A}} \boldsymbol{f}(\boldsymbol{\vartheta}) \right)^T \left(\tilde{\boldsymbol{s}} - \tilde{\boldsymbol{A}} \boldsymbol{f}(\boldsymbol{\vartheta}) \right),$$
 (4.21)

with $\vartheta \in \mathbb{R}^{N_{\text{par}}N_{r} \times 1}$ the N_{par} HR parameter maps lexicographically ordered, N_{par} the number of different parameters in the signal model function, $f(\vartheta) \in \mathbb{R}^{NN_{r} \times 1}$ the N predicted HR images obtained from letting the signal model function operate on the parameters ϑ , which describe the physiological conditions for each of the N acquired LR images, and $\tilde{A} \in \mathbb{R}^{NN_{s} \times NN_{r}}$ the matrix projecting the model predicted HR images to the model predicted LR images.

Similar as for Eq. (4.15), regularization can be added to the equation of the model-based reconstruction problem:

$$\hat{\boldsymbol{\vartheta}} = \arg\min_{\boldsymbol{\vartheta}} \left(\tilde{\boldsymbol{s}} - \tilde{\boldsymbol{A}} \boldsymbol{f}(\boldsymbol{\vartheta}) \right)^{T} \left(\tilde{\boldsymbol{s}} - \tilde{\boldsymbol{A}} \boldsymbol{f}(\boldsymbol{\vartheta}) \right) + \boldsymbol{\vartheta}^{T} \boldsymbol{K}^{T} \boldsymbol{K} \boldsymbol{\vartheta}, \qquad (4.22)$$

with $K \in \mathbb{R}^{N_{\text{par}}N_r \times N_{\text{par}}N_r}$ the regularization matrix, which is applied to the lexicographically ordered parameter maps ϑ in model-based SRR. Now, in this case, a **separate regularization term for each parameter** is required to reduce the variance of the solution:

$$\hat{\boldsymbol{\vartheta}} = \arg\min_{\boldsymbol{\vartheta}} \|\tilde{\boldsymbol{s}} - \tilde{\boldsymbol{A}}\boldsymbol{f}(\boldsymbol{\vartheta})\|_{2}^{2} + \sum_{q=1}^{N_{\text{par}}} \lambda_{q} \|\Delta \boldsymbol{\vartheta}_{q}\|_{2}^{2}, \qquad (4.23)$$

where $\vartheta_q \in \mathbb{R}^{N_r \times 1}$ is an HR parameter map, lexicographically ordered, for one of the N_{par} parameters in the model function. The balanced selection of multiple hyperparameters λ_q poses an added complexity to model-based SRR compared to the selection of a single regularization hyperparameter λ in anatomical SRR. In addition, the often nonlinear coupling of the HR tissue parameter maps in the MRI signal models, e.g. in exponential T_1 or T_2 relaxation models (see preceding chapters), also creates an extra difficulty in the selection. Some multiple regularization parameter selection approaches have been proposed, often extensions of single parameter selection strategies (Belge et al., 2002; Gazzola & Reichel, 2016). However, at the time of writing, the application of these methods for quantitative model-based SRR has not yet been explored. Furthermore, with the advent of AI and learning-based techniques, there is also a noticeable shift towards reconstruction techniques that avoid the use of regularization and associated hyperparameter selection altogether, by learning prior information directly from the training data provided (Lønning et al., 2019; Sabidussi et al., 2021, 2023).

4.3 The need for robust motion compensation

A crucial condition for SRR is that no motion is present between the separately acquired LR images. These LR images should be accurately aligned to a small fraction of a voxel in a common reference frame. However, patient motion typically occurs when the scanned subject cannot remain still during imaging. This is notably the case for awake neonates, involuntary moving adult subjects, and fetuses. Moreover, the risk of motion increases when many LR images are required, such as for model-based SRR for qMRI, which demands the acquisition of multiple LR images with varying contrast weightings to sample the parametric signal model.

While some model-based SRR approaches have been proposed without any motion estimation (Bano et al., 2020), most model-based SRR methods compensate for motion artefacts by using a pre-processing routine in which the motion parameters of each LR image are updated

once after registration (Van Steenkiste et al., 2017):

$$\begin{cases} \hat{\theta}_{\text{reg}} = \arg\min_{\boldsymbol{\vartheta}} \mathcal{C}(\boldsymbol{\theta}, \tilde{\boldsymbol{s}}_{\text{moving}}, \tilde{\boldsymbol{s}}_{\text{ref,fixed}}) \\ \hat{\boldsymbol{\vartheta}} = \arg\min_{\boldsymbol{\vartheta}} \|\tilde{\boldsymbol{s}} - \tilde{\boldsymbol{A}}(\hat{\boldsymbol{\theta}}_{\text{reg}})\boldsymbol{f}(\boldsymbol{\vartheta})\|_{2}^{2} + \sum_{q=1}^{N_{\text{par}}} \lambda_{q} \|\Delta\boldsymbol{\vartheta}_{q}\|_{2}^{2} \end{cases}$$
(4.24)

where $C(\theta, \tilde{s}_{\text{moving}}, \tilde{s}_{\text{ref,fixed}})$ represents a registration cost function that aligns the (moving) LR images, $\tilde{s}_{\text{moving}}$, to a fixed reference image, $\tilde{s}_{\text{ref,fixed}}$, updating the motion parameter set $\theta \rightarrow \hat{\theta}_{\text{reg}}$. In a following step, $\hat{\theta}_{\text{reg}}$ is then fixed in the model-based SRR problem, as elaborated in Eq. (4.4). As such, the risk of artifacts occurring from incorrect alignment at the voxel level can be avoided.

However, a downside to the pre-registration approach is the lack of a feedback mechanism that connects the motion compensation routine with the final estimation of the biophysical parameters. Once the motion parameters $\hat{\theta}_{reg}$ are fixed, potential propagating registration errors in the SRR step can no longer be corrected for. Clearly, it is more beneficial to use integrated methods that combine model-based SRR with the joint estimation of motion parameters so as to reduce potential propagating errors and to allow the biophysical parameter maps of interest to be estimated with optimal accuracy and precision, i.e.,

$$\{\hat{\vartheta}, \hat{\theta}\} = \arg\min_{\vartheta, \theta} \|\tilde{s} - \tilde{A}(\theta)f(\vartheta)\|_{2}^{2} + \sum_{q=1}^{N_{\text{par}}} \lambda_{q} \|\Delta \vartheta_{q}\|_{2}^{2}.$$
(4.25)

References

- Abramowitz, M. (1974). Handbook of Mathematical Functions, With Formulas, Graphs, and Mathematical Tables. USA: Dover Publications, Inc.
- Aja-Fernández, S., Vegas-Sánchez-Ferrero, G., & Tristán-Vega, A. (2014). Noise estimation in parallel MRI: GRAPPA and SENSE. *Magn. Reson. Imaging*, 32(3), 281–290.
- Andersen, A. H., & Kirsch, J. E. (1996). Analysis of noise in phase contrast MR imaging. *Med. Phys.*, 23(6), 857–869.
- Andrews, H. C., & Hunt, B. R. (1979). Digital image restoration. Pattern Recognition, 11(1), 75.
- Bano, W., et al. (2020). Model-based super-resolution reconstruction of T2 maps. *Magn. Reson. Med.*, 83(3), 906–919.
- Barzilai, J., & Borwein, J. M. (1988). Two-point step size gradient methods. IMA Journal of Numerical Analysis, 8(1), 141–148.
- Beirinckx, Q., Bladt, P., van der Plas, M. C. E., van Osch, M. J. P., Jeurissen, B., den Dekker, A. J., & Sijbers, J. (2024). Model-based super-resolution reconstruction for pseudo-continuous arterial spin labeling. *NeuroImage*, 286, 120506.
- Beirinckx, Q., Jeurissen, B., Nicastro, M., Poot, D. H. J., Verhoye, M., den Dekker, A. J., & Sijbers, J. (2022). Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping. *Computerized Medical Imaging and Graphics*, 100, 102071.
- Beirinckx, Q., Jeurissen, B., Verhoye, M., den Dekker, A. J., & Sijbers, J. (2019). Super-resolution T1 mapping with integrated motion compensation in a joint maximum likelihood framework. *Magn. Reson. Mater. Phys.*, 32 (Suppl. 1), S14.05.
- Beirinckx, Q., Ramos-Llordén, G., Jeurissen, B., Poot, D. H. J., Parizel, P. M., Verhoye, M., Sijbers, J., & den Dekker, A. J. (2020). Joint maximum likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework: a simulation study. *Fundam. Inform.*, 172, 105–128.
- Belge, M., Kilmer, M. E., & Miller, E. L. (2002). Efficient determination of multiple regularization parameters in a generalized L-curve framework. *Inverse Problems*, 18(4), 1161.

- Ben-Ezra, A., Greenspan, H., & Rubner, Y. (2009). Regularized super-resolution of brain MRI. In 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Vols 1 and 2, (pp. 254+).
- Betzig, E., et al. (2006). Imaging intracellular fluorescent proteins at nanometer resolution. *Science*, 313(5793), 1642–1645.
- Bladt, P., Beirinckx, Q., van der Plas, M. C. E., Schmid, S., Teeuwisse, W. M., Jeurissen, B., den Dekker, A. J., Sijbers, J., & van Osch, M. J. P. (2020). Super-resolution reconstruction of single-PLD pseudo-continuous ASL images. In *Proc. Intl. Soc. Mag. Reson. Med.*, vol. 28, (p. 3293).
- Bladt, P., Beirinckx, Q., Van Steenkiste, G., Jeurissen, B., Achten, E., den Dekker, A. J., & Sijbers, J. (2017). Super-resolution multi-PLD PCASL: a simulation study. *Magn. Reson. Mater. Phy.*, 30 (Suppl. 1)(S396).
- Boulant, N., et al. (2023). Analysis of noise in phase contrast MR imaging. *Commissioning of the Iseult* CEA 11.7 T whole-body MRI: current status, gradient–magnet interaction tests and first imaging experience, 36(2), 175–189.
- Bredies, K., Kunisch, K., & Pock, T. (2010). Total generalized variation. SIAM Journal on Imaging Sciences, 3(3), 492–526.
- Candocia, F., & Principe, J. (1999). Super-resolution of images based on local correlations. *IEEE Transactions on Neural Networks*, *10*(2), 372–380.
- Carmi, E., Liu, S., Alon, N., Fiat, A., & Fiat, D. (2006). Resolution enhancement in mri. *Magnetic Resonance Imaging*, 24(2), 133–154.
- Chambolle, A., Levine, S. E., & Lucier, B. J. (2011). An upwind finite-difference method for total variation-based image smoothing. *SIAM J. Imaging Sci.*, 4(1), 277–299.
- Chen, Z., et al. (2023). Deep Learning for Image Enhancement and Correction in Magnetic Resonance Imaging State-of-the-Art and Challenges. *J. Digit. Imaging*, *36*(1), 204–230.
- Cocosco, C. A., Kollokian, V., Kwan, R. K. S., & Evans, A. C. (1997). BrainWeb: Online interface to a 3D MRI simulated brain database. *Neuroimage*, 5(4), 425.
- Constantinides, C. D., Atalar, E., & McVeigh, E. R. (1997). Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magn. Reson. Med.*, *38*(5), 852–857.
- den Dekker, A. J., & Sijbers, J. (2014). Data distributions in magnetic resonance images: A review. *Physica Med.*, *30*(7), 725–741.
- Ding, P. L. K., Li, Z., Zhou, Y., & Li, B. (2020). Deep Residual Dense U-Net for Resolution Enhancement in Accelerated MRI Acquisition.
- Dong, C., et al. (2015). Image super-resolution using deep convolutional networks.
- Donoho, D. (2006). Compressed sensing. IEEE Trans. Inf. Theory, 52(4), 1289-1306.
- Elad, M., & Feuer, A. (1996). Super-resolution reconstruction of an image. In *Proceedings of 19th Convention of Electrical and Electronics Engineers in Israel*, (pp. 391–394).
- Elad, M., & Feuer, A. (1997). Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE Trans. Image Process.*, 6(12), 1646–1658.
- Farsiu, S., et al. (2004). Fast and Robust Multiframe Super Resolution. *IEEE Trans. Image Process.*, 13(10), 1327–1344.
- Feinberg, D. A., Beckett, A., & Chen, L. (2013). Arterial spin labeling with simultaneous multi-slice echo planar imaging. *Magn. Reson. Med.*, 70(6), 1500–1506.
- Feinberg, D. A., & Setsompop, K. (2013). Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. *Journal of Magnetic Resonance*, 229, 90–100. Frontiers of In Vivo and Materials MRI Research.
- Feinberg, D. A., et al. (2023). Next-generation MRI scanner designed for ultra-high-resolution human brain imaging at 7 Tesla. *Nature Methods*, *20*(12), 2048–2057.
- Fiat, D. (1997). Method of enhancing an MRI signal. U.S. Patent 6294914 B1.
- Galatsanos, N., & Katsaggelos, A. (1992). Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation. *IEEE Trans. Image Process.*, 1(3), 322–336.
- Gamper, U., Boesiger, P., & Kozerke, S. (2008). Compressed sensing in dynamic MRI. *Magn. Reson. Med.*, 59(2), 365–373.

- Gazzola, S., & Reichel, L. (2016). A new framework for multi-parameter regularization. *BIT Numerical Mathematics*, *56*(3), 919–949.
- Gholipour, A., Estroff, J. A., & Warfield, S. K. (2010). Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans. Med. Imag.*, 29(10), 1739–1758.
- González Ballester, M. A., Zisserman, A. P., & Brady, M. (2002). Estimation of the partial volume effect in MRI. *Med. Image Anal.*, 6(4), 389–405.
- Greenspan, H. (2009). Super-resolution in medical imaging. Comput. J., 52, 43-63.
- Greenspan, H., Oz, G., Kiryati, N., & Peled, S. (2002). MRI inter-slice reconstruction using superresolution. Magn. Reson. Imaging, 20(5), 437–446.
- Griswold, M. A., et al. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn. Reson. Med.*, 47(6), 1202–1210.
- Gudbjartsson, H., & Patz, S. (1995). The Rician distribution of noisy MRI data. *Magn. Reson. Med.*, 34(6), 910–914.
- Guerquin-Kern, M., Haberlin, M., Pruessmann, K. P., & Unser, M. (2011). A Fast Wavelet-Based Reconstruction Method for Magnetic Resonance Imaging. *IEEE Trans. Med. Imag.*, *30*(9), 1649–1660.
- Guerquin-Kern, M., Van De Ville, D., Vonesch, C., Baritaux, J.-C., Pruessmann, K. P., & Unser, M. (2009). Wavelet-regularized reconstruction for rapid MRI. In 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, (pp. 193–196).
- Hadamard, J. (1902). Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, (pp. 49–52).
- Hansen, P. C. (1992). Analysis of Discrete III-Posed Problems by Means of the L-Curve. *SIAM Review*, 34(4), 561–580.
- Heidemann, R. M., et al. (2003). A brief review of parallel magnetic resonance imaging. *European Radiology*, *13*(10), 2323–2337.
- Hell, S. W., & Wichmann, J. (1994). Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy. *Opt. Lett.*, *19*(11), 780–782.
- Irani, M., & Peleg, S. (1993). Motion analysis for image enhancement: Resolution, occlusion, and transparency. J. Vis. Commun. Image Represent., 4(4), 324–335.
- Ivanov, D., et al. (2023). Magnetic resonance imaging at 9.4 T: the Maastricht journey. Magn. Reson. Mater. Phys. Biol. Med., 36(2), 159–173.
- Kang, M., & Chaudhuri, S. (2003). Super-resolution image reconstruction. IEEE Signal Processing Magazine, 20(3), 19–20.
- Kim, S., & Su, W.-Y. (1993). Recursive high-resolution reconstruction of blurred multiframe images. IEEE Trans. Image Process., 2(4), 534–539.
- Kim, S., et al. (1990). Recursive reconstruction of high resolution image from noisy undersampled multiframes. IEEE Transactions on Acoustics, Speech, and Signal Processing, 38(6), 1013–1027.
- Knoll, F., Bredies, K., Pock, T., & Stollberger, R. (2011). Second order total generalized variation (TGV) for MRI. *Magn. Reson. Med.*, *65*(2), 480–491.
- Ladd, M. E., et al. (2018). Pros and cons of ultra-high-field MRI/MRS for human application. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 109, 1–50.
- Lin, Z., & Shum, H.-Y. (2004). Fundamental limits of reconstruction-based superresolution algorithms under local translation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(1), 83–97.
- Liu, B., Abdelsalam, E., Sheng, J., & Ying, L. (2008). Improved spiral sense reconstruction using a multiscale wavelet model. In 2008 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, (pp. 1505–1508).
- Lønning, K., Putzky, P., Sonke, J.-J., Reneman, L., Caan, M. W., & Welling, M. (2019). Recurrent inference machines for reconstructing heterogeneous MRI data. *Med. Image Anal.*, 53, 64–78.
- Luong, Q. H. (2009). Advanced Image and Video Resolution Enhancement Techniques. Phd thesis, Ghent University.
- Lustig, M., Donoho, D., & Pauly, J. M. (2007). Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.*, 58(6), 1182–1195.
- Mead, J. (2020). χ^2 test for total variation regularization parameter selection. Inverse Probl. Imaging,

14(3), 401-421.

- Moerner, W. E., & Kador, L. (1989). Optical detection and spectroscopy of single molecules in a solid. *Phys. Rev. Lett.*, *62*, 2535–2538.
- Morozov, V. A. (1966). Regularization of incorrectly posed problems and the choice of regularization parameter. USSR Computational Mathematics and Mathematical Physics, 6(1), 242–251.
- Nasrollahi, K., & Moeslund, T. B. (2014). Super-resolution: a comprehensive survey. *Machine Vision and Applications*, 25(6), 1423–1468.
- Nguyen, N., Milanfar, P., & Golub, G. (2001). A computationally efficient superresolution image reconstruction algorithm. *IEEE Transactions on Image Processing*, 10(4), 573–583.
- Nicastro, M., Jeurissen, B., Beirinckx, Q., Smekens, C., Poot, D. H. J., Sijbers, J., & den Dekker, A. J. (2022). To shift or to rotate? comparison of acquisition strategies for multi-slice super-resolution magnetic resonance imaging. *Frontiers in Neuroscience*, 16.
- Noll, D. C., Boada, F. E., & Eddy, W. F. (1997). A spectral approach to analyzing slice selection in planar imaging: Optimization for through-plane interpolation. *Magn. Reson. Med.*, 38(1), 151–160.
- Paige, C. C., & Saunders, M. A. (1982). LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares. ACM Trans. Math. Softw., 8(1), 43–71.
- Park, S. C., Park, M. K., & Kang, M. G. (2003). Super-resolution image reconstruction: a technical overview. IEEE Signal Processing Magazine, 20(3), 21–36.
- Peeters, R. R., et al. (2004). The use of super-resolution techniques to reduce slice thickness in functional MRI. *Int. J. Imaging Syst. Technol.*, 14(3), 131–138.
- Peled, S., & Yeshurun, Y. (2001). Superresolution in MRI: Application to human white matter fiber tract visualization by diffusion tensor imaging. *Magn. Reson. Med.*, 45(1), 29–35.
- Peled, S., & Yeshurun, Y. (2002). Superresolution in MRI—perhaps sometimes. *Magn. Reson. Med.*, 48(2), 409–409.
- Pipe, J. (1998). Asymmetric sampling along k(slice-select) in two-dimensional multislice MRI. Magn. Reson. Med., 39(4), 625–634.
- Pipe, J. G. (1999). Motion correction with PROPELLER MRI: Application to head motion and free-breathing cardiac imaging. *Magn. Reson. Med.*, 42(5), 963–969.
- Plenge, E., et al. (2012). Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magn. Reson. Med.*, 68(6), 1983–1993.
- Poot, D. H. J., Jeurissen, B., Bastiaensen, Y., Veraart, J., Van Hecke, W., Parizel, P. M., & Sijbers, J. (2013). Super-resolution for multislice diffusion tensor imaging. *Magn. Reson. Med.*, 69(1), 103–113.
- Poot, D. H. J., Van Meir, V., & Sijbers, J. (2010). General and efficient super-resolution method for multi-slice MRI. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, vol. 6361, (pp. 615–622).
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). Numerical Recipes in C (2nd ed.): The Art of Scientific Computing. USA: Cambridge University Press.
- Pruessmann, K. P. (2006). Encoding and reconstruction in parallel MRI. *NMR in Biomedicine*, 19(3), 288–299.
- Regatte, R. R., & Schweitzer, M. E. (2007). Ultra-high-field MRI of the musculoskeletal system at 7.0T. J. Magn. Reson. Imaging, 25(2), 262–269.
- Robinson, D., & Milanfar, P. (2006). Statistical performance analysis of super-resolution. *IEEE Trans. Image Process.*, *15*(6), 1413–1428.
- Robson, P. M., et al. (2008). Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn. Reson. Med.*, 60(4), 895–907.
- Rousseau, F., Kim, K., Studholme, C., Koob, M., & Dietemann, J. L. (2010). On super-resolution for fetal brain MRI. Med. Image. Comput. Comput. Assist. Interv., 13 (Pt 2), 355–362.
- Rousseau, F., et al. (2006). Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images. *Acad. Radiol.*, *13*(9), 1072–1081.
- Roy, M., Upadhyaya, B., Rai, J., & Sharma, K. (2023). Resolution Enhancement of Brain MRI Images Using Deep Learning. *Engineering Proceedings*, 59(1).
- Rudin, L. I., Osher, S., & Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms.

Physica D: Nonlinear Phenomena, 60(1), 259–268.

- Sabidussi, E. R., Klein, S., Caan, M. W. A., Bazrafkan, S., den Dekker, A. J., Sijbers, J., Niessen, W. J., & Poot, D. H. J. (2021). Recurrent inference machines as inverse problem solvers for MR relaxometry. *Med. Image Anal.*, 74, 102220.
- Sabidussi, E. R., Klein, S., Jeurissen, B., & Poot, D. H. J. (2023). dtiRIM: A generalisable deep learning method for diffusion tensor imaging. *NeuroImage*, *269*, 119900.
- Scheffler, K. (2002). Superresolution in MRI? Magn. Reson. Med., 48(2), 408.
- Setsompop, K., et al. (2012). Improving diffusion MRI using simultaneous multi-slice echo planar imaging. NeuroImage, 63(1), 569–580.
- Shilling, R. Z., et al. (2006). Merging multiple stacks MRI into a single data volume. In *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006.*, (pp. 1012–1015).
- Shilling, R. Z., et al. (2009). A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI. *IEEE Trans. Med. Imag.*, 28(5), 633–644.
- Souza, A., & Senn, R. (2008). Model-based super-resolution for mri. In 2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, (pp. 430–434).
- Tian, J., & Ma, K.-K. (2011). A survey on super-resolution imaging. *Signal Image Video Process.*, 5(3), 329–342.
- Tieng, Q. M., et al. (2011). MRI resolution enhancement: How useful are shifted images obtained by changing the demodulation frequency? *Magn. Reson. Med.*, *65*(3), 664–672.
- Tournier, J. D., et al. (2019). Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *NeuroImage*, 202, 116137.
- Tsai, R., & Huang, T. (1980). Moving image restoration and registration. In *ICASSP '80. IEEE* International Conference on Acoustics, Speech, and Signal Processing, vol. 5, (pp. 418–421).
- Uecker, M., et al. (2011). Reply to: MRI resolution enhancement: How useful are shifted images obtained by changing the demodulation frequency? *Magn. Reson. Med.*, *66*(6), 1511–1512.
- Van Reeth, E., Tham, I. W., Tan, C. H., & Poh, C. L. (2012). Super-resolution in magnetic resonance imaging: A review. *Concepts Magn. Reson.*, 40A(6), 306–325.
- Van Steenkiste, G., et al. (2016). Super-resolution reconstruction of diffusion parameters from diffusionweighted images with different slice orientations. *Magn. Reson. Med.*, 75(1), 181–195.
- Van Steenkiste, G., et al. (2017). Super-resolution T1 estimation: Quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. *Magn. Reson. Med.*, 77(5), 1818–1830.
- Vogel, C. R. (2002). *Computational Methods for Inverse Problems*. USA: Society for Industrial and Applied Mathematics (SIAM).
- Walsh, D. O., Gmitro, A. F., & Marcellin, M. W. (2000). Adaptive reconstruction of phased array MR imagery. Magn. Reson. Med., 43(5), 682–690.
- Yue, L., Shen, H., Li, J., Yuan, Q., Zhang, H., & Zhang, L. (2016). Image super-resolution: The techniques, applications, and future. *Signal Processing*, 128, 389–408.
- Zhang, L., & Wu, X. (2006). An edge-guided image interpolation algorithm via directional filtering and data fusion. IEEE Trans. Image Process., 15(8), 2226–2238.
- Zhang, X., Lam, E. Y., Wu, E. X., & Wong, K. K. Y. (2008). Application of Tikhonov Regularization to Super-Resolution Reconstruction of Brain MRI Images. In X. Gao, et al. (Eds.) *Medical Imaging* and Informatics, (pp. 51–56). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Zimmerman, R. A., Wendell, A. G., & Carmody, R. F. (2000). *Neuroimaging: Clinical and Physical Principles*. Springer-Verlag, New York, Inc.

Part III

Contributions

5

Joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework : a simulation study

Contents

5.1	Introduction	
5.2	Theory	
5.3	Simulation experiments 128	
5.4	Results and discussion	
5.5	Conclusion	
Appe	endices	
5.A	Downsampling and upsampling operators	
Refe	rences	

ABSTRACT

Magnetic resonance imaging (MRI) based T_1 mapping allows spatially resolved quantification of the tissue-dependent spin-lattice relaxation time constant T_1 , which is a potential biomarker of various neurodegenerative diseases, including Multiple Sclerosis, Alzheimer disease, and Parkinson's disease. In conventional T_1 MR relaxometry, a quantitative T_1 map is obtained from a series of T_1 -weighted MR images. Acquiring such a series, however, is time consuming. This has sparked the development of more efficient T_1 mapping methods, one of which is a super-resolution reconstruction (SRR) framework in which a set of low resolution (LR) T_1 -weighted images is acquired and from which a high resolution (HR) T_1 map is directly estimated.

In this chapter, the SRR T_1 mapping framework is augmented with motion estimation. That is, motion between the acquisition of the LR T_1 -weighted images is modeled and the motion parameters are estimated simultaneously with the T_1 parameters. Based on Monte Carlo simulation experiments, we show that such an integrated motion/relaxometry estimation approach yields more accurate T_1 maps compared to a previously reported SRR based T_1 mapping approach.

The work in this chapter has been published as:

Beirinckx, Q., Ramos-Llordén, G., Jeurissen, B., Poot, D. H. J., Parizel, P. M., Verhoye, M., Sijbers, J., den Dekker, A. J., "Joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework: a simulation study", in <u>Fundamenta Informaticae</u>, Vol. 172, no. 2, pp. 105-128, 2020. doi: 10.3233/FI-2020-1896

5.1 Introduction

 T_1 mapping is a quantitative Magnetic Resonance Imaging (MRI) technique that generates maps of the tissue-specific spin-lattice relaxation time T_1 (Deoni et al., 2003). There is growing evidence that T_1 mapping can be applied to detect subtle microscopic tissue damage, with potential for earlier diagnosis of various brain diseases including multiple sclerosis (Larsson et al., 1989; Vrenken et al., 2006; Papadopoulos et al., 2010), epilepsy (Conlon et al., 1988) and Alzheimer's disease (Erkinjuntti et al., 1987). Despite these promising results, T_1 mapping currently remains a research tool and is not yet part of routine clinical assessment. The main obstacle for clinical adaptation of T_1 mapping is that conventional T_1 mapping techniques require long scan times to achieve adequate accuracy, precision and spatial resolution.

The gold standard method for T_1 mapping, the inversion recovery method, is presented in Fig. 5.1. When an object is placed in a strong magnetic field, its nuclear spins align to this magnetic field, resulting in a net magnetic moment oriented in the so-called longitudinal direction (i.e., the direction parallel to the external magnetic field, corresponding with the z-direction in Fig. 5.1). Next, this equilibrium state is disturbed by applying a 180° radio frequency (RF) pulse, which inverts the longitudinal magnetization. After this pulse, the spins start to relax back towards the equilibrium state with a time constant T_1 (Taylor et al., 2016). After inversion time TI, the longitudinal component is tipped into the transverse plane by a 90° RF pulse, after which the (T_1 -weighted) MR signal is measured. In this way, T_1 -weighted images are acquired at different inversion times. Subsequently, a T_1 map is estimated by voxel-wise fitting a parametric model to these images. Since many images are required in such an acquisition scheme, conventional T_1 mapping suffers from long acquisition times.

A simple way to reduce acquisition time is to lower the number of T_1 -weighted images. This, however, results in a loss of precision in the estimated T_1 map. Alternatively, a large number of T_1 -weighted images can be acquired in a short acquisition time by reducing the acquisition time of each individual T_1 -weighted image by lowering their spatial resolution. Commonly, this is done by acquiring multi-slice images where the slice-thickness is much larger than the spatial resolution within the slice, i.e., the through-plane resolution is much lower than the in-plane resolution. Additionally, increasing the slice thickness increases the signal-to-noise ratio (SNR) of the T_1 -weighted images, as signal strength scales linearly with imaged volume. However, thicker slices also suffer from increased partial volume effects, which arise when different tissues occur within a single voxel. In summary, reducing the acquisition time in conventional T_1 mapping is clearly a trade-off in which faster scanning comes at the cost of either a lower precision or a lower spatial resolution and increased partial volume effects, of the resulting T_1 map.

To improve this trade-off, a super-resolution reconstruction (SRR) method was recently proposed that estimates a 3D high resolution T_1 map with isotropic voxel size from a set of low resolution T_1 -weighted multi-slice images with different slice orientations and anisotropic voxel size (Van Steenkiste et al., 2017). These low resolution images are acquired at a high in-plane resolution and a low through-plane resolution. It was shown that this method indeed provides a better trade-off between resolution, precision and acquisition time than direct high-resolution acquisition (Poot et al., 2010). In this approach, motion was compensated for by adjusting the transformation parameters constituting the motion operator


Figure 5.1: A: Inversion Recovery Sequence: The longitudinal net nuclear magnetization vector is inverted by a 180° pulse. After inversion time TI, the longitudinal component is tipped into the transverse plane by a 90° pulse, after which the (T_1 -weighted) MR signal is measured. By acquiring multiple MR signals (images) at different inversion times, the recovery of the longitudinal magnetization towards its equilibrium value can be sampled. The time between two repetitions of the sequence, i.e. the time between the inversion pulses, is called the repetition time TR. **B**: Effect of the inversion recovery sequence on the net nuclear longitudinal magnetization vector $M_{z'}$ as seen in the RF-rotating frame. B_0 represents the external magnetic field vector. (a) Initial net nuclear longitudinal magnetization in alignment with B_0 , (b) the 180° inverts the longitudinal magnetization $M_{z'}$, (c)-(d) the longitudinal magnetization $M_{z'}$ relaxes and recovers to equilibrium, (e) after an inversion time TI the relaxing longitudinal magnetization $M_{z'}$ is tipped into the transverse plane by a 90° pulse before readout.

in a preprocessing step, and fixing these parameters in the SR-T1 estimation routine that followed. Fixing the motion parameters, however, may lead to inaccurate (i.e., biased) T_1 maps since no feedback mechanism is present in the SR-T1 estimation routine that can undo incorrect fixation of motion parameters. As such, errors that potentially exist in the motion estimation step might propagate into the T1 estimation. At the same time, another recent work has proposed a unified Maximum Likelihood framework for simultaneous motion and T_1 estimation in non-super-resolution T_1 mapping (Ramos-Llordén et al., 2017). It was demonstrated that the joint incorporation of the relaxation model, the motion model as well as the data statistics provide substantially more accurate motion and T_1 parameter estimates. In the present chapter, we explore, by means of simulation experiments, the potential of combining both approaches, resulting into joint Maximum Likelihood estimation of T_1 and motion in a super-resolution framework.

The remainder of this chapter is organized as follows. In Section 5.2, the image acquisition model, the proposed joint Maximum Likelihood estimator (MLE) and its implementation are described. Section 5.3 describes the simulation experiments, of which the results are presented and discussed in Section 5.4. Finally, in Section 5.5 conclusions are drawn.

5.2 Theory

The proposed method starts from a set of $N T_1$ -weighted multi-slice images, each with a different slice direction. The multi-slice images, which are assumed to have a high in-plane resolution and a low through-plane resolution, will be referred to as the low resolution (LR) images. That is, the slice thickness, or through-plane voxel size, is larger than the in-plane voxel size, leading to anisotropic voxels. The method then estimates a high-resolution (HR) T_1 map with isotropic voxels from a set of LR multi-slice T_1 -weighted images and simultaneously estimates the motion between the acquisition of these LR images.

In the derivation of the imaging model, we will assume that the LR T_1 -weighted images are acquired with a multi-slice inversion recovery (IR) conventional spin echo (SE) sequence, being the gold standard sequence for T_1 mapping (Hahn, 1949; Drain, 1949; Crawley & Henkelman, 1988).

5.2.1 MR imaging model

Let $T_1 = (T_1(j)) \in \mathbb{R}^{N_r \times 1}$ be the vector containing the values of the unknown T_1 map at the HR grid points $\{x_j\}$ (with $x_j \in \mathbb{R}^{3\times 1}$ and j the HR voxel index, $j = 1, ..., N_r$). Furthermore, let $s_n \in \mathbb{R}^{N_s \times 1}$, with n = 1, ..., N, denote the vector containing the intensities of the noiseless LR T_1 -weighted multi-slice image with slice direction n, acquired with inversion time TI_n , and sampled at the LR grid points $\{y_{nl}\}$ (with $y_{nl} \in \mathbb{R}^{3\times 1}$ and l the LR voxel index, $l = 1, ..., N_s$). To derive the mathematical relation between the HR T_1 map of interest, T_1 , and the LR image s_n , we now first introduce the virtual, noise free, HR T_1 -weighted image $r_n = (r_n(j)) \in \mathbb{R}^{N_r \times 1}$, which is assumed to be acquired with the same inversion time TI_n as s_n and sampled at the (nonrotated) HR grid points of T_1 .

Then, r_n can be modeled as a function of T_1 and a quantity $\rho = (\rho_j) \in \mathbb{R}^{N_r \times 1}$, which is proportional to the proton density (Bernstein et al., 2004):

$$r_n(j) = \rho(j) \left(1 - 2e^{-\frac{Tl_n}{T_1(j)}} \right),$$
 (5.1)

where we have assumed a perfect inversion pulse of 180° and a repetition time $TR \gg T_1$. In the remainder of this chapter, T_1 and ρ will be referred to as *relaxation model parameters* to be estimated.

Mathematically, we can now express the LR image s_n as the result of applying a sequence of operators on the virtual HR image r_n :

$$\boldsymbol{s}_n = \boldsymbol{D} \boldsymbol{B} \boldsymbol{G}_n \boldsymbol{M}_{\boldsymbol{\theta}_n} \boldsymbol{r}_n, \tag{5.2}$$

with $M_{\theta_n} \in \mathbb{R}^{N_r \times N_r}$, $G_n \in \mathbb{R}^{N_r \times N_r}$, $B \in \mathbb{R}^{N_r \times N_r}$ and $D \in \mathbb{R}^{N_s \times N_r}$ linear operators that describe, respectively, unintended motion, a known geometric transformation, spatially invariant blurring, and downsampling. In this contribution, we assume the unintended motion M_{θ_n} to be rigid, parameterized by

$$\boldsymbol{\theta}_n = (t_{xn}, t_{yn}, t_{zn}, \alpha_n, \beta_n, \gamma_n)^T, \qquad (5.3)$$

with t_{xn} , t_{yn} , t_{zn} the translation parameters and α_n , β_n , γ_n the Euler angles of three elementary rotation matrices that describe rotation around the x, y and z axis, respectively. The superscript T in Eq. (5.3) denotes the transpose operation. In the present contribution, we use the same implementation of the rigid motion operator M_{θ} as in (Ramos-Llordén et al., 2017), where they used the fact that rotation matrices can be decomposed as the product of three shear matrices. Each of the shearings can be implemented efficiently with Fast Fourier Transforms (FFT). Translation is implemented using the FFT as well. The operator G_n applies a known geometric transformation that models the image acquisition with a specific slice direction. More specifically, operator G_n models the SRR acquisition, in which multiple LR T_1 -weighted images at different orientations are acquired by rotation of the acquisition plane for each image around one fixed encoding axis. In our implementation, operator G_n is a simplified version of M_{θ} that models rotation around one fixed encoding axis. However, whereas M_{θ_n} models the effect of unintended motion, which is unknown and has to be estimated from the data, the geometric transformation G_n is known and determined by the prescribed slice direction of the LR image s_n .

The blurring operator \boldsymbol{B} describes the point spread function (PSF) of the MRI acquisition process, which can be modeled as tensor product of the PSF in three orthogonal directions: the through-plane (i.e., slice-selection) direction and the two in-plane directions, which are known as the phase- and frequency-encoding direction. We currently consider the PSF in the through-plane direction only. This through-plane PSF depends on the slice selection method. In this incipient contribution, the operators \boldsymbol{B} and \boldsymbol{D} describing spatially invariant blurring and downsampling, respectively, are combined into one operator $\overline{\boldsymbol{D}}$ that performs downsampling by averaging along the through-plane direction (Li et al., 2014). Details about this specific $\overline{\boldsymbol{D}}$ operator (and corresponding adjoint operator) are given in Appendix 5.A.

For convenience of expression, we define $A_n = \overline{D}G_n M_{\theta_n}$, and rewrite (5.2) as

$$\boldsymbol{s}_n = \boldsymbol{A}_n \boldsymbol{r}_n, \tag{5.4}$$

with $A_n = (a_n(I,j)) \in \mathbb{R}^{N_s \times N_r}$. By combining Eqs. (5.1) and (5.4), the noiseless signal in voxel *I* of the LR T_1 -weighted image can be described in terms of the HR maps $T_1 = (T_1(j))$ and $\rho = (\rho(j))$:

$$s_n(l; \mathbf{T_1}, \boldsymbol{\rho}) = \sum_{j=1}^{N_r} a_n(l, j) \rho(j) \left(1 - 2e^{-\frac{T \ln n}{T_1(j)}} \right).$$
(5.5)

In this contribution, we will assume that magnitude images are acquired, as is common for spin echo IR sequences. The voxel intensities of magnitude images reflect only the magnitude of the longitudinal magnetization, disregarding polarity (Tofts, 2004). Hence, in the absence of noise, the magnitude images are described by $|s_n|$, with $|\cdot|$ the point-wise modulus operator.

Obviously, real-world images will be subject to noise. In this contribution, the noise is assumed to be additive, zero mean Gaussian noise. It has been shown that this is a valid assumption when the signal-to-noise ratio of the magnitude data is sufficiently high (> 3) (Gudbjartsson & Patz, 1995; Andersen & Kirsch, 1996; Constantinides et al., 1997; den Dekker & Sijbers, 2014), which is typically the case for the LR images. Hence, if we denote the acquired LR magnitude images by $\tilde{s}_n \in \mathbb{R}^{N_s \times 1}$, our image acquisition model can be described as

$$\tilde{\boldsymbol{s}}_n = |\boldsymbol{A}_n \boldsymbol{r}_n| + \boldsymbol{e}_n, \quad n = 1, \dots N, \tag{5.6}$$

with $e_n \in \mathbb{R}^{N_s \times 1}$ a vector containing zero mean Gaussian noise contributions.

5.2.2 The joint Maximum Likelihood estimator

Having derived the imaging model in section 5.2.1, this section will describe the model-based SRR framework that estimates an HR ρ and T_1 map (ρ and T_1 , respectively) simultaneously with the motion parameters θ_n , n = 1, ..., N, from a set of LR images { $\tilde{s}_1, \tilde{s}_2, ..., \tilde{s}_N$ }, using a joint Maximum Likelihood estimator (MLE). The MLE is chosen because it is asymptotically unbiased, efficient (i.e., most precise) and consistent (van den Bos, 2007). The MLE fully exploits prior knowledge on the statistical distribution of the data. In our case, the data consists of the voxels of the LR images. These voxels can be modeled as random variables that, due to the presence of noise, fluctuate about their expected values which are described by the model $|A_n r_n|$ that was derived in section 5.2.1. Assuming additive, zero-mean Gaussian distributed noise, the PDF of a voxel $\tilde{s}_n(l)$, with $l = 1, ..., N_s$, of the image \tilde{s}_n is given by:

$$p_{\tilde{s}_n(l)}(\tilde{s}_n(l); \boldsymbol{T_1}, \boldsymbol{\rho}, \boldsymbol{\theta}_n) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\left(\tilde{s}_n(l) - s_n(l); \boldsymbol{T_1}, \boldsymbol{\rho}, \boldsymbol{\theta}_n\right)\right)^2}{2\sigma^2}},$$
(5.7)

with σ the standard deviation of the noise, which in this contribution is assumed to be spatially and temporally invariant. Assuming all voxels of all T_1 -weighted LR images statistically independent, the joint PDF of all voxels is given by:

$$p_{\tilde{\mathbf{s}}}(\tilde{\mathbf{s}}; \boldsymbol{T}_{1}, \boldsymbol{\rho}, \boldsymbol{\theta}) = \prod_{n=1}^{N} \prod_{l=1}^{N_{s}} p_{\tilde{s}_{n}(l)}(\tilde{s}_{n}(l); \boldsymbol{T}_{1}, \boldsymbol{\rho}, \boldsymbol{\theta}_{n}), \qquad (5.8)$$

with $\mathbf{\tilde{s}} = (\mathbf{\tilde{s}}_1^T, \dots, \mathbf{\tilde{s}}_N^T)^T$ and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^T, \dots, \boldsymbol{\theta}_N^T)^T$. To simplify the notation, let us define the parameter vector $\boldsymbol{\tau} = (\boldsymbol{T}_1^T, \boldsymbol{\rho}^T, \boldsymbol{\theta}^T)^T$.

To construct the MLE of τ , the *likelihood function* $L(\tau|\mathbf{\tilde{s}})$, which is the joint PDF of (5.8) regarded as a function of the unknown parameter vector τ (with $\mathbf{\tilde{s}}$ fixed), is needed. The MLE $\hat{\tau}_{ML}$ of the parameter vector τ from measured data $\mathbf{\tilde{s}}$ is that value of τ that maximizes the likelihood function $L(\tau|\mathbf{\tilde{s}})$, or equivalently, the so-called log-likelihood function $\mathcal{L}_{\mathbf{\tilde{s}}}(\tau|\mathbf{\tilde{s}}) \triangleq \log L(\tau|\mathbf{\tilde{s}})$, with respect to τ , i.e.,

$$\hat{\boldsymbol{\tau}}_{\mathsf{ML}} = \arg\max_{\boldsymbol{\tau}} \mathcal{L}_{\tilde{\mathbf{s}}}(\boldsymbol{\tau}|\tilde{\mathbf{s}}).$$
 (5.9)

It follows from Eq. (5.8) that the log-likelihood function can be written as

$$\mathcal{L}_{\tilde{\mathbf{s}}}(\boldsymbol{\tau}|\tilde{\mathbf{s}}) = -NN_{s}\ln\left(\sqrt{2\pi}\sigma\right) - \frac{1}{2\sigma^{2}}\sum_{n=1}^{N}\sum_{l=1}^{N_{s}}\left(\tilde{s}_{n}(l) - s_{n}(l;\boldsymbol{T}_{1},\boldsymbol{\rho},\boldsymbol{\theta}_{n})\right)^{2}.$$
(5.10)

Hence, the ML estimator $\hat{\tau}_{\rm ML}$ is equal to the ordinary (unweighted) least-squares estimator:

$$\hat{\boldsymbol{\tau}}_{\mathsf{ML}} = \arg\min J(\boldsymbol{\tau}),$$
 (5.11)

with

$$J(\tau) = \sum_{n=1}^{N} \sum_{l=1}^{N_s} (\tilde{s}_n(l) - s_n(l; T_1, \rho, \theta_n))^2.$$
 (5.12)

The non-linear optimization problem in Eq. (5.11) can be solved using the alternating minimization method, also known as the cyclic block-coordinate descent (cBCD) method (Fessler & Kim, 2011; Beck & Tetruashvili, 2013). In this method, the parameter vector τ is split into blocks and the cost function $J(\tau)$ is successively minimized with respect to each block in a cyclic order. In our case, we use a split into two blocks that contain the motion parameters and relaxation model parameters, respectively. In this way, the large-scale optimization problem in Eq. (5.11) is separated into more easily solvable problems (Ramos-Llordén et al., 2017). Moreover, it can be shown that this cBCD method assures a convergence property where $J(\tau)$ decreases at every iteration (Fessler & Kim, 2011). Convergence to at least a local minimum is guaranteed (Fan et al., 1998). In summary, the joint MLE is obtained by the following iterative recursive procedure:

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \arg\min_{\boldsymbol{\theta}} J(\hat{\boldsymbol{T}}_{1}^{(t)}, \hat{\boldsymbol{\rho}}^{(t)}, \boldsymbol{\theta})$$
(P.1)

$$\hat{T}_{1}^{(t+1)}, \hat{\rho}^{(t+1)} = \arg\min_{T_{1},\rho} J(T_{1},\rho,\hat{\theta}^{(t+1)})$$
(P.2)

with $\hat{\theta}^{(0)} = \theta_{\text{ini}}$, $\hat{\rho}^{(0)} = \rho_{\text{ini}}$ and $\hat{T}_1^{(0)} = T_{1\text{ini}}$ the initial values of the parameters θ , ρ and T_1 , respectively. By its definition, this procedure produces a nonincreasing sequence of cost function values (Beck & Tetruashvili, 2013). The procedure is terminated when the number of iterations exceeds t_{max} or when $\mathcal{E}^{(t)} < \mathcal{E}_{\text{min}}$, where $\mathcal{E}^{(t)} = J(\hat{\tau}^{(t-1)}) - J(\hat{\tau}^{(t)})$, and consecutive iterations are started from $\mathcal{E}^{(0)} = r\mathcal{E}_{\text{min}}$, with $r \in \mathbb{R}_{>1}$. The pseudo-code of the joint MLE algorithm is presented in Algorithm 1.

5.2.3 Implementation

For the proposed joint MLE algorithm, in which the non-linear optimization problem is solved in alternating fashion between problems (P.1) and (P.2), the motion estimation problem (P.1) adopts a particularly simple structure when the relaxation model parameters are fixed. Assuming no dependence of $\{\theta_n\}_{n=1}^N$ through index *n*, as is done here, the motion estimation problem can be decoupled into *N* independent minimization problems that can be evaluated in parallel. In the absence of additional information, a natural choice for the initialization of the motion parameters in the first iteration is a zero-motion initial condition such that the rigid motion operator $M_{\theta_{ini}} = I$, with *I* the identity matrix. Initial values ρ_{ini} and T_{1ini} were obtained by voxel-wise NLLS fitting the modulus of the relaxation model in Eq. (5.1) to the upsampled LR images with a Levenberg-Marquardt (Nocedal & Wright, 2006) algorithm,

Algorithm 1: Joint MLE

 $\begin{array}{l} \hline \text{Input: LR images } \tilde{s} \text{ and initial values } \theta_{\text{ini}}, \rho_{\text{ini}} \text{ and } T_{1\text{ini}} \\ \hline \text{Output: ML estimates } \hat{\theta}_{\text{ML}}, \hat{T}_{1\text{ML}} \text{ and } \hat{\rho}_{\text{ML}} \\ \hline \text{Set } t \leftarrow 0 \text{ and } \hat{\theta}^{(0)}, \hat{T}_{1}^{(0)}, \hat{\rho}^{(0)} \leftarrow \theta_{\text{ini}}, \rho_{\text{ini}}, T_{1\text{ini}}; \\ \hline \mathcal{E}^{(0)} = r\mathcal{E}_{\text{min}}, \text{ with } r \in \mathbb{R}_{>1}; \\ \hline \text{while } \mathcal{E}^{(t)} \geq \mathcal{E}_{\text{min}} \text{ and } t < t_{\text{max}} \text{ do} \\ & \text{Solve (P.1) to get } \hat{\theta}^{(t+1)}; \\ & \hat{\theta}^{(t+1)} = \arg\min_{\theta} J(\hat{T}_{1}^{(t)}, \hat{\rho}^{(t)}, \theta), \text{ started from } \theta \leftarrow \hat{\theta}^{(t)}; \\ & \text{Solve (P.2) to get } \hat{\rho}^{(t+1)}, \hat{T}_{1}^{(t+1)}; \\ & \hat{T}_{1}^{(t+1)}, \hat{\rho}^{(t+1)} = \arg\min_{T_{1},\rho} J(T_{1}, \rho, \hat{\theta}^{(t+1)}), \text{ started from } T_{1}, \rho \leftarrow \hat{T}_{1}^{(t)}, \hat{\rho}^{(t)}; \\ & \text{Calculate } \mathcal{E}^{(t+1)} = J(\hat{T}_{1}^{(t)}, \hat{\rho}^{(t)}, \hat{\theta}^{(t)}) - J(\hat{T}_{1}^{(t+1)}, \hat{\rho}^{(t+1)}, \hat{\theta}^{(t+1)}); \\ & \text{Set } t \leftarrow t+1; \\ \hline \text{end} \\ & \hat{\theta}_{\text{ML}} = \hat{\theta}^{(t)}, \hat{T}_{1\text{ML}} = \hat{T}_{1}^{(t)} \text{ and } \hat{\rho}_{\text{ML}} = \hat{\rho}^{(t)}; \\ & \text{return } \hat{\theta}_{\text{ML}}, \hat{T}_{1\text{ML}}, \hat{\rho}_{\text{ML}}; \\ \hline \end{array}$

using the MATLAB routine lsqnonlin. Upsampling was performed using the adjoint operator sequence $G_n^T \overline{D}^T$ acting on the LR images \tilde{s}_n , followed by application of $|\cdot|$, the point-wise modulus operator, to regain magnitude images.

The cost functions of LS problems (P.1) and (P.2) were minimized with a trust-region Newton algorithm using the MATLAB routine fminunc. Explicit analytical gradients were supplied for LS problem (P.1), while both explicit analytical gradients and implicit Hessian matrix elements (in the form of a matrix multiplication routine) were supplied for LS problem (P.2). The matrix multiplications in Eq. (5.2) were implemented by splitting the transformation operators M_{θ_n} and G_n in sets of shear operations, each of which can be efficiently applied as a filtering operation in the frequency domain (Ramos-Llordén et al., 2017).

The computational complexity of the joint MLE algorithm is primarily defined by the FFT operations that are part of the implementation of operators G_n and M_{θ_n} in Eq. (5.2). Using that a Q element 1D FFT has computational complexity of $\mathcal{O}(Q \log_2(Q))$, we derived that a single step of problem (P.1) required $\mathcal{O}(69M^3 \log_2(M^2) + 5M^6 \log_2(M^6))$ floating point operations, given that G_n and M_{θ_n} operate on HR images with isotropic dimensions $M \times M \times M$. This includes the operations introduced by the explicit analytical expressions for the gradient of the objective function of problem (P.1) w.r.t. motion parameters θ_n . Furthermore, the given number of floating point operations should be multiplied by a factor N-1, since problem (P.1) is optimized in parallel manner. In addition, a single step of problem (P.2) requires $\mathcal{O}(N \cdot (48M^3 \log_2(M^2) + 4M^6 \log_2(M^6)))$ floating point operations. This number also includes the additional FFT operations introduced by the analytical gradient and Hessian expression for the objective function w.r.t. the T_1 parameters. The computational requirements of operator \overline{D} were less demanding with one call of operator \overline{D} calculated up to 10 times faster than G_n , and up to 50 times faster than M_{θ_n} , for the considered phantom size.

5.3 Simulation experiments

In this section, we describe the Monte Carlo (MC) simulation experiments that were carried out to evaluate the performance of the proposed joint MLE and to compare it with:

- SRR-T1: SR least squares (LS) estimation without motion correction. In this approach, the motion is simply ignored. That is, the least-squares criterion (Eq. (5.12)) is minimized with respect to the relaxation model parameters only, while fixing the motion parameters at $\theta = 0$.
- SRR-T1-MI: Mutual Information (MI) based registration prior to SR LS estimation. In this approach, the LR images $\{\tilde{s}_1, \tilde{s}_2, \ldots, \tilde{s}_N\}$ are first upsampled by applying the adjoint operator A_n^T . Next, the upsampled (HR) images are registered using a mutual information image similarity metric (Mattes et al., 2003; Pluim et al., 2003). The motion parameters that result from this procedure are then substituted in the LS criterion (Eq. (5.12)), which is then minimized with respect to the relaxation model parameters only.
- SRR-T1-PRE: SRR T1 mapping described in (Van Steenkiste et al., 2017). In this approach, the motion parameters are estimated in a preprocessing step prior to the estimation of the HR T_1 and ρ map. For this purpose, an iterative model-based motion correction scheme is used. First, the LR images $\{\tilde{s}_1, \tilde{s}_2, \ldots, \tilde{s}_N\}$ are upsampled by applying the adjoint operator A_n^T . The first time that this upsampling is performed, the motion operator M_{θ_n} that co-constitutes A is set equal to the identity matrix. Next, a T_1 and ρ map are estimated by voxel-wise fitting the modulus of model (5.1) to these upsampled images, using the Levenberg-Marquardt algorithm. Based on the thus obtained HR T_1 and ρ maps, LR images are generated using model (5.5). These LR images are then rigidly aligned with the LR images $\{\tilde{s}_1, \tilde{s}_2, \ldots, \tilde{s}_N\}$ by minimizing their mean squared difference. The resulting motion parameter estimates are then used to update M_{θ_n} . All steps are repeated until the stopping criterion is met with $\mathcal{E}_{\min} = 10^{-4}$. The motion parameters are then fixed and a HR T_1 and ρ map is estimated using a LS estimator, with as initial values for T_1 and ρ the values that resulted from the motion correction procedure.

Details about the specific simulation settings for each of the described MC simulation experiments are summarized in Table 5.1. Information about the initialization of (P.1) and (P.2) is also summarized in Table 5.1 for each simulation experiment. To perform realistically adequate simulation experiments, a set of 2D multi-slice IR-SE T_1 -weighted LR images affected by inter-image motion (as in Eq. (5.6)) and noise was modeled from the ground truth T_1 and proton density maps. These ground truth maps were based on a simple cubic $12 \times 12 \times 12$ numerical phantom that was adopted from (Van Steenkiste et al., 2017). The phantom maps consisted of distinct regions, representing grey and white matter tissue parameters that were characterized using reported T_1 and ρ (proton density) values in human brain tissue at 3T (Wright et al., 2008). The ground truth T_1 values for grey and white matter were 1607 ms and 838 ms, respectively, while those for ρ were set at 0.86 and 0.77 for the respective regions. An overview of this numerical phantom, which we refer to as Phantom 1, is shown in Fig. 5.2.

From these ground truth T_1 and ρ maps, a set of N = 14 noiseless LR T_1 -weighted magnitude images was simulated using the forward model (5.6). The dimensions of each LR image were

Table 5.1: Simulation settings for the different Monte Carlo experiments: SR LS estimation without motion correction (SRR-T1), Mutual information based registration prior to SR LS estimation (SRR-T1-MI), SRR T1 mapping with preprocessing loop (SRR-T1-PRE), the proposed Joint MLE (SRR-T1-JMLE).

	SRR-T1	SRR-T1-MI	SRR-T1-PRE	SRR-T1-JMLE
Dimension of HR maps	$12 \times 12 \times 12$	$12 \times 12 \times 12$	$12 \times 12 \times 12$	$12 \times 12 \times 12$
Dimension of LR images	$12 \times 12 \times 6$	$12 \times 12 \times 6$	$12 \times 12 \times 6$	$12 \times 12 \times 6$
Spatial SNR	SNR∈[20, 30,, 100]ª	snr∈[20, 30, , 100]ª	snr∈[20, 30, , 100]ª	SNR∈[20, 30,, 100] ^a
Number of images N	14	14	14	14
Inversion times TI _n [s]	$TI_n \in [0.1, \ldots, 8]^b$	$TI_n \in [0.1, \ldots, 8]^b$	$TI_n \in [0.1, \ldots, 8]^b$	$TI_n \in [0.1, \ldots, 8]^b$
# slice orientations	7	7	7	7
# TI per slice orientation	2	2	2	2
Slice orientation angles [°]	0:(180/7):154.28	0:(180/7):154.28	0:(180/7):154.28	0:(180/7):154.28
Initialization of (P.1)	n/a	MI registration ^c	preprocessing loop ^d	zero-motion IC ^e
Initialization of (P.2)	vw-NLLS-LM ^f	vw-NLLS-LM ^f	vw-NLLS-LM ^f	vw-NLLS-LM ^f
Optim. algorithm of (P.1)	n/a	n/a, fixed ^g	n/a, fixed ^g	trust-region Newton ^h
Optim. algorithm of (P.2)	trust-region Newton ⁱ	trust-region Newton ^{g,i}	trust-region Newton ^{g, i}	trust-region Newton ⁱ

^aDuring simulations, nine different SNR values were studied ranging from 20 to 100, sampled with step size 10.

^bEach of the LR images $\{\theta_n\}_{n=1}^N$, with N = 14, had a unique inversion time TI_n . The logarithms of these inversion times $\{\mathsf{TI}_n\}_{n=1}^N$ were equidistantly spaced between $T_1 = 0.1s$ and $T_{14} = 8s$.

^cThe upsampled LR images are pairwise registered using MATLAB's imregtform function (MATLAB, 2017). During the rigid registration process, the number of multi-level image pyramid levels is equal to two, and the first image of the series is chosen as a reference, hence $\theta_1 = 0$. The one-plus-one evolutionary optimizer configuration is used, for which the number of iterations is set to a very high value (> 5000) to ensure convergence of the motion parameter estimation. The remaining MI registration parameters are set to the default values of the MATLAB built-in code. Assuming no dependence of the motion parameters $\{\theta_n\}_{n=1}^N$ through the index *n*, the different pairwise registration problems can be decoupled into N-1 sub-problems, which can be implemented very efficiently with MATLAB parallel computing tools.

^dPreprocessing loop: the iterative model-based motion correction scheme is used, as described in (Van Steenkiste et al., 2017).

^ezero-motion IC: In the absence of additional information, a natural choice for the initialization of the motion parameters in the first iteration is a zero-motion initial condition (IC) such that the rigid motion operator $M_{\theta_{ini}} = I$, with I the identity matrix.

^fvw-NLLS-LM: voxel-wise NLLS fitting the modulus of relaxation model in (5.1) to upsampled LR images with Levenberg-Marquardt algorithm, using MATLAB's lsqnonlin routine, with the initial estimate per voxel chosen equal to $[\rho, T_1] = [0.5, 1.5]$. Upsampling was performed using the adjoint operator sequence $G_n^T \overline{D}^T$ followed by application of $|\cdot|$, the point-wise modulus operator.

 $^{g}\mbox{Problem}$ (P.2) is solved for the relaxation model parameters only, while fixing the motion parameters at those that result from the initialization procedure.

^hMATLAB's fminunc routine, implemented with explicit analytical expressions for the gradient of the objective function.

MATLAB's fminunc routine, implemented with explicit analytical expressions for the gradient of the objective function and implicit Hessian matrix elements (in the form of a matrix multiplication routine).



Figure 5.2: Phantom 1: Overview of the ground truth HR maps, and visualization of the downsampling along the slice dimension for one LR image.

equal to $12 \times 12 \times 6$. The anisotropy factor (AF) was equal to 2. This AF is defined as the ratio between the through-plane slice thickness and the (isotropic) in-plane voxel size in the frequency encoding and phase encoding direction, see also Fig. 5.3. For each LR image N, the translational shifts t_{xn} , t_{yn} , t_{zn} and Euler angles α_n , β_n , γ_n that define the rigid motion parameters $\{\theta_n\}_{n=1}^N$ in Eq. (5.3) were generated randomly from a uniform distribution on the interval [-1, 1] (voxel units) and [-5, 5] degrees respectively. The reference image was chosen to be $\mathbf{\tilde{s}}_1$, hence $\theta_1 = \mathbf{0}$. The same set of randomly generated rigid motion parameters was used for all simulation experiments. Furthermore, MATLAB's intrinsic coordinate system is used to represent 3D images, for which the origin is chosen at the center of the 3D image.

To account for wraparound artefacts that stem from the use of FFT to perform rotations, appropriate zeropadding was performed in each direction. Furthermore, it should be noted that since in our simulations the $M \times M \times M$ discrete sampled image volume to be rotated had an even number of sampling points in each direction (i.e., M was even), the FFT based procedure required an extra multiplication with an exponential phase factor to obtain real values after rotation (Larkin et al., 1997).

To fully recover the HR information, the LR images need to contain complementary information about the phantom. Rotation in image space corresponds to a rotation in frequency domain. As previously argued (Plenge et al., 2012), acquiring the LR images with different slice orientations ensures that each LR image covers a different part of k-space (Fig. 5.4). In this way, the LR data will contain high spatial frequencies in all three dimensions. This approach results in more effective sampling of k-space than shifting the LR images by subpixel distances along the slice selection direction. In the latter case, the SR reconstruction result relies heavily on the success of recovering the aliased high frequency in the slice direction, since the narrow slice selection frequency band covers exactly the same part of the k-space for each LR image. Similar to the acquisition protocol in (Van Steenkiste et al., 2016, 2017), the rotation of the LR images was performed around the virtual phase encoding axis in increments of $180/N_o$ degrees, with N_o the number of slice orientations. Aiming at a short acquisition time, the number of slice orientations was kept low, but sufficiently high to ensure that the k-space is maximally covered. In particular, the number of slice orientations was fixed to 7, corresponding to having two different LR T_1 -weighted magnitude images per slice orientation. An overview of the slice orientations and corresponding inversion times (TI), combined with their coverage of k-space is shown in Fig. 5.4. Each LR T_1 -weighted magnitude image had a unique inversion time, with $\mathsf{TI}_n \in [0.1, \ldots, 8]s$, where the TIs were sampled equidistantly in log-space.

After fixing the motion parameters and acquisition geometry, downsampling by averaging and the application of the point-wise modulus operator for magnitude images in accordance with Eq. 5.6, zero mean white Gaussian noise was added to the LR images. The noise level was chosen to obtain SNR values between 20 and 100, where SNR is the ratio of the spatial mean of the LR image with the highest TI and the standard deviation of the noise. For each SNR, $N_{\rm MC} = 140$ realizations of sets of LR images were generated. For each realization, a HR T_1 and ρ map as well as the motion parameters θ were estimated.

The different simulations experiments were implemented in MATLAB (MATLAB, 2017), and run on a computer with an Intel i7-6850K hexa-core CPU @ 3.6 GHz and 32 GB of RAM. The proposed joint MLE simulation experiment for the numerical phantom required around 4 GB RAM (allocated memory usage), and for a fixed tolerance of $\mathcal{E}_{min} = 10^{-4}$, on average



Figure 5.3: Schematic comparison between image space and k-space (2D and 3D view) for a multi-orientation low resolution acquisition. The anisotropy of the voxels in image space is defined by the anisotropy factor, $AF = \frac{b}{a}$ with *b*, the slice thickness, and *a*, the voxel size in the frequency encoding direction (and phase encoding direction). Since we choose to rotate only around the phase encoding axis, the k-space can only be sampled in a cylinder with diameter $\frac{1}{2}$.



Figure 5.4: Overview of the slice orientations of the LR images and corresponding inversion times TI_n (top row) and the *k*-space sampling strategy (middle and bottom row). The middle row shows the *k*-space sampling of seven individual LR images, each having a different slice orientation, whereas the bottom row shows the overlap in *k*-space when those images are combined. The shaded area denotes the sampled *k*-space, while the white region is not sampled.

about 14 alternating MLE iterations were needed to ensure convergence of the optimization process. In order to run the series of MC simulations quickly and efficiently, the University of Antwerp's High Performance Computing core (HPC) facility CalcUA was used.

To assess the performance of each method to estimate the T_1 map, the following performance measures were used (Ramos-Llordén et al., 2017):

- (a) Relative bias. The bias quantifies the accuracy or, equivalently, the systematic error of the estimator (van den Bos, 2007). For each voxel, the relative sample bias was calculated as $(\overline{\hat{T}}_1 T_1)/T_1$, where $\overline{\hat{T}}_1$ is the sample mean of the N_{MC} estimates \widehat{T}_1 and T_1 is the true value. A measure of the overall accuracy of the T_1 map was obtained by calculating the spatial mean of the absolute value of the relative sample bias.
- (b) Relative standard deviation. The standard deviation quantifies the precision, or, equivalently, the non-systematic error of the estimator (van den Bos, 2007). For each voxel, the relative sample standard deviation was calculated as $std(\hat{T}_1)/T_1$, and an overall precision measure was obtained by taking the spatial mean of these relative sample standard deviations.
- (c) Relative root-mean-square error (relative RMSE). The RMSE is a measure that incorporates both accuracy and precision. For each voxel, the relative sample RMSE was calculated as $\sqrt{(\hat{T}_1 T_1)^2}/T_1$. An overall RMSE measure was obtained by calculating the spatial mean of these relative sample RMSE values.

By substitution of ρ for T_1 , the performance of each method to estimate the ρ map was assessed in an identical way using measures (a)-(c).

To assess the ability of the proposed method to estimate motion, the following performance measure was used:

(d) Motion component root-(mean)-mean-square error (RMMSE), defined as

$$\sqrt{\frac{1}{N-1}\sum_{n=2}^{N}\overline{\left([\hat{\boldsymbol{\theta}}_{n}]_{j}-[\boldsymbol{\theta}_{n}]_{j}\right)^{2}}},$$
(5.13)

with $[\theta_n]_j$ the *j*th component of θ_n and $\overline{[\hat{\theta}_n]_j}$ the sample mean of the N_{MC} estimates $[\hat{\theta}_n]_j$.

To supplement the results of Phantom 1, also a second, more challenging phantom was created (Fig. 5.5), which consisted of distinct grey and white matter tissue regions that included both uniform regions and checkerboard patterns, combined with horizontal and vertical planar structures. This second phantom will be referred to as Phantom 2. The same ground truth T_1 values as for Phantom 1 were used to characterize grey and white matter voxels. The different Monte Carlo experiments were repeated for Phantom 2 for a fixed spatial SNR = 50, keeping the other simulation settings identical as for Phantom 1. The reconstruction results of both phantoms were visually compared w.r.t. their respective ground truth T_1 parameter maps. For each phantom and for each method, an average T_1 map was calculated by voxel-wise averaging over all N_{MC} reconstruction results for SNR = 50.



Figure 5.5: Phantom 2: Overview of the ground truth HR maps, and visualization of the downsampling along the slice dimension for one LR image.

5.4 Results and discussion

Figs. 5.6-5.8 summarize the statistical performance results that were obtained from the simulation experiments for Phantom 1. Fig. 5.6(a-c) and Fig. 5.7(a-c) show the overall relative sample bias, standard deviation and RMSE for all considered estimators of, respectively, T_1 and ρ , as a function of the SNR. The results clearly show that for the SRR-T1 method without motion correction, the estimation of the relaxation model parameters performs poorly in terms of accuracy and precision, with high relative bias and high relative standard deviation, thereby underlining the importance of a proper motion estimation framework. Furthermore, SRR-T1-MI also shows a poor performance. In terms of precision, SRR-T1-MI performs even worse than SRR-T1. This observation is quite remarkable. It may reflect the limitations of intensity-based image registration when it comes to registering images of largely different contrast, which have been reported earlier (Xue et al., 2012; Roujol et al., 2015). The poor performance of the SRR-T1-MI method may also be partly due to its implementation. In our implementation, the LR magnitude images were naively upsampled using the adjoint operator A_{a}^{T} followed by application of the point-wise modulus operator. The latter preserves the magnitude characteristic of the resulting HR images after upsampling. However, information loss is inherent and unavoidable in this upsampling process, partially due to the non-existence of an adjoint modulus operation, and this might impede correct intensity-based registration of these HR images in the next step. Finally, the small size of the images considered in our simulation experiment may also play a role, as SRR-T1-MI may be more sensitive to the image size than the other methods considered. Next, the SRR-T1-PRE method clearly improves the estimation results, both in terms of accuracy and precision. For low SNR values, this method shows even a better relative standard deviation than the SRR-T1-JMLE approach, as can be observed from Fig. 5.6(b) and Fig. 5.7(b). However, this lower precision of the proposed SRR-T1-JMLE is more than compensated by its higher accuracy, resulting in the superior performance of SRR-T1-JMLE in terms of relative RMSE over the whole range of SNR values.

Fig. 5.8 shows the motion component RMMSE for each of the six rigid motion components as a function of the SNR. SRR-T1-JMLE clearly outperforms the other methods (SRR-T1, SRR-T1-MI, and SRR-T1-PRE) in terms of the motion component RMMSE. This is particularly visible for the rotation parameter components α , β , γ , in the lower bottom half of Fig. 5.8. It is worth emphasizing once more that the motion parameter problem in the SRR-T1-JMLE method was initialized from a zero-motion initial condition, i.e. $M_{\theta_{ini}} = I$, with I the identity matrix. This also highlights the robustness of the SRR-T1-JMLE method for poor motion initialization scenarios. The quantitative performance measures for Phantom 2, for a fixed spatial SNR = 50, are summarized in Table 5.2. Results with this new phantom are very similar to those obtained with Phantom 1, underlining the superior performance of SRR-T1-JMLE in terms of relative RMSE and motion component RMMSE.



Figure 5.6: Results of the simulation experiments for Phantom 1: (a) relative T_1 bias, (b) relative T_1 standard deviation, (c) relative T_1 RMSE, as a function of SNR. Error bars correspond with the standard error of the spatial mean, but are omitted when their size matches the order of the graph symbol size.



Figure 5.7: Results of the simulation experiments for Phantom 1: (a) relative ρ bias, (b) relative ρ standard deviation, (c) relative ρ RMSE, as a function of SNR. Error bars correspond with the standard error of the spatial mean, but are omitted when their size matches the order of the graph symbol size.



Figure 5.8: Results of the simulation experiments for Phantom 1, showing the motion component RMMSE for each of the six rigid motion components, as a function of SNR. Error bars are omitted when their size matches the order of the graph symbol size.

Table 5.2: Quantitative performance measures for Phantom 2, calculated over all $N_{MC} = 140$ reconstruction results for SNR = 50, for the different Monte Carlo experiments: SR LS estimation without motion correction (SRR-T1), Mutual information based registration prior to SR LS estimation (SRR-T1-MI), SRR T1 mapping with preprocessing loop (SRR-T1-PRE), the proposed Joint MLE (SRR-T1-JMLE).

	SRR-T1	SRR-T1-MI	SRR-T1-PRE	SRR-T1-JMLE
Overall rel. T_1 bias [%]	71 ± 3	42 ± 2	6.3 ± 0.1	1.83 ± 0.03
Overall rel. T_1 std. dev. [%]	8.0 ± 0.5	46 ± 2	0.51 ± 0.01	2.50 ± 0.03
Overall rel. T_1 RMSE [%]	72 ± 3	67 ± 3	6.4 ± 0.1	3.21 ± 0.04
Overall rel. ρ bias [%]	16.1 ± 0.3	15.2 ± 0.3	2.23 ± 0.04	0.309 ± 0.005
Overall rel. ρ std. dev. [%]	1.44 ± 0.05	9.6 ± 0.1	0.204 ± 0.001	0.686 ± 0.005
Overall rel. ρ RMSE [%]	16.4 ± 0.3	19.1 ± 0.2	2.25 ± 0.04	0.772 ± 0.006
RMMSE of t_x [voxel units]	n/a	0.59 ± 0.06	0.06 ± 0.01	0.014 ± 0.002
RMMSE of t_y [voxel units]	n/a	0.53 ± 0.06	0.10 ± 0.02	0.011 ± 0.002
RMMSE of t_z [voxel units]	n/a	0.54 ± 0.05	0.05 ± 0.02	0.010 ± 0.001
RMMSE of α [degrees]	n/a	2.1 ± 0.1	0.93 ± 0.07	0.055 ± 0.004
RMMSE of β [degrees]	n/a	1.9 ± 0.1	0.91 ± 0.07	0.068 ± 0.004
RMMSE of γ [degrees]	n/a	2.4 ± 0.1	1.54 ± 0.09	0.060 ± 0.005



Figure 5.9: Visual comparison of the performance of the different SRR methods: SRR-T1 (row 1), SRR-T1-MI (row 2), SRR-T1-PRE (row 3), and our proposed SRR-T1-JMLE (row 4). On the left, three orthogonal slices of the T_1 map obtained by voxel-wise averaging over all $N_{MC} = 140$ reconstruction results for SNR=50. On the right, histograms showing the voxel data distribution of the corresponding full 3D T_1 parameter map estimates for the different simulation experiments. The ground truth T_1 values for grey and white matter, 1607 ms and 838 ms respectively, are marked with vertical lines.



Figure 5.10: SRR-T1-JMLE reconstruction result for Phantom 2: Ground Truth (row 1), SRR-T1-PRE (row 2), SRR-T1-JMLE (row 3). On the left, three orthogonal slices of the T_1 map obtained by voxel-wise averaging over all $N_{MC} = 140$ reconstruction results for SNR=50. On the right, histograms showing the voxel data distribution of the corresponding full 3D T_1 parameter map estimate. The ground truth T_1 values for grey and white matter, 1607 ms and 838 ms respectively, are marked with vertical lines.

Finally, to provide a visual comparison of the performance of the different methods, for each method an average T_1 map was calculated by voxel-wise averaging over all N_{MC} reconstruction results for SNR = 50. Fig. 5.9 shows the resulting average T_1 maps for three orthogonal slices (sagittal, axial and coronal planes) of Phantom 1, accompanied by histograms of the voxel data distribution of the corresponding full 3D average T_1 maps. The number of histogram bins was specified at 200. Note that the grey and white matter phantoms considered have histograms consisting of two distinct peaks corresponding with the ground truth T_1 values of both tissues. From Fig. 5.9, for Phantom 1, it is clear that only with SRR-T1-PRE and SRR-T1-JMLE, these peaks can be distinguished. Moreover, Fig. 5.9 clearly shows that the block-wise homogeneous structure of the phantom is best reconstructed by SRR-T1-JMLE. The same observations can be made for Phantom 2, for which the results of the visual comparison are shown in Fig. 5.10.

Overall, the results clearly demonstrate the superior performance of the SRR-T1-JMLE method in terms of accuracy and relative RMSE for both motion and T_1 and ρ estimation. This is also supported by the visual comparison presented in Fig. 5.9 and Fig. 5.10, where SRR-T1-JMLE clearly outperforms the other estimation frameworks.

In the outline of the MR imaging model under paragraph 5.2.1, the voxel intensity values of the HR images r_n were modeled by a three-parameter T_1 model (Eq. (5.1)) that depends on T_1 and ρ , for given inversion times TI_n (Bernstein et al., 2004). It should be noted that if the assumptions that substantiate the choice of this model, i.e. perfect inversion pulse of 180° and a repetition time TR $\gg T_1$, are invalid in practice, the proposed joint MLE method can still be used, but the model should be extended so as to avoid biased results. Such an extension may include the introduction of additional unknown model parameters to be estimated from the data (Barral et al., 2010), which may have a negative influence on the precision. This well-known trade-off between bias and precision should always been taken into account in model selection.

According to the computational complexity of the SRR-T1-JMLE algorithm described in section 5.2.3, increasing the volume size of the images, i.e. choosing a larger region-ofinterest (ROI), will result in longer computation times, as the number of floating point operations increases. As a possible solution to this problem, the ROI could be split in several blocks (with overlap to avoid edge artifacts), where the HR relaxation parameters are reconstructed in each block separately. This would allow parallelization of estimation problem (P.2), which would considerably reduce the computational complexity, memory consumption and computation time of the SRR-T1-JMLE method.

5.5 Conclusion

Quantitative MR T_1 mapping suffers from long acquisition times with high risk for patient motion artefacts, resulting in poor accuracy of estimated T_1 relaxometry parameters. In this contribution, we explored the potential of augmenting a recently proposed super-resolution reconstruction method for MRI T_1 mapping with simultaneous motion estimation in a maximum likelihood framework. Super-resolution reconstruction provides a better trade-off between resolution, precision and acquisition time than conventional direct high-resolution acquisition. By extending super-resolution reconstruction with simultaneous motion estimation, potential bias in the estimated T_1 map caused by motion can be substantially reduced compared to motion correction by preprocessing. By means of Monte Carlo simulation experiments, our newly proposed method was quantitatively compared against a ground-truth T_1 map together with three other approaches. The results were analysed using statistical performance measures and by performing a visual comparison. In conclusion, the results of these simulation experiments demonstrate that our newly proposed joint relaxometry and motion estimation approach yields more accurate T_1 maps than a previously reported SRR based T_1 mapping approach, in which motion registration is applied as a preprocessing step prior to T_1 mapping. Future work will focus on the validation of the proposed joint MLE method on real data scenarios, the development of advanced blurring operators for the acquisition model, and the extension of the motion model to non-rigid or affine motion. In addition, it is also worthwhile to investigate intra-image motion correction strategies and to further customize the alternating optimization scheme.

Acknowledgments

The authors gratefully acknowledge support of the Industrial Research Fund of the Antwerp University Association, the European Space Agency and BELSPO Prodex (BrainDTI), the Research Foundation Flanders (FWO Belgium) through project funding G084217N and 12M3116N.

Appendices

5.A Downsampling and upsampling operators

In the buildup of the forward model (5.4) encapsulated in operator A, operators B and D describing spatially invariant blurring and downsampling, respectively, are combined into one operator \overline{D} that performs downsampling by averaging (Li et al., 2014). Conventionally, downsampling keeps one sample out of a block and discards the remaining samples, whereas blurring takes into consideration the point spread functions of the MRI acquisition process. Downsampling by averaging is used here. For example, downsampling by a factor of 2 in 1D has matrix form,

$$\begin{bmatrix} \ddots & \ddots & & & \\ & \frac{1}{2} & \frac{1}{2} & & \\ & & & \frac{1}{2} & \frac{1}{2} & \\ & & & & \ddots & \ddots \end{bmatrix} \begin{bmatrix} \vdots \\ \boldsymbol{x}[0] \\ \boldsymbol{x}[1] \\ \boldsymbol{x}[2] \\ \boldsymbol{x}[3] \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ \frac{\boldsymbol{x}[0] + \boldsymbol{x}[1]}{2} \\ \frac{\boldsymbol{x}[2] + \boldsymbol{x}[3]}{2} \\ \vdots \end{bmatrix} .$$
(5.A.1)

In the spatial (i.e. image) domain, this downsampling from a discrete vector x[n] to y[n] can be written in more compact form as

$$\boldsymbol{y}[n] = \overline{\boldsymbol{D}}\boldsymbol{x}[n] = \frac{1}{M} \sum_{m=0}^{M-1} \boldsymbol{x}[nM+m], \qquad (5.A.2)$$

where *M* corresponds with the anisotropy factor AF, defined as the ratio of the through-plane slice thickness and the (isotropic) in-plane voxel size.

Implementing the iterative recursive procedure described by problems (P.1) and (P.2) benefits from having adjoint operators. Interestingly, upsampling is the adjoint of downsampling. Conventionally, upsampling with zero insertion is used, we use upsampling with replication. For example, upsampling by a factor of 2 in 1D has matrix form,

The upsampling matrix is the transpose of the downsampling matrix in Eq. (5.A.2). We can write the upsampled x[n] from y[n] as

$$\boldsymbol{x}[n] = \overline{\boldsymbol{D}}^{T} \boldsymbol{y}[n] = \frac{1}{M} \boldsymbol{y}\left[\left\lfloor \frac{n}{M} \right\rfloor\right],$$
 (5.A.4)

where the floor function $|\cdot|$ gives the largest integer less than or equal to its argument.

References

- Andersen, A. H., & Kirsch, J. E. (1996). Analysis of noise in phase contrast MR imaging. *Med. Phys.*, 23(6), 857–869.
- Barral, J. K., et al. (2010). A robust methodology for in vivo T1 mapping. *Magn. Reson. Med.*, 64(4), 1057–1067.
- Beck, A., & Tetruashvili, L. (2013). On the convergence of block coordinate descent type methods. *SIAM J. Optim.*, 23(4), 2037–2060.
- Bernstein, M. A., King, K. F., & Zou, X. J. (2004). Handbook of MRI pulse sequences. London: Elsevier Academic press.
- Conlon, P., Trimble, M. R., Rogers, D., & Callicott, C. (1988). Magnetic resonance imaging in epilepsy: a controlled study. *Epilepsy Res.*, 2(1), 37–43.
- Constantinides, C. D., Atalar, E., & McVeigh, E. R. (1997). Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magn. Reson. Med.*, 38(5), 852–857.
- Crawley, A. P., & Henkelman, R. M. (1988). A comparison of one-shot and recovery methods in T1 imaging. *Magn. Reson. Med.*, 7(1), 23–34.
- den Dekker, A. J., & Sijbers, J. (2014). Data distributions in magnetic resonance images: A review. *Physica Med.*, 30(7), 725–741.
- Deoni, S. C. L., Rutt, B. K., & Peters, T. M. (2003). Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Magn. Reson. Med.*, 49(3), 515–526.
- Drain, L. E. (1949). A direct method of measuring nuclear spin-lattice relaxation times. Proc. Phys. Soc. London, Sect. A, 62, 301–306.
- Erkinjuntti, T., et al. (1987). Do white matter changes on MRI and CT differentiate vascular dementia from Alzheimer's disease? *J. Neurol. Neurosurg. Psychiatry*, 50(1), 37–42.
- Fan, J., Farmen, M., & Gijbels, I. (1998). Local maximum likelihood estimation and inference. J. R. Stat. Soc. Series B. Stat. Methodol., 60(3), 591–608.
- Fessler, J. A., & Kim, D. (2011). Axial block coordinate descent (ABCD) algorithm for X-ray CT image reconstruction. In Proc. Intl. Mtg. on Fully 3D image Recon. in Rad. and Nuc. Med., (pp. 262–265).
- Gudbjartsson, H., & Patz, S. (1995). The Rician distribution of noisy MRI data. *Magn. Reson. Med.*, 34(6), 910–914.

- Hahn, E. L. (1949). An accurate nuclear magnetic resonance method for measuring spin-lattice relaxation times. *Phys. Rev.*, *76*(1), 145–146.
- Larkin, K. G., Oldfield, M. A., & Klemm, H. (1997). Fast Fourier method for the accurate rotation of sampled images. *Opt. Commun.*, 139(1-3), 99–106.
- Larsson, H. B. W., et al. (1989). Assessment of demyelination, edema, and gliosis by in vivo determination of T1 and T2 in the brain of patients with acute attack of multiple sclerosis. *Magn. Reson. Med.*, *11*(3), 337–348.
- Li, Y., Matej, S., & Metzler, S. D. (2014). Image reconstructions from super-sampled data sets with resolution modeling in PET imaging. *Med. Phys.*, 41(12), 121912.
- MATLAB (2017). MATLAB and Image Processing Toolbox Release 2017b. Natick, Massachusetts, United States: The Mathworks, Inc.
- Mattes, D., et al. (2003). PET-CT image registration in the chest using free-form deformations. *IEEE Trans. Med. Imag.*, 22(1), 120–128.
- Nocedal, J., & Wright, S. (2006). Numerical optimization. New York, NJ, USA: Springer-Verlag.
- Papadopoulos, K., et al. (2010). T1-relaxation time changes over five years in relapsing-remitting multiple sclerosis. *Mult. Scler. J.*, *16*(4), 427–433.
- Plenge, E., et al. (2012). Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magn. Reson. Med.*, 68(6), 1983–1993.
- Pluim, J. P. W., Maintz, J. B., & Viergever, M. A. (2003). Mutual-information-based registration of medical images: a survey. IEEE Trans. Med. Imag., 22(8), 986–1004.
- Poot, D. H. J., Van Meir, V., & Sijbers, J. (2010). General and efficient super-resolution method for multi-slice MRI. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, vol. 6361, (pp. 615–622).
- Ramos-Llordén, G., et al. (2017). A unified maximum likelihood framework for simultaneous motion and T_1 estimation in quantitative MR T_1 mapping. *IEEE Trans. Med. Imag.*, 36(2), 433–446.
- Roujol, S., et al. (2015). Adaptive registration of varying contrast-weighted images for improved tissue characterization (ARCTIC): Application to T1 mapping. *Magn. Reson. Med.*, 73(4), 1469–1482.
- Taylor, A. J., Salerno, M., Dharmakumar, R., & Jerosch-Herold, M. (2016). T1 Mapping Basic Techniques and Clinical Applications. JACC: Cardiovasc. Imaging, 9(1, SI), 67–81.
- Tofts, P. (2004). *Quantitative MRI of the brain: measuring changes caused by disease*. Chichester, England: John Wiley & Sons.
- van den Bos, A. (2007). Parameter Estimation for Scientists and Engineers. Hoboken, New Yersey, USA: John Wiley & Sons.
- Van Steenkiste, G., et al. (2016). Super-resolution reconstruction of diffusion parameters from diffusionweighted images with different slice orientations. *Magn. Reson. Med.*, 75(1), 181–195.
- Van Steenkiste, G., et al. (2017). Super-resolution T1 estimation: Quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. *Magn. Reson. Med.*, 77(5), 1818–1830.
- Vrenken, H., et al. (2006). Whole-brain T1 mapping in multiple sclerosis: global changes of normalappearing gray and white matter. *Radiology*, 240(3), 811–820.
- Wright, P. J., et al. (2008). Water proton T1 measurements in brain tissue at 7, 3, and 1.5 T using IR-EPI, IR-TSE, and MPRAGE: results and optimization. *Magn. Reson. Mater. Phys. Biol. Med.*, 21(1-2), 121–130.
- Xue, H., et al. (2012). Motion correction for myocardial T1 mapping using image registration with synthetic image estimation. *Magn. Reson. Med.*, 67(6), 1644–1655.

6

Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping

Contents

6.1	Introduction
6.2	Theory
6.3	Materials and Methods
6.4	Results
6.5	Discussion
6.6	Conclusion
Арр	endices
6.A	Analytical derivatives for joint MAP optimization
6.B	Model operators: implementation and derivatives
6.C	Applications in musculoskeletal imaging
Refe	rences

ABSTRACT

Quantitative Magnetic Resonance (MR) imaging provides reproducible measurements of biophysical parameters, and has become an essential tool in clinical MR studies. Unfortunately, 3D isotropic high resolution (HR) parameter mapping is hardly feasible in clinical practice due to prohibitively long acquisition times. Moreover, accurate and precise estimation of quantitative parameters is complicated by inevitable subject motion, the risk of which increases with scanning time. In this chapter, we present a model-based superresolution reconstruction (SRR) method that jointly estimates HR quantitative parameter maps and inter-image motion parameters from a set of 2D multi-slice contrast-weighted images with a low through-plane resolution. The method uses a Bayesian approach, which allows to optimally exploit prior knowledge of the tissue and noise statistics. To demonstrate its potential, the proposed SRR method is evaluated for a T1 and T2 quantitative mapping protocol. Furthermore, the method's performance in terms of precision, accuracy, and spatial resolution is evaluated using simulated as well as real brain imaging experiments. Results show that our proposed fully flexible, quantitative SRR framework with integrated motion estimation outperforms state-of-the-art SRR methods for quantitative MRI.

The work in this chapter has been published as:

Beirinckx, Q., Jeurissen, B., Nicastro, M., Poot, D. H. J., Verhoye, M., den Dekker, A. J., Sijbers, J., "Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping", <u>Computerized Medical Imaging & Graphics</u>, Vol. 100, pp. 102071, 2022. doi: 10.1016/j.compmedimag.2022.102071

6.1 Introduction

In recent decades, magnetic resonance imaging (MRI) has evolved from a qualitative imaging tool to a quantitative measurement method. Whereas qualitative MRI relies on the subjective interpretation of tissue contrast, quantitative MRI (gMRI) aims to measure reproducible and objective maps of biophysical parameters, which allows the comparison of measurements across subjects and sites, or over time (e.g., longitudinal follow-up of patients). Indeed, biophysical parameters measured by qMRI, such as relaxation times and diffusion metrics, are increasingly used as biomarkers for neurological diseases (Seiler et al., 2021), in quantitative musculoskeletal imaging (de Mello et al., 2019), or in qMRI-guided radiotherapy (van Houdt et al., 2021). Unfortunately, despite its broad range of potential applications, gMRI is not widely used in clinical practice. This is mainly because gMRI requires a series of MR images with different contrast weightings to estimate the biophysical parameter maps of interest and suffers from long scan times to provide accurate and precise parameter maps at 3D isotropic high spatial resolution. Methods have been proposed that enable reconstruction from highly under-sampled images and hence speed up image acquisition, such as modelbased reconstruction (Maier et al., 2019), low-rank approaches (Zhang et al., 2015), or the imposition of sparsity constraints (Zhao et al., 2012). However, they generally come at the cost of either a lower precision or a lower spatial resolution of the reconstructed parameter maps.

To break the trade-off between resolution, precision and acquisition time, super-resolution reconstruction (SRR) has been put forward (Greenspan et al., 2002; Van Reeth et al., 2012). In this approach, high-resolution (HR) 3D isotropic images are estimated from a set of multi-slice images with a high in-plane but low through-plane resolution, where the multi-slice images are acquired with either sub-voxels shifts in the through-plane direction (Greenspan et al., 2002), three orthogonal slice orientations (Rousseau et al., 2006; Gholipour et al., 2010; Scherrer et al., 2012; Sui et al., 2021), slice orientations rotated around a common frequency encoding axis (Shilling et al., 2009), or arbitrary slice orientations (Poot et al., 2010b). SRR has indeed been shown to provide a better trade-off between acquisition time, spatial resolution, and signal-to-noise ratio (SNR) than conventional direct HR acquisition (Plenge et al., 2012). Meanwhile, SRR has also been successfully applied to different qMRI modalities, including diffusion MRI (Poot et al., 2013; Fogtmann et al., 2014; Van Steenkiste et al., 2016), relaxometry (Van Steenkiste et al., 2017; Bano et al., 2020; Lajous et al., 2020) and arterial spin labeling (Bladt et al., 2020). In some of these approaches, HR images are individually reconstructed from a set of equally contrast-weighted LR images, prior to voxel-wise fitting a parametric qMRI signal model (e.g., a diffusion model or relaxation model) to these reconstructed HR images (Poot et al., 2013; Lajous et al., 2020), whereas in other approaches the qMRI signal model is included in the reconstruction and HR parameter maps are estimated directly from the LR images, without first reconstructing the individual HR contrast weighted images (Fogtmann et al., 2014; Van Steenkiste et al., 2016, 2017; Bano et al., 2020).

In addition to the challenge of 3D isotropic HR parameter mapping, the long qMRI scan times come with an increased risk of patient motion. If this motion is not properly accounted for, the spatial resolution of the obtained parameter maps will be negatively affected. Like conventional qMRI methods, SRR methods for qMRI usually correct for motion by performing image registration as a pre-processing step, prior to the estimation of the HR parameter maps (Van Steenkiste et al., 2016, 2017), where the latter step is often preceded by an

intermediate step of HR image reconstruction (Scherrer et al., 2012; Poot et al., 2013). A downside to this multi-step approach is the lack of a feedback mechanism that connects the motion compensation routine with the final estimation of the HR parameter maps. As a result, registration errors may propagate into the parameter estimation step, introducing a bias (Nachmani et al., 2019).

To avoid error propagation, image registration can be integrated in a joint motion/qMRI parameter estimation framework. This strategy has already been successfully applied to correct inter-scan motion in T_1 mapping (Ramos-Llordén et al., 2017) or to correct for motion in multi-shell diffusion MRI (Christiaens et al., 2021). At the same time, methods have been proposed that combine SRR with joint motion estimation for anatomical (qualitative) MRI (Rousseau et al., 2010; Jiang et al., 2007; Gholipour et al., 2010; Fogtmann et al., 2012; Kainz et al., 2015; Ebner et al., 2020). However, until now, the development of a unified motion estimation/SRR approach for qMRI has received little attention (Fogtmann et al., 2014; Beirinckx et al., 2020).

In the present contribution, we propose a multi-frame model-based SRR method for multiparametric quantitative MRI with integrated inter-image motion estimation in a Bayesian Maximum a Posteriori (MAP) estimation framework. As a guiding application, we focus on MR relaxometry, but the method's modular construction ensures an easy adaption to other qMRI modalities. The novelty of the method lies in its unique combination of properties that makes it stand out from existing SRR methods in qMRI. First, by combining superresolution image reconstruction and quantitative parameter estimation in a single integrated model-based approach, 3D HR biophysical parameter maps are estimated directly from a set of multi-slice differently contrast-weighted LR images, which distinguishes our method from two-step qMRI SRR approaches that reconstruct individual HR images from equally contrast-weighted LR images prior to voxel-wise fitting a qMRI signal model (e.g. a relaxation model or diffusion model) to these reconstructed images (Poot et al., 2013; Lajous et al., 2020). Second, the joint estimation of the motion and the biophysical parameters of interest allows our method to outperform state-of-the-art qMRI SRR algorithms that either do not correct for motion (Bano et al., 2020), or work with decoupled motion estimation algorithms (Van Steenkiste et al., 2017). Third, unlike state-of-the-art SRR methods in qMRI that rely on orthogonal slice orientations and use the same set of contrast weightings for each slice orientation (e.g., Fogtmann et al., 2014), our method allows for arbitrary slice orientations and a different contrast weighting for each LR image, offering a much-increased imaging flexibility. Finally, its Bayesian estimation approach allows our method to optimally exploit prior knowledge of tissue properties and noise statistics, as opposed to standard regularized least-squares methods (Poot et al., 2010b; Van Steenkiste et al., 2017; Bano et al., 2020; Lajous et al., 2020).

To demonstrate its potential, the proposed unified quantitative SRR method is evaluated for T_1 mapping and T_2 mapping. Its performance in terms of accuracy, precision and mean squared error is extensively validated using synthetic whole brain simulations. Finally, the applicability of the SRR method is demonstrated on *in-vivo* brain data and its performance on brain structure delineation (spatial resolution) is evaluated.

6.2 Theory

This section introduces the forward model of the SRR problem considered in this contribution. It describes the relation between the LR images and the HR parameter maps to be reconstructed and accounts for unintended motion. Furthermore, the Bayesian Maximum a Posteriori (MAP) estimator is introduced that is used to estimate the HR maps jointly with the motion parameters, accounting for the MR data distribution and using a total variation (TV) prior for the HR maps and a non-informative prior for the motion parameters.

Remarks on notation - In the following paragraphs, we slightly deviate from the more conventional notation often used in Bayesian statistics. Typically, random variables are denoted by uppercase letters (e.g., X), while their realizations are represented by lowercase letters (e.g., x). However, in this work, we have adopted different conventions, which should be clear from the context and accompanying explanations. We hope this slight deviation does not cause any confusion and appreciate your understanding.

6.2.1 Forward model

Let $s = \{s_n\}_{n=1}^N$ be the set of N vectorized noiseless anisotropic LR multi-slice contrastweighted magnitude images, where $s_n = \{s_{nl}\}_{l=1}^{N_s} \in \mathbb{R}^{N_s \times 1}$ is sampled at the LR grid points $y_n = \{y_{nl}\}_{l=1}^{N_s} \in \mathbb{R}^{3 \times N_s}$ with N_s the number of voxels per LR image. Then, each s_n can be modelled as:

$$\boldsymbol{s}_n = |\boldsymbol{D}\boldsymbol{B}\boldsymbol{G}_n\boldsymbol{M}_{\boldsymbol{\theta}_n}\boldsymbol{r}_n|\,, \qquad (6.2.1)$$

where $|\cdot|$ denotes the pointwise modulus operator and $\mathbf{r}_n = \{r_{nj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ represents the virtual, noise-free HR image assumed to be acquired with the same contrast-weighting settings as \mathbf{s}_n and defined at the targeted isotropic HR grid points $\mathbf{x} = \{\mathbf{x}_j\}_{j=1}^{N_r} \in \mathbb{R}^{3 \times N_r}$, with N_r the number of voxels of the HR image. Furthermore, $\mathbf{M}_{\theta_n} \in \mathbb{R}^{N_r \times N_r}$, $\mathbf{G}_n \in \mathbb{R}^{N_r \times N_r}$, $\mathbf{B} \in \mathbb{R}^{N_r \times N_r}$, and $\mathbf{D} \in \mathbb{R}^{N_s \times N_r}$ are linear operators that describe motion, a known geometric transformation that maps the grid coordinates of the HR image \mathbf{r}_n to those of the LR image \mathbf{s}_n , spatially invariant blurring, and down-sampling, respectively. The motion operator \mathbf{M}_{θ_n} is modeled as a parametric function of θ_n . Assuming rigid inter-image motion, the parameter vector $\theta_n \in \mathbb{R}^{6 \times 1}$ is given by

$$\boldsymbol{\theta}_{n} = \left[t_{xn}, t_{yn}, t_{zn}, \alpha_{n}, \beta_{n}, \gamma_{n} \right]^{T}, \qquad (6.2.2)$$

with t_{xn} , t_{yn} , t_{zn} the translation parameters and α_n , β_n , γ_n the Euler angles of three elementary rotation matrices that describe rotation around the x, y and z axis, respectively. The operator G_n models the SRR acquisition scheme. In the SRR acquisition scheme considered in this work, the LR images each have a different slice orientation, where the different orientations are obtained by rotating around a fixed encoding axis. Detailed descriptions of the warping operator M_{θ_n} , which is analytically differentiable w.r.t. θ_n , as well as the operators G_n , B, and D are included as part of the supplementary material.

In qMRI, one is not so much interested in the voxel intensities of the HR images r_n , but rather in the values of the underlying biophysical tissue parameters in those voxels, such as the proton densities and T_1 and T_2 relaxation times. Let $\vartheta = \{\vartheta_q\}_{q=1}^Q \in \mathbb{R}^{N_r \times Q}$ be the biophysical parameter maps to be inferred, with $\vartheta_q = \{\vartheta_{qj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ the q^{th} tissue parameter map and $\vartheta_{\bullet j} \in \mathbb{R}^{Q \times 1}$ all tissue parameters of the j^{th} voxel of ϑ_q . Then, the j^{th}

voxel of the HR image r_n , can be modelled as

$$r_{nj} = f_n(\boldsymbol{\vartheta}_{\bullet j}) \tag{6.2.3}$$

with $f_n(\vartheta_{\bullet j}) : \mathbb{R}^{Q \times 1} \mapsto \mathbb{R}$ a relaxometry, diffusion, or perfusion model, or any other qMRI model that describes the relation between r_{nj} and the underlying biophysical tissue parameters $\vartheta_{\bullet j}$ in the corresponding voxel. In the current contribution, the proposed SRR method is evaluated for T_1 and T_2 relaxometry, using the signal models described in section 6.3. The forward model of the SRR reconstruction problem considered in this contribution is obtained by substituting Eq. (6.2.3) in Eq. (6.2.1). The parameters to be estimated are the parameter maps ϑ and motion parameters $\theta = \{\theta_n\}_{n=1}^N$.

6.2.2 Joint Bayesian estimation framework

6.2.2.1 Bayes theorem

Let $\tilde{s} = {\tilde{s}_n}_{n=1}^N \in \mathbb{R}^{N_s \times N}$ denote the set of N measured LR multi-slice images with $\tilde{s}_n = {\tilde{s}_{nl}} \in \mathbb{R}^{N_s \times 1}$. Following a Bayesian approach, both the data \tilde{s} and the parameters $\{\vartheta, \theta\}$ to be estimated are modeled as random variables, where Bayes' theorem gives an expression for the *posterior* distribution of the parameters given the data:

$$p(\boldsymbol{\vartheta}, \boldsymbol{\theta} | \tilde{\boldsymbol{s}}) = \frac{p(\tilde{\boldsymbol{s}} | \boldsymbol{\vartheta}, \boldsymbol{\theta}) p(\boldsymbol{\vartheta}) p(\boldsymbol{\theta})}{p(\tilde{\boldsymbol{s}})} \quad , \tag{6.2.4}$$

with $p(\tilde{s}|\vartheta, \theta)$ the *likelihood* function of the data, $p(\vartheta)$ and $p(\theta)$ the *prior* distributions that encapsulate the prior knowledge about ϑ and θ , respectively, and $p(\tilde{s})$ a scaling factor that can be ignored since it does not affect the estimator that will be described below.

6.2.2.2 Maximum a posteriori estimator

The MAP estimator maximizes $p(\vartheta, \theta | \tilde{s})$ w.r.t. the parameters $\{\vartheta, \theta\}$:

$$\{\hat{\boldsymbol{\vartheta}}, \hat{\boldsymbol{\theta}}\} = \arg\max_{\boldsymbol{\vartheta}, \boldsymbol{\theta}} p(\boldsymbol{\vartheta}, \boldsymbol{\theta} | \tilde{\boldsymbol{s}}).$$
(6.2.5)

Eq. (6.2.5) is typically solved by minimizing the negative logarithm of $p(\vartheta, \theta | \tilde{s})$.

6.2.2.3 Likelihood function

Without loss of generalization, the measured LR images \tilde{s} are assumed to be Rician distributed, which is a valid noise model for magnitude images reconstructed from single-coil *k*-space data (den Dekker & Sijbers, 2014), for images reconstructed from multi-coil data with SENSE (Aja-Fernández et al., 2014), or with GRAPPA jointly with a spatial-matched-filter or the Adaptive Combine method (Walsh et al., 2000). Then, the probability density function (PDF) of \tilde{s}_{nl} is given by:

$$p(\tilde{s}_{nl}|\boldsymbol{\vartheta},\boldsymbol{\theta}_n) = \frac{\tilde{s}_{nl}}{\sigma_{nl}^2} e^{-\frac{\tilde{s}_{nl}^2 + s_{nl}^2(\boldsymbol{\vartheta},\boldsymbol{\theta}_n)}{2\sigma_{nl}^2}} I_0\left(\frac{\tilde{s}_{nl}s_{nl}(\boldsymbol{\vartheta},\boldsymbol{\theta}_n)}{\sigma_{nl}^2}\right) u(\tilde{s}_{nl}),$$
(6.2.6)

with $I_0(\cdot)$ the zeroth order modified Bessel function of the first kind, and σ_{nl} the nonstationary (i.e. spatially-dependent) standard deviation of the Gaussian noise disturbing the complex data underlying the magnitude MR data. The unit step function $u(\cdot)$ is used to indicate that (6.2.6) is non-zero for non-negative values of \tilde{s}_{nl} only. Assuming all voxels to be statistically independent, the joint PDF of \tilde{s} is given by

$$p(\tilde{\boldsymbol{s}}|\boldsymbol{\vartheta},\boldsymbol{\theta}) = \prod_{n=1}^{N} \prod_{l=1}^{N_s} p(\tilde{\boldsymbol{s}}_{nl}|\boldsymbol{\vartheta},\boldsymbol{\theta}_n).$$
(6.2.7)

When (6.2.7) is viewed as a function of the unknown parameters $\{\vartheta, \theta\}$ given the data \tilde{s} , it is called the *likelihood* function. It follows from (6.2.6) and (6.2.7) that the negative log-likelihood function $\mathcal{L}_{\tilde{s}} \equiv -\log p(\tilde{s}|\vartheta, \theta)$ can be written as (Sijbers et al., 1998)

$$\mathcal{L}_{\tilde{s}}(\vartheta, \theta | \tilde{s}) = \sum_{n=1}^{N} \sum_{l=1}^{N_{s}} \left[-\log \tilde{s}_{nl} + \log \sigma_{nl}^{2} + \frac{\tilde{s}_{nl}^{2}}{2\sigma_{nl}^{2}} + \frac{s_{nl}^{2}(\vartheta, \theta_{n})}{2\sigma_{nl}^{2}} - \log I_{0} \left(\frac{\tilde{s}_{nl} s_{nl}(\vartheta, \theta_{n})}{\sigma_{nl}^{2}} \right) \right].$$
(6.2.8)

Furthermore, it is assumed that the noise standard deviations can be estimated prior to the construction of the MAP estimator of $\{\vartheta, \theta\}$ using tailored noise estimation routines (Aja-Fernández et al., 2015; Pieciak et al., 2017; Maitra & Faden, 2009; Bouhrara et al., 2017).

6.2.2.4 Prior distributions

For each of the Q HR tissue parameter maps associated with ϑ , a discretized upwind TV prior (Chambolle et al., 2011) is chosen. TV is renowned for preserving edges and reducing noise by penalizing large intensity variations, promoting smooth regions while maintaining essential structures. In MRI, preserving sharp edges is crucial for accurate diagnosis and interpretation, and denoising is vital in SRR to enhance resolution without amplifying noise. The upwind scheme, considering gradient direction, further enhances edge preservation by accurately capturing discontinuities and avoiding artificial smoothing along edges:

$$p(\boldsymbol{\vartheta}_q) \propto \exp\{-\frac{2}{\lambda_q} \mathsf{TV}(\boldsymbol{\vartheta}_q)\}, \text{ with } q = 1, \dots, Q,$$
 (6.2.9)

where $\lambda_q > 0$ denotes the hyperparameter to be selected by the user, as will be discussed in section 6.2.2.6, and with

$$\mathsf{TV}(\boldsymbol{\vartheta}_{q}) = \sum_{j} \left[\sqrt{ \boldsymbol{\varepsilon}^{2} + (\Delta^{x,+}(\boldsymbol{\vartheta}_{qj}))^{2} + (\Delta^{x,-}(\boldsymbol{\vartheta}_{qj}))^{2} + (\Delta^{y,-}(\boldsymbol{\vartheta}_{qj}))^{2} - \boldsymbol{\varepsilon} } \right], \qquad (6.2.10)$$
$$+ (\Delta^{z,+}(\boldsymbol{\vartheta}_{qj}))^{2} + (\Delta^{z,-}(\boldsymbol{\vartheta}_{qj})^{2} \right],$$

where $\Delta^{x,+}(\vartheta_{qj})$, $\Delta^{x,-}(\vartheta_{qj})$, $\Delta^{y,+}(\vartheta_{qj})$, $\Delta^{y,-}(\vartheta_{qj})$, $\Delta^{z,+}(\vartheta_{qj})$, and $\Delta^{z,-}(\vartheta_{qj})$ represent the forward (+) and backward (-) first order differences, in the *x*-, *y*-, and *z*-direction, at the *j*th HR voxel of the parameter map ϑ_q . Furthermore, a small value $\epsilon > 0$ is introduced, to avoid derivative singularities of TV when ϑ_q is locally constant.

For the motion parameters θ , a non-informative prior $p(\theta)$ is adopted, assuming $p(\theta)$ to be uniform over the range of values for which the likelihood function is non-negligible.

6.2.2.5 Alternating minimization

The nonlinear optimization problem (6.2.5) is solved using the *alternating minimization* method, also known as the cyclic block-coordinate descent (cBCD) method (Fessler & Kim, 2011; Beck & Tetruashvili, 2013). In this method, the parameters $\{\vartheta, \theta\}$ are split into two blocks that contain the motion parameters θ and the tissue parameters ϑ , respectively, and the cost function is successively minimized with respect to each block in a cyclic order:

$$\hat{\theta}^{(t+1)} = \arg\min_{\theta} \mathcal{L}_{\tilde{s}}(\hat{\vartheta}^{(t)}, \theta | \tilde{s})$$
(P.1)

$$\hat{\boldsymbol{\vartheta}}^{(t+1)} = \arg\min_{\boldsymbol{\vartheta}} \left[\mathcal{L}_{\tilde{\boldsymbol{s}}}(\boldsymbol{\vartheta}, \hat{\boldsymbol{\theta}}^{(t+1)} | \tilde{\boldsymbol{s}}) + \sum_{q=1}^{Q} \frac{2}{\lambda_{q}} \mathsf{TV}(\boldsymbol{\vartheta}_{q}) \right]$$
(P.2)

with $\hat{\vartheta}^{(0)} = \vartheta_{\text{ini}}$, and $\hat{\theta}^{(0)} = \theta_{\text{ini}}$ the initial values of the HR tissue parameters ϑ and the motion parameters θ , respectively. The procedure is terminated when a maximum number of iterations, t_{max} , is exceeded, or when a convergence tolerance on the relative difference of the tissue parameter estimates between consecutive iterations is reached. The pseudo-code of our proposed MAP estimation framework is presented in Algorithm 1. The initial values ϑ_{ini} , and θ_{ini} are obtained using a *while*-loop routine consisting of three main steps. First, a HR magnitude contrast-weighted image is approximated from each LR contrast-weighted image by applying the adjoint operator $M_{\theta_n}^T G_n^T B^T D^T$ of the SRR forward model (6.2.1) to each LR image, followed by the application of the pointwise modulus operator $|\cdot|$, to regain magnitude images. Second, initial tissue parameter values ϑ_{ini} are obtained by voxel-wise nonlinear least-squares (NLLS) fitting the modulus of the signal model (6.3.2) to these upsampled LR images with a Levenberg-Marquardt algorithm. In a third step, problem (P.1) is solved to obtain initial estimates for the motion parameters θ_{ini} , where the tissue parameter estimates of the second step are kept fixed in the cost function $\mathcal{L}_{\tilde{s}}(\hat{\vartheta}^{(t+1)}, \hat{\theta}|\tilde{s})$.

The *inter-image* motion estimation problem (P.1) adopts a particularly simple structure when the signal model parameters remain fixed. If no dependence of $\{\theta_n\}_{n=1}^N$ through index *n* is assumed, the minimization can be decoupled into *N* optimization problems, which can be implemented very efficiently by parallel operations. Each of these decoupled problems is minimized using a trust-region Newton algorithm (Coleman & Li, 1994), with analytical expressions for the Jacobian to speed up convergence. The derivation of these expressions is included as part of the supplementary material provided with this work. To solve the large-scale optimization problem (P.2), a trust-region-reflective Newton algorithm is used (Coleman & Li, 1994), with analytical expressions for the Jacobian and Hessian, which have also been included as part of the supplementary material.

6.2.2.6 Regularization parameter selection

The hyperparameters $\lambda_1, \ldots, \lambda_Q$ of the prior distribution (6.2.9) act as regularization parameters that balance data consistency (as quantified by the likelihood function) against the requirement that the parameter maps be smooth (as imposed by the TV prior). The selection of the regularization parameters of a nonlinear optimization problem like the one at hand is a challenging task for which no standard procedure exists. Poorly chosen regularization parameters may lead to either over-smoothing or under-smoothing. In this work, the individual regularization parameters were determined such that the corresponding TV terms contribute equally to the cost function of (P.2). To this end, each TV term of (P.2)

Algorithm 1: Model-based SRR with joint motion estimation (SRR-joint)

Input: LR images \tilde{s} and initial values ϑ_{ini} and θ_{ini} Output: MAP estimates $\hat{\vartheta}_{MAP}$ and $\hat{\theta}_{MAP}$ Set $t \leftarrow 0$ and $\hat{\vartheta}^{(0)}$, $\hat{\theta}^{(0)} \leftarrow \vartheta_{ini}$, θ_{ini} ; $\mathcal{E}^{(0)} = r\mathcal{E}_{min}$, with $r \in \mathbb{R}_{>1}$; while $\mathcal{E}^{(t)} \ge \mathcal{E}_{min}$ and $t < t_{max}$ do \triangleright Solve (P.1) to get $\hat{\theta}^{(t+1)}$: $\hat{\theta}^{(t+1)} = \arg\min_{\vartheta} \mathcal{L}_{\tilde{s}}(\hat{\vartheta}^{(t)}, \theta|\tilde{s})$, started from $\theta \leftarrow \hat{\theta}^{(t)}$; \triangleright Solve (P.2) to get $\hat{\vartheta}^{(t+1)}$: $\hat{\vartheta}^{(t+1)} = \arg\min_{\vartheta} \left[\mathcal{L}_{\tilde{s}}(\vartheta, \hat{\theta}^{(t+1)}|\tilde{s}) + \sum_{q=1}^{Q} \frac{2}{\lambda_{q}} \mathsf{TV}(\vartheta_{q}) \right]$, started from $\vartheta \leftarrow \hat{\vartheta}^{(t)}$; \triangleright Calculate^a $\mathcal{E}^{(t+1)} = \|\hat{\vartheta}^{(t+1)} - \hat{\vartheta}^{(t)}\|_{2} / \|\hat{\vartheta}^{(t+1)}\|_{2}$; \triangleright Set $t \leftarrow t + 1$; end $\hat{\vartheta}_{MAP} = \hat{\vartheta}^{(t)}$ and $\hat{\theta}_{MAP} = \hat{\theta}^{(t)}$; return $\hat{\vartheta}_{MAP}, \hat{\theta}_{MAP}$;

^aVectorization of $\hat{artheta}^{(t+1)}$ and $\hat{artheta}^{(t)}$ is performed before taking the norm.

is evaluated in the initial estimates of the respective tissue parameter map, and the ratio of each TV value to the TV value of the first tissue parameter is used to determine the regularization parameters $\lambda_2, \ldots, \lambda_Q$ as a function of λ_1 . More specifically, this leads to $\lambda_q = [TV(\hat{\vartheta}_{q,\text{ini}})/TV(\hat{\vartheta}_{1,\text{ini}})] \cdot \lambda_1$. As such, the multi-parameter regularization selection problem is cast into a single-parameter regularization problem. The remaining regularization parameter, λ_1 , is chosen empirically from repeated reconstructions for progressively increasing values of λ_1 , as a compromise between noise removal and image resolution. Special care is taken to ensure that no in-plane resolution loss occurs due to over-smoothing, thereby maintaining the integrity of the 2D multi-slice data.

6.3 Materials and Methods

The proposed model-based Bayesian SRR method with joint motion estimation was validated in whole brain simulations choosing T_1 relaxometry as a showcase example. Next, to demonstrate the ability of the proposed method to improve the quality of reconstructed parameter maps, a proof-of-concept evaluation was performed for a T_1 and T_2 quantitative mapping protocol using two contrast-weighted *in vivo* brain datasets.

The following parametric signal models were adopted in this contribution:

• T_1 relaxation signal model of the gold standard inversion recovery (IR) sequence (Barral et al., 2010):

$$f_n(\boldsymbol{\vartheta}_{\bullet j}) = \rho_j \left(1 - (1 - \cos \alpha) e^{-\frac{\mathrm{TR}}{T_{1,j}}} + e^{-\frac{\mathrm{TR}}{T_{1,j}}} \right), \qquad (6.3.1)$$

with TI_n the n^{th} inversion time, α the inversion pulse angle, TR the repetition time and $\vartheta_{\bullet j} = [\rho_j, \mathcal{T}_{1,j}]^{\mathcal{T}}$ the tissue parameter vector at position x_j , in which ρ_j is a parameter

proportional to the proton density and receiver gain and $T_{1,j}$ is the longitudinal relaxation time. Assuming $\alpha = 180^{\circ}$ and TR $\gg T_1$, Eq. (6.3.1) simplifies to

$$f_n(\boldsymbol{\vartheta}_{\bullet j}) = \rho_j \left(1 - 2 \, e^{-\frac{\mathrm{T} I_n}{T_{1,j}}} \right). \tag{6.3.2}$$

 T₂ relaxation signal model of a conventional Multi-Echo Spin Echo (MESE) sequence (Carr & Purcell, 1954):

$$f_n(\boldsymbol{\vartheta}_{\bullet j}) = \rho_j e^{-\frac{\mathsf{T} E_n}{\mathcal{T}_{2,j}}}, \qquad (6.3.3)$$

with TE_n the n^{th} echo time, and $\vartheta_{\bullet j} = [\rho_j, T_{2,j}]^T$ the tissue parameter vector at position \boldsymbol{x}_j , in which ρ_j is again a parameter proportional to the proton density and receiver gain and $T_{2,j}$ is the transverse relaxation time. Note that we have assumed a perfect 90° excitation pulse to tilt the magnetization vector in the transverse plane, and perfect 180° refocusing pulses to recover multiple spin echoes corresponding with T_2 estimates along the signal envelope.

The proposed SRR method was compared with an SRR approach without motion estimation, and one in which SRR is preceded by a motion compensation step. To sum up, the following three frameworks were compared against each other:

- 1. **SRR-static**: a model-based SRR framework without motion estimation. This approach consists of three steps. First, a HR magnitude image is approximated from each LR image using the adjoint operator $G_n^T B^T D^T$ of the SRR forward model (6.2.1), followed by application of the pointwise modulus operator $|\cdot|$. Second, voxel-wise NLLS fitting of the modulus of the signal model is performed using a Levenberg-Marquardt algorithm to obtain initial parameter map estimates. Finally, problem (P.2) is solved assuming $\hat{\theta} = 0$.
- 2. **SRR-reg**: a model-based SRR framework in which the inter-image motion parameters are estimated prior to the SRR by means of an advanced registration routine (Van Steenkiste et al., 2017). In this approach, a registration routine is performed consisting of four steps, where the first two steps correspond with the first two steps of SRR-static. In a third step, LR images are simulated using the estimated HR parameter maps from the previous step and the forward model (6.2.1). As a fourth step, rigid motion parameter estimates θ_{REG} are obtained from pairwise rigid registration using a mean squared error metric and a regular step gradient descent optimization algorithm (MaxIter = 800, GradientMagnitudeTolerance = 10^{-12}). In order to obtain rigid motion parameters that can be used as input parameters of the motion operator M_{θ_n} , which is part of the forward model in problem (P.2), registration needs to be performed on the HR grid. As such, the simulated and acquired LR image datasets to be co-registered are transformed to the HR grid using the adjoint operator $G_n^T B^T D^T$ of the SRR forward model (6.2.1). Next, steps 1-4 are repeated until a convergence tolerance $\mathcal{E}_{min}=10^{-4}$ on the relative difference of the tissue parameter estimates between consecutive iterations is met. Finally, problem (P.2) is solved in which motion parameter estimates θ_{REG} obtained from the registration routine remain fixed.
- 3. **SRR-joint**: the proposed SRR framework with joint motion estimation, as described in section 6.2.2. The pseudo code of this framework is described in Algorithm 1. The maximum number of iterations and the tolerance criterion to halt the algorithm were chosen to be $t_{\text{max}} = 80$ and $\mathcal{E}_{\text{min}} = 10^{-4}$, respectively.

For the *in vivo* experiments described in section 6.3.2, the regularization parameters of SRRjoint were chosen following the procedure described in section 6.2.2.6, yielding $2/\lambda_1 = 1.1 \times 10^{-2}$ and $2/\lambda_2 = 5.6 \times 10^{-3}$ for the *in vivo* T_1 mapping experiment, and $2/\lambda_1 = 2.4 \times 10^{-2}$ and $2/\lambda_2 = 1.0 \times 10^{-2}$ for the *in vivo* T_2 mapping experiment, respectively. Next, the same regularization parameters were used for SRR-static and SRR-reg, to guarantee a fair comparison. Finally, the same regularization weights as for the *in vivo* T_1 mapping experiment were also used for the whole brain Monte Carlo simulation experiments described in the next section.

6.3.1 Whole Brain Simulations

Ground truth T_1 and ρ parameter maps for a synthetic whole brain Monte Carlo simulation experiment were generated from parameter maps obtained after model-based SRR on the T_1 -weighted *in vivo* dataset, further described in section 6.3.2. Both HR parameter maps were of size $160 \times 160 \times 160$, with an isotropic voxel size of 1.6 mm.

From these ground truth parameter maps, $N_{MC} = 8$ Rician distributed realisations of a LR T_1 -weighted dataset were simulated. Each dataset consisted of N = 14 images with log(Tl_n) equidistant between log(100 ms) and log(3000 ms) (Van Steenkiste et al., 2017). The LR images were synthesized using the forward model (6.2.1), with an image size of $160 \times 160 \times 40$ and with an anisotropic voxel size of $1.6 \times 1.6 \times 6.4$ mm³. The inter-image motion parameters $\{\theta_n\}_{n=1}^N$ were chosen equal to an estimated set of motion parameters obtained from model-based SRR with SRR-joint on the T_1 -weighted *in vivo* dataset to guarantee realistic head movement. The extreme and mean values for each of the motion parameters are reported in Table 6.3.1, where the mean value of the k^{th} motion component, i.e. $\overline{\theta}_k$, was calculated as

$$\overline{\theta}_k = \frac{1}{N} \sum_{n=1}^N \theta_{nk}.$$
(6.3.4)

Table 6.3.1: Extreme and mean values for each of the motion parameters that were used in the synthetic whole brain simulation experiments.

	t _x	t _y	t _z	α	β	γ
	[mm]	[mm]	[mm]	[degree]	[degree]	[degree]
extremum $\overline{\theta}_k$	0.517	2.486	2.082	2.890	-0.538	-0.836
	0.28	1.69	0.38	0.67	-0.03	-0.54

Similar to the acquisition protocol used in (Van Steenkiste et al., 2016, 2017) and the acquisition protocol in the *in vivo* experiments (cfr. section 6.3.2), the LR images were simulated with different slice orientations, where the rotation was performed around the phase encoding axis in increments of $180/N_0$ degrees, with $N_0 = 7$ the number of slice orientations. Since rotation in image space corresponds to rotation in frequency domain, this acquisition scheme ensures that each LR image covers a different part of the *k*-space (as shown in the top row of Fig. 6.3.1) (Plenge et al., 2012). Two LR images were simulated for each slice orientation, where each of the thus resulting N = 14 images had a unique inversion time. An overview of the slice orientations and inversion times of the T_1 -weighted LR images, along with their *k*-space coverage, is given in Fig. 6.3.1.

CHAPTER 6

Finally, the LR T_1 -weighted images were corrupted with spatially variant Rician noise, where the spatially variant noise pattern corresponded with an isotropic Gaussian function to model the gradual deterioration of the head coil detection towards the center of the brain (Pieciak et al., 2017). The level of the noise map was adjusted to match that of the *in vivo* T_1 -weighted dataset that will be described in the next section. To this end, the overall SNR, defined as the ratio of the spatial mean of the signal to the standard deviation of the noise, where the latter is estimated using the method of Coupé *et al.* (Coupé et al., 2010), was calculated in a small homogeneous region of the corpus callosum of the *in vivo* T_1 -weighted image acquired with TI = 100 ms, and observed to be 16. Next, the level of the noise map of the simulated T_1 -weighted images was adjusted to match this SNR value in the corpus callosum of the simulated image sampled at the same inversion time. In all simulation SRR experiments, the noise standard deviation maps were assumed to be known.



Figure 6.3.1: Overview of the different slice orientations for the LR T_1 -weighted *in vivo* (and simulated) dataset(s), together with a schematic representation of the overlap in *k*-space when images are combined. Fourteen 2D IR TSE T_1 -weighted LR images were acquired with large slice thickness and a high in-plane resolution. The slice orientation was consecutively altered by rotation over a specified angle (0°, 25.7°, 51.4°, ..., 154.2°) around the phase-encoding direction. As indicated, each T_1 -weighted LR image was acquired with a unique inversion time.

6.3.2 In Vivo Data

The proposed SRR method was validated using two *in vivo* human brain datasets suffering from involuntary patient motion. Both healthy volunteers (adult, male, 28 and 32 years old) were scanned after written informed consent and approval by the institutional ethics committee using a 3T MRI scanner (Magnetom PrismaFit, Siemens Healthcare, Erlangen, Germany) with VE11B software, a maximum gradient amplitude of 80 mT/m, a maximum slew rate of 200 T/m/s, and a dedicated head-coil with 32 receiver channels. Magnitude

data was reconstructed from the complex coil images using the adaptive combine algorithm (Walsh et al., 2000).

The first *in vivo* LR dataset consisted of a series of T_1 -weighted LR images with anisotropic voxel size. In total, 14 repetitions of an interleaved multi-slice IR TSE with low through-plane resolution (voxel size, $1.0 \times 1.0 \times 4.0 \text{ mm}^3$), with turbo factor 10, without slice gap, and with 100% sampling, were acquired. The slice thickness of the LR dataset was chosen to have whole brain coverage without exceeding SAR limits. The acquisition matrix was equal to 256 × 256, with a total number of slices equal to 40. Furthermore, the bandwidth was fixed at 305 Hz/pixel, and the TR and echo time (TE) were equal to 5000 ms and 8.8 ms, respectively. No in-plane acceleration was used. Each acquisition was characterized by a specific rotation around the phase-encoding axis and a unique inversion time, where the rotation angles and inversion times agree with those used in the simulation study, as summarized in Fig. 6.3.1. The scan time per anisotropic 2D slice stack was 2 minutes and 3 seconds, resulting in a total scan time of 28 minutes and 44 seconds. SRR was performed at an isotropic HR grid with a voxel size of $1.0 \times 1.0 \times 1.0 \text{ mm}^3$. Spatially variant noise standard deviation maps were estimated using the method of (Aja-Fernández et al., 2015).



Figure 6.3.2: Schematic representation of a multi-echo spin echo (**MESE**) sequence. Compared to a standard spin echo sequence, the MESE sequence stimulates the spin system with additional 180° pulses. As long as T_2 -relaxation is not complete and MR signal is present, this allows to generate extra echoes, i.e. additional T_2 -weighted images, within a given repetition time. The amplitude of each echo is progressively smaller due to the T_2 decay. Also, the echo time (TE) spacing, i.e. the time between consecutive echoes, is inherently fixed.

A second *in vivo* T_2 -weighted anisotropic LR dataset was acquired, using 7 repetitions of an interleaved multislice MESE acquisition with low through-plane resolution (voxel size, $1.75 \times 1.75 \times 7.0 \text{ mm}^3$), without slice gap, using a 3-fold in-plane GRAPPA acceleration factor with 24 reference lines. A schematic representation of the MESE sequence is given in Fig. 6.3.2. The acquisition matrix was equal to 128×128 , with a total number of slices equal to 26. The bandwidth was fixed at 227 Hz/pixel. Each MESE acquisition was characterized by a unique rotation around the phase-encoding axis (rotations similar as in Fig. 6.3.1), and consisted of 4 unique echo times (Fig. 6.3.3). An overview of the sampled TEs per MESE acquisition is given in Table 6.3.2. The echo time spacing Δ TE in each MESE was chosen as such to ensure full coverage of the T2 relaxation curve when all 7 acquisitions are combined. This echo time selection is illustrated in more detail in Fig. 6.3.4. The TR = 4320 ms was kept constant for each MESE acquisition to avoid differences in T_1 -weighting. In addition, the first echo of each MESE acquisition was ignored in the SRR reconstruction, which is a common consideration for MESE acquisitions (Petrovic et al., 2015), to avoid protruding errors from imperfect refocusing and stimulated (secondary) echoes that disrupt the T_2 decay of the primary SEs. In this way, the total count of sampled TEs was limited to 21. The scan time per anisotropic 2D MESE acquisition was 4 minutes 11 seconds, resulting in a total scan time for this proof-of-concept protocol of 29 minutes 17 seconds. SRR was performed at an isotropic HR grid with a voxel size of $1.75 \times 1.75 \times 1.75 \text{ mm}^3$, and non-stationary noise standard deviation maps were again estimated using the method of (Aja-Fernández et al., 2015).



Figure 6.3.3: Schematic representation of a **multi-MESE** acquisition strategy for which multiple MESE acquisitions (see Fig. 6.3.2) are combined. Each individual MESE number is characterized with a different rotated acquisition geometry. Furthermore, each MESE sequence uses the same repetition time (TR) to avoid T1-weighting effects.

Table 6.3.2: Distribution of echo times per MESE acquisition for the *in vivo* T_2 mapping experiment. Slice orientation angles corresponds with those given in Fig. 6.3.1. Gray coloured cells indicate the first echo times that were ignored to alleviate the effect of stimulated secondary echoes.

	Slice orientation angle [°]	TE1	TE_2	TE ₃	TE4
MESE 1	0	10.0	20.0	30.0	40.0
MESE 2	25.7	11.8	23.6	35.4	47.2
MESE 3	51.4	19.2	38.4	57.6	76.8
MESE 4	77.1	22.6	45.2	67.8	90.4
MESE 5	102.8	34.0	68.0	102.0	136.0
MESE 6	128.5	36.9	73.8	110.7	147.6
MESE 7	154.2	40.0	80.0	120.0	160.0



Figure 6.3.4: Echo time selection for the *in vivo* T_2 mapping experiment: Seven different MESE acquisitions were used to acquire a total of 28 T_2 -weighted LR images. Each MESE acquisition was characterized by a unique rotation around the phase-encoding axis (rotation angles similar as in Fig. 6.3.1). Furthermore, each MESE consisted of 4 unique echo times, which are tabulated (bottom right), and visualized with the corresponding LR image number (top). The first echo time of each MESE was ignored in the model-based SRR to alleviate the effect of stimulated secondary echoes. This effect is clearly distinguishable when plotting the mean signal intensity for each LR T_2 -weighted image as a function of the echo time (bottom left).
6.3.3 Quantitative Image Analysis

The results of the synthetic whole brain Monte Carlo simulation experiment were assessed quantitatively using the following performance measures (Ramos-Llordén et al., 2017; Beirinckx et al., 2020):

- (a) Relative bias. The bias quantifies the accuracy of an estimator (van den Bos, 2007). Relative bias maps were calculated for each framework as $(\bar{\vartheta}_q \vartheta_q) \oslash \vartheta_q$, where $\bar{\vartheta}_q$ and ϑ_q refer to the tissue parameter maps which contain the element-wise sample mean of the $N_{\rm MC}$ estimates $\hat{\vartheta}_q$, and the true reference values, respectively, and where \oslash denotes the element-wise division operator.
- (b) Relative standard deviation. The standard deviation quantifies the precision of an estimator (van den Bos, 2007). Relative standard deviation maps were calculated for each framework as $\left(\frac{N_{MC}}{N_{MC}-1}(\hat{\vartheta}_q \overline{\hat{\vartheta}}_q) \circ (\hat{\vartheta}_q \overline{\hat{\vartheta}}_q)\right)^{\circ \frac{1}{2}} \otimes \vartheta_q$, where \circ and the superscript $\circ \frac{1}{2}$ denote the Hadamard product and element-wise square-root operator, respectively.
- (c) Relative root-mean-squared error (relative RMSE). The RMSE is a measure that incorporates both accuracy and precision. Relative RMSE maps were calculated as $\left(\overline{(\hat{\vartheta}_q \vartheta_q) \circ (\hat{\vartheta}_q \vartheta_q)}\right)^{\circ \frac{1}{2}} \otimes \vartheta_q$.

Additionally, the spatial means of the relative bias, standard deviation and RMSE maps were calculated inside a brain mask, which was extracted from the reference ρ map using the Brain Extraction Tool (BET) (Smith, 2002).

To assess the ability of the different frameworks to estimate motion, the following performance measure was used:

(d) Motion component root-(mean)-mean-squared-error (RMMSE), defined as

$$\left(\frac{1}{N}\sum_{n=1}^{N}\overline{(\hat{\theta}_{n}-\theta_{n})\circ(\hat{\theta}_{n}-\theta_{n})}\right)^{\circ\frac{1}{2}},$$
(6.3.5)

where θ_n refers to the true reference values and the operator $\overline{(\cdot)}$ denotes the element-wise sample mean over the N_{MC} estimates $\hat{\theta}_n$.

For the *in vivo* T_1 mapping experiment, results were quantitatively assessed in terms of spatial resolution and SNR efficiency. Spatial resolution of the obtained parameter maps was assessed in all 3 image dimensions by measuring the average width over 15 edge profiles. The sample of edge profiles was selected in one parameter map (Fig. 6.3.5), and then consistently compared across all the parameter maps of the respective frameworks. The edge width, defined as the width (in high resolution voxels) from 10% to 90% of the edge height, was measured by least squares fitting with a sigmoid function:

$$\eta(x) = a_1 + \frac{a_2}{1 + \exp(-a_3(x - a_4))},$$
(6.3.6)

from which the edge width can be derived, given by $4.4/a_3$ (Greenspan et al., 2002).

Furthermore, SNR measurements were obtained from the *in vivo* reconstruction results for each framework. First, volumes-of-interest (VOIs) were manually delineated in uniform

regions of white matter, CSF, and the caudate nucleus of the ρ map reconstructed with the SRR-joint framework. For the aforementioned tissue types, the VOIs had volumes equal to 100 mm³, 21 mm³, and 48 mm³, respectively. Next, the same VOIs were selected in the T_1 map reconstructed with SRR-joint, and in the ρ and T_1 maps reconstructed with SRR-static and SRR-reg. Subsequently, the SNR was calculated in each VOI as the ratio of the spatial mean to the standard deviation.



Figure 6.3.5: Spatial resolution assessment by means of edge profile fitting, drawn across 3 linesof-interest for different inserts along each orthogonal plane, and compared for the three SRR frameworks.

6.3.4 Implementation

All algorithms were written in MATLAB and partially in C++, and run on a computer with an Intel[®] CoreTM i7-6850K hexa-core CPU with 15MB of cache clocked at 3.60 GHz, with 32 GB of RAM. The computational complexity of the proposed SRR-joint algorithm is primarily defined by the Fast Fourier Transform (FFT)-based image warping operators M_{θ_n} and G_n in the forward model (6.2.1). The FFT-based implementation allows to solve the inverse SRR problem using exact adjoint image warping, and avoids inaccuracies caused by an approximate inverse of the motion. Furthermore, M_{θ_n} is analytically differentiable w.r.t. θ_n . To speed up reconstruction, the FFT's of these image warping operators are executed on the GPU, reducing reconstruction time by a factor of 2-6 compared to pure MATLAB code, mainly dependent on the number of LR images and corresponding image dimensions. In addition, as mentioned in section 6.2.2.5, MATLAB parallel computing tools were used to estimate θ_n for each value of n separately when solving problem (P.2) of the alternating minimization method. Similarly, voxel-wise NLLS model fitting during the initialization step of the different SRR frameworks was performed in a parallel manner. The modified Bessel functions required to calculate the negative log-likelihood function with Rician PDF and the upwind TV prior term, as described in sections 6.2.2.3-6.2.2.4, were implemented using custom C++ MEX-files for use with MATLAB. Also, to avoid excessive memory usage, the Hessian matrix of problem (P.2) was implemented using a Hessian multiply function, which gives the result of a Hessian-times-vector product without computing the Hessian directly. Bearing in mind these implementation details, and given the (rather strict) tolerance criteria described in section 6.3, the reconstruction using SRR-joint took approximately 8.67 hours for a simulated LR T_1 -weighted dataset, 6.43 hours for the *in vivo* T_2 -weighted dataset, and 14.02 hours for the in vivo T1-weighted dataset, respectively. Overall, it is expected that a more advanced implementation of the framework using only C/C++ and GPU/CUDA programming will lead to further reduction of the reconstruction time. In particular, we would like to highlight a CUDA implementation for exact adjoint image warping designed to run on NVIDIA GPUs (Renders et al., 2021), which could potentially be used to speed up the present implementation of the SRR-joint framework. Finally, this proof-of-concept implementation treats $M_{m{ heta}_n}$ and G_n as separate operators. However, the input of both operators could be combined to limit the number of FFT's and improve the computational efficiency.

6.4 Results

6.4.1 Whole brain simulations

Table 6.4.1 summarizes the quantitative performance measures that were obtained from the whole brain simulation experiment for the frameworks SRR-static, SRR-reg, and SRR-joint. For each performance measure, the best performing framework is highlighted in shaded green. It follows from Table 6.4.1 that in terms of accuracy SRR-joint clearly outperforms SRR-static (with a factor 2) and SRR-reg. In terms of precision, SRR-static outperforms the other two approaches, as indicated by the lower overall standard deviation. However, in terms of the overall RMSE, SRR-joint performs best, both for T_1 and ρ mapping.

The absence of motion estimation in the SRR framework becomes evident by looking at maps of the relative RMSE (Fig. 6.4.1), for each of the three SRR frameworks. A closer look at these maps, shows the improved performance in terms of accuracy of the SRR-joint framework compared to the other two approaches. Here, the joint estimation of motion parameters allows for a more accurate estimation of tissue parameters at tissue interfaces, in particular for interfaces at tissue types with longer T_1 relaxation times such as the corpus callosum, and voxels at the periphery of the brain. Additionally, maps of the absolute value of the relative bias and of the relative standard deviation are shown in Fig. 6.4.2 and Fig. 6.4.3.

Table 6.4.1: Quantitative performance measures with standard error (SE) and 95% confidence intervals (CI) for the whole brain simulations, calculated over $M_{MC} = 8$ reconstruction results, for each SRR framework.

		SRR-	static		SRR	t-reg		SRR	-joint
	value	SE	CI	value	SE	CI	value	SE	CI
Overall rel. bias \mathcal{T}_1 [%] ρ [%]	16.840 64.530	0.030 4.014	(16.782,16.899) (56.663,72.398)	10.121 31.875	0.020 2.083	(10.081,10.161) (26.793,34.957)	8.698 24.075	0.016 1.712	(8.667,8.731) (20.719,27.431)
Overall rel. std. de \mathcal{T}_1 [%] ρ [%]	o.316 0.316 0.842	0.001 0.039	(0.314,0.318) (0.765,0.919)	0.591 1.072	0.002 0.034	(0.589,0.595) (1.006,1.139)	0.686 1.313	0.002 0.052	(0.682,0.692) (1.211,1.415)
Overall rel. RMSE \mathcal{T}_1 [%] ρ [%]	16.855 64.547	0.030 4.014	(16.796,16.914) (56.679,72.415)	10.176 30.945	0.020 2.083	(10.136,10.216) (26.863,35.028)	8.799 24.204	0.016 1.713	(8.767,8.831) (20.846,27.561)
RMMSE									
$oldsymbol{t}_{ imes}$ [mm]	0.374	0	n/a	0.095	0.004	(0.085, 0.105)	0.063	0.006	(0.049, 0.079)
$oldsymbol{t}_{Y}$ [mm]	1.381	0	n/a	0.322	0.018	(0.280, 0.364)	0.027	0.009	(0.006,0.049)
$t_z [{\sf mm}]$	1.543	0	n/a	0.289	0.037	(0.203, 0.375)	0.043	0.017	(0.003, 0.083)
α [degree]	1.207	0	n/a	0.335	0.049	(0.219, 0.453)	0.019	0.006	(0.005, 0.035)
eta [degree]	0.245	0	n/a	0.109	0.013	(0.079,0.139)	0.015	0.003	(0.007,0.023)
\sim [dearee]	0.261	С	n/a	0.079	0.004	(0.069.0.089)	0.008	0.003	(0.001.0.017)



Figure 6.4.1: Relative RMSE maps for T_1 and ρ , calculated from the reconstruction results of the synthetic whole brain simulations. For each of the different model-based SRR frameworks orthogonal mid-slice views are shown. Numbers at the bottom of the images indicate the overall relative RMSE measure, which was obtained by calculating the spatial mean of the corresponding relative RMSE map.



Figure 6.4.2: Absolute value of the relative bias maps for T1 and ρ , calculated from the reconstruction results of the synthetic whole brain simulations. For each of the different model-based SRR frameworks orthogonal mid-slice views are shown. Numbers at the bottom of the images indicate the overall relative bias measure, which was obtained by calculating the spatial mean of the absolute value of the corresponding relative bias map.



Figure 6.4.3: Relative standard deviation maps for T1 and ρ , calculated from the reconstruction results of the synthetic whole brain simulations. For each of the different model-based SRR frameworks orthogonal mid-slice views are shown. Numbers at the bottom of the images indicate the overall relative standard deviation measure, which was obtained by calculating the spatial mean of the corresponding relative standard deviation map.

6.4.2 In Vivo Data

Fig. 6.4.4 shows orthogonal mid-slice views of a directly acquired IR TSE T_1 -weighted image with low through-plane resolution sampled at TI = 100 ms, and a synthesized T_1 -weighted image with high through-plane resolution that was produced from the SRR T_1 and ρ parameter map estimates at the same TI. To ease qualitative comparison, zoomed image regions are shown indicating noticeable resolution improvements. The corresponding quantitative T_1 relaxation and ρ parameter map estimates that were obtained using the proposed SRR-joint framework are also shown in Fig. 6.4.4. The improved resolution in each orthogonal plane is clearly visible. In particular, SRR manages to recover the fine details lost due to the acquisition with low through-plane resolution.

Next, to compare the reconstruction results for the LR-T1w *in vivo* dataset, Fig. 6.4.5 shows the estimated T_1 and ρ parameter maps obtained using SRR-static, SRR-reg, and SRR-joint, respectively. Fig. 6.4.5 also shows the absolute value of the relative difference between the reconstructed parameter maps obtained with SRR-static and SRR-reg, taking the corresponding parameter maps obtained with the SRR-joint framework as a reference. Based on Fig. 6.4.5, it can be deduced that the joint estimation of motion parameters yields visible differences at the tissue interfaces, with a noticeably better delineation of the various brain structures. This is also confirmed by the edge width measurements for the T_1 and ρ parameter maps summarized in Table 6.4.2, where SRR-joint achieves smaller edge widths, i.e. a higher spatial resolution, for all parameter maps as compared to SRR-static and SRR-reg. Furthermore, SNR measurements for the selected VOIs in the reconstructed tissue parameter maps of the *in vivo* data experiment were consistently higher for SRR-joint as compared to the other two frameworks, except for the SNR value of CSF in the ρ parameter map (Table 6.4.2).

		SRR-	-static		SRF	₹-reg		SRR	-joint
	value	SE	CI	value	SE	CI	value	SE	CI
Average edge widt	h								
$T_1 map [mm]$	3.878	0.389	(2.995,4.762)	3.759	0.423	(2.801, 4.718)	3.671	0.412	(2.738,4.605)
ρ map [mm]	2.920	0.357	(2.111,3.729)	2.499	0.402	(1.587,3.413)	2.389	0.468	(1.328,3.450)
$\overline{SNR_{VOI}}$ in \mathcal{T}_1 map	J								
white matter	97.103	5.094	(88.064,107.246)	99.656	5.336	(90.073,109.803)	133.730	8.161	(119.206,150.323)
CSF	30.131	7.861	(19.287,43.518)	32.412	6.393	(21.860,43.253)	34.769	8.944	(22.384,48.906)
caudate nucleus	35.211	16.113	(17.852,65.443)	38.697	14.549	(21.356,64.944)	45.074	16.729	(24.307,75.419)
$\overline{SNR_{VOI}}$ in ρ map									
white matter	66.084	4.097	(57.727,73.382)	85.208	4.153	(77.596,93.347)	93.667	4.996	(84.195,102.947)
CSF	19.632	3.688	(14.403,25.933)	22.897	4.189	(16.213,29.394)	19.456	4.046	(14.295, 26.428)
caudate nucleus	42.494	2.946	(37.241, 48.165)	55.883	4.904	(47.988,65.498)	59.989	5.136	(51.068, 69.679)

summarized for each SRR framework.	Table 6.4.2: Quantitative performance measures with standard errc
	or (SE)
) and 95% confidence intervals (
	(CI) for the <i>ir</i>
	VIVO
	T_1 mapping
	experiment,



Figure 6.4.4: Orthogonal mid-slice views with zoomed close-ups showing the resolution improvement for a directly acquired IR TSE T_1 -weighted image with low through-plane resolution sampled at TI₁ (first column), compared to a synthesized T_1 -weighted image with high through-plane resolution (second column), that was produced from the SRR T_1 and ρ parameter map estimates (columns 3 and 4) sampled at the same inversion time. Note that for the LR-T1w *in vivo* data set, SRR-joint can recover the fine details lost to the acquisition with low through-plane resolution. Dashed lines indicate the slice locations.



Figure 6.4.5: Reconstruction results for the LR-T1w *in vivo* dataset showing orthogonal mid-slice views of the quantitative T_1 and ρ parameter maps obtained using SRR-static (left column), SRR-reg (middle column), and SRR-joint (right column), respectively. For comparison reasons, the absolute value of the relative difference maps for T_1 and ρ is shown, which is calculated using the SRR-joint reconstruction result as relative reference. Numbers in boxes represent the overall relative difference maps.

In addition, Fig. 6.4.6 shows a directly acquired IR TSE T_1 -weighted image with low throughplane resolution compared to synthesized T_1 -weighted images with high through-plane resolution, that were produced from the SRR T_1 and ρ parameter map estimates for each framework. Note that SRR-joint outperforms SRR-static and SRR-reg, showing enhanced delineation of brain structures, as indicated by the yellow arrows for different regions of interest.



Figure 6.4.6: Orthogonal mid-slice views with zoomed close-ups showing the resolution improvement for a directly acquired IR TSE T_1 -weighted image with low through-plane resolution sampled at TI₁ (first column), compared to synthesized T_1 -weighted images with high through-plane resolution (columns 2-4), that were produced from the SRR T_1 and ρ parameter map estimates for each framework, sampled at the same echo time.

The reconstruction results for the SRR-joint framework on the *in vivo* T_2 -weighted dataset are summarized in Fig. 6.4.7. This figure shows orthogonal mid-slice views of a directly acquired MESE T_2 -weighted image with low through-plane resolution sampled at TE = 42.7 ms, a synthesized T_2 -weighted image with high through-plane resolution that was produced from the SRR T_2 and ρ parameter map estimates sampled at the same TE, and the obtained T_2 and ρ parameter map estimates, respectively. From Fig. 6.4.7, it can be appreciated that SRR-joint enhances the spatial resolution, and reduces the partial volume effects present in the acquired MESE T_2 -weighted images with low through-plane resolution. As a result, the interfaces and fine structural details of the different tissue types appear more clear in the quantitative T_2 and ρ parameter maps. Furthermore, to visually compare how the different

CHAPTER 6

SRR frameworks arrive at different parameter map estimates for the T_2 -weighted dataset, Fig. 6.4.9 shows the absolute value of the relative difference between the reconstructed parameter maps obtained with SRR-static and SRR-reg, taking the parameter maps obtained with the SRR-joint framework as a reference.



Figure 6.4.7: Orthogonal mid-slice views with zoomed close-ups showing the resolution improvement for a directly acquired MESE T_2 -weighted image with low through-plane resolution sampled at TE = 47.2 ms (first column), compared to a synthesized T_2 -weighted image with high through-plane resolution (second column), that was produced from the SRR T_2 and ρ parameter map estimates (columns 3 and 4) sampled at the same echo time. Dashed lines indicate the slice locations.

In addition, Fig. 6.4.8 shows a directly acquired MESE T_2 -weighted image with low throughplane resolution compared to synthesized T_2 -weighted images with high through-plane resolution, that were produced from the SRR T_2 and ρ parameter map estimates for each framework. As can be appreciated from Fig. 6.4.8, SRR-joint outperforms SRR-static and SRR-reg, showing enhanced delineation of brain structures and a reduction in noise artifacts. We recall that identical regularization weights were used for the three SRR frameworks. Furthermore, it follows from Fig. 6.4.8 that the SRR-joint-T2w image shows improved detail in the axial view, which is the in-plane orientation of the LR image. This is probably due to reduced through-slice blurring.

Motion parameter estimates obtained using the SRR-reg and SRR-joint framework on the *in vivo* datasets are reported in Fig. 6.4.10. In particular, graphs of the translation and rotation parameters estimated for each LR image number are plotted. LR image numbers were ranked in order of acquisition. As indicated by the order of magnitude of the estimated motion parameters, inter-image rigid motion was less present in the T_1 -weighted dataset as compared to the T_2 -weighted dataset. Although the motion parameter traces look very similar for SRR-joint and SRR-reg, small differences can still be observed that likely contribute to the superior performance of SRR-joint compared to SRR-reg.

Furthermore, by construction of the MESE sequence no inter-image motion should exist between the different LR images (i.e. different echoes) of the same MESE scan in the T_2 -weighted dataset. Indeed, it follows from Fig. 6.4.10 that motion parameter estimates obtained using SRR-joint are consistent for the three LR image numbers corresponding with each MESE number. For SRR-reg, on the other hand, one can observe nonphysical



Figure 6.4.8: Orthogonal mid-slice views with zoomed close-ups showing the resolution improvement for a directly acquired MESE T_2 -weighted image with low through-plane resolution sampled at TE = 47.2 ms (first column), compared to synthesized T_2 -weighted images with high through-plane resolution (columns 2-4), that were produced from the SRR T_2 and ρ parameter map estimates for each framework, sampled at the same echo time. Dashed lines indicate the slice locations.

differences of the motion parameter estimates for the LR image numbers per MESE number. This motion stability property can be further quantified by calculating the mean across MESE scans of the standard deviations across the echoes per individual MESE scan, for each motion parameter. Table 6.4.3 summarizes these values for SRR-reg and SRR-joint. As summarized in Table 6.4.3, SRR-joint arrives at significantly lower sample mean values compared to SRR-reg, indicating superior motion stability performance. Note that, both for SRR-joint and SRR-reg, a rigid motion parameter set was estimated per individual LR image.

Table 6.4.3: Quantification of motion stability performance for the *in vivo* T_2 -weighted data set. Tabulated values indicate the sample mean across the MESE scans of the standard deviations calculated across the echoes per individual MESE scan, for each motion parameter, using the SRR-reg and SRR-joint framework, respectively. Translation values are reported in millimeters, rotation values in degrees. Lower values indicate better performance.

	t _x	t _y	t _z	α	β	γ
	[mm]	[mm]	[mm]	[degree]	[degree]	[degree]
SRR-reg	0.064	0.205	0.958	0.111	0.123	0.052
SRR-joint	0.004	0.023	0.075	0.020	0.005	0.007

Finally, to evaluate the convergence behavior of the three SRR methods in the *in vivo* T_1 and T_2 mapping experiments, Fig. 6.4.11 shows the cost function value and the 2-norm of the residual between the measured LR images and their predictions based on the estimated tissue and motion parameters as a function of the number of iterations. For both *in vivo* experiments, it can be observed that SRR-joint arrives at lower cost function values and lower residual values than SRR-static and SRR-reg.



Figure 6.4.9: Reconstruction results for the LR-T2w *in vivo* dataset showing orthogonal mid-slice views of the quantitative T_2 and ρ parameter maps obtained using SRR-static (left column), SRR-reg (middle column), and SRR-joint (right column), respectively. For comparison reasons, the absolute value of the relative difference maps for T_2 and ρ is shown, which is calculated using the SRR-joint reconstruction result as relative reference. Numbers in boxes represent the overall relative difference maps.



Figure 6.4.10: Graphs of the motion parameter estimates that were obtained for the *in vivo* T_1 -weighted dataset (left column) and T_2 -weighted dataset (right column), using the SRR-reg and SRR-joint framework, respectively. The LR image numbers are ranked in order of acquisition. For the T_2 -weighted dataset, the MESE numbers are indicated (right column, bottom graph) with their corresponding LR image numbers. Translation parameters are reported in millimeters, rotation parameters in degrees.



Figure 6.4.11: Convergence plots showing the cost function value and the residual norm as function of the iterations for the *in vivo* T_1 mapping experiment (left column) and T_2 mapping experiment (right column), respectively.

6.5 Discussion

In this contribution, we presented a Bayesian framework for model-based motion-corrected SRR in qMRI. The framework allows the joint estimation of 3D isotropic HR tissue parameter maps and inter-image motion parameters from a set of multi-slice magnitude images with a low through-plane resolution. The framework's potential was demonstrated in both simulations and real data experiments, using T_1 and T_2 mapping as carrying examples. As follows from Table 6.4.1, the proposed SRR framework with joint motion estimation (SRR-joint) showed superior motion parameter estimation and, at the same time, improved tissue parameter mapping RMSE compared to previously published approaches without (SRR-static) and with (SRR-reg) motion pre-compensation. More specifically, the motion component RMMSE of SRR-joint was about an order of magnitude smaller compared to that of SRR-reg and even more for SRR-static. Furthermore, the overall relative RMSE of the tissue parameters T_1 and ρ for SRR-joint was about 20% smaller compared to that of SRR-reg and about 50% smaller compared to SRR-static. Finally, the proposed SRR-joint framework revealed sharper edges in the real data experiments, providing a noticeably better delineation of brain structures, as compared to SRR-reg and SRR-static.

Our proposed framework is modular with respect to the signal and noise model describing the MR data. That is, the T_1 or T_2 relaxation model used in this work can easily be replaced by any other quantitative signal model. Examples may include SRR strategies for quantification T_2^* -relaxation times of the knee (Smekens et al., 2021), blood flow in single post labeling delay pseudo-Continuous Arterial Spin Labeling (Bladt et al., 2020), or diffusion (Van Steenkiste et al., 2016). In addition, the framework is modular with respect to the assumed distribution of the MR data. Indeed, MR data can be characterized by various noise distributions (other than the Rice distribution), considering either single-coil or multi-coil acquisition systems (den Dekker & Sijbers, 2014). Examples include the noncentral chi distribution, which is valid for magnitude images reconstructed from multi-coil data using the sum-of-squares method

(Constantinides et al., 1997), or data distributions that occur for parallel MRI techniques which perform undersampling of the *k*-space to reduce the acquisition time, such as SENSE or GRAPPA (Aja-Fernández et al., 2016). Our Bayesian joint motion and tissue parameter estimation framework can be easily adapted towards any of these data distributions.

Unlike various SRR methods in the literature that rely on orthogonal slice orientations, our method allows for arbitrary slice orientations, which offers much more flexibility with respect to sampling of the *k*-space and setting the contrast weightings. This increased *k*, *q*-space sampling flexibility is a key asset for optimal experiment design studies aimed at the estimation of quantitative tissue parameters with the highest precision (Poot et al., 2010a; Zhao et al., 2019; Morez et al., 2023). In future work, we intend to investigate, given a fixed acquisition time and relying on Cramér-Rao lower bound analysis, the best slice direction and contrast weighting combination of each of the LR images in terms of the precision with which qMRI parameters can be estimated with our proposed SRR-joint framework. Preliminary results of this study for SRR-static have recently been reported by Nicastro *et al.* (Nicastro et al., 2020).

The current framework has some limitations. First, while the framework corrects for motion between the LR multislice images, intra- and inter-slice motion is not yet accounted for. To compensate for intra-slice motion, our framework could be combined with prospective motion correction strategies (Gao et al., 2021; Maclaren et al., 2013). Furthermore, inter-slice motion could be accounted for by adding motion parameters for each individual slice of the LR images and estimating these parameters jointly with the HR tissue parameter maps. Note, however, that although such an approach may improve the accuracy of the estimated maps, the addition of extra parameters to be estimated comes at the expense of a reduced precision. Hence, both effects should be carefully weighed against each other. The extension of our framework to include inter-slice motion and the trade-off between accuracy and precision that comes with it are subject of future investigation.

Second, in this contribution the hyperparameters of the prior distributions (6.2.9) are selected by casting the Bayesian MAP estimation problem as a regularized optimization problem of which the regularization weights are chosen empirically, aiming at equal contributions of the different regularization terms. This approach may be sub-optimal. To the best of our knowledge, however, there is no consensus on the optimal selection strategy of regularization parameters in a multi-parameter nonlinear regression problem like the one at hand. Nevertheless, we hypothesize that choosing the (hyperparameters of the) prior distributions based on prior acquisitions or learning them from available (q)MRI databases may be promising alternative approaches, which are subject of ongoing research.

Finally, the existence of fast 2D multi-slice protocols for MR relaxometry parameter mapping is crucial to fully exploit the benefits of model-based SRR, and to allow for clinically acceptable scan times. The proof-of-concept acquisitions in this chapter aim to illustrate the advantages of joint motion estimation. The combination of model-based SRR with state-of-the-art sequences applying undersampling strategies to further reduce acquisition time is subject of future work. As an example, recent work for T_2 mapping discussed the use of GRAPPATINI (Hilbert et al., 2018), a fast prototype sequence allowing for block-based Cartesian undersampling of *k*-space combined with additional GRAPPA acceleration. Its potential for model-based SRR has been previously reported (Bano et al., 2020), albeit without any appropriate motion estimation routine for SRR. We are convinced that our contribution can serve as an extension to such an approach and to other model-based SRR frameworks that only account for motion using pre-compensation routines.

6.6 Conclusion

In conventional model-based SRR approaches for qMRI, it is common practice to compensate for motion prior to the SRR, e.g. by using a pre-registration routine. However, as demonstrated in this chapter, this conventional two-step approach lacks high accuracy motion estimation and leads to biased parameter estimates. Hence, we have proposed a rigorous unified framework for model-based SRR with joint motion estimation using a Bayesian Maximum A Posteriori (MAP) estimator. The framework allows the joint estimation of 3D isotropic HR tissue parameter maps and inter-image motion parameters from a set of multi-slice magnitude images with a low through-plane resolution. Our SRR framework, which is modular with respect to the quantitative signal model and the assumed distribution of the MR data, has been validated in synthetic whole brain simulations and also with two *in vivo* human brain data sets, for T_1 and T_2 mapping, respectively. It has been demonstrated that the proposed SRR framework provides a more detailed delineation of brain structures and shows superior motion parameter estimation and improved tissue parameter mapping RMSE compared to state-of-the-art SRR approaches.

Acknowledgments

The authors gratefully acknowledge support from Belgian Science Policy Prodex (Grant ISLRA 2009–1062), and the Research Foundation Flanders (FWO) through project funding G084217N and 12M3119N. In addition, this project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 764513. Finally, the authors would like to thank Floris Vanhevel and Céline Smekens for their help during the acquisition of the *in-vivo* datasets and Prof. Dr. Paul M. Parizel for his kind support in the initial phase of the project.

Appendices

6.A Analytical derivatives for joint MAP optimization

The proposed joint MAP estimation consists of the following iterative recursive procedure (see section 6.2.2.5):

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \arg\min_{\boldsymbol{\theta}} \mathcal{L}_{\tilde{\boldsymbol{s}}}(\hat{\boldsymbol{\vartheta}}^{(t)}, \boldsymbol{\theta} | \tilde{\boldsymbol{s}})$$
(P.1)

$$\hat{\boldsymbol{\vartheta}}^{(t+1)} = \arg\min_{\boldsymbol{\vartheta}} \left[\mathcal{L}_{\tilde{\boldsymbol{s}}}(\boldsymbol{\vartheta}, \hat{\boldsymbol{\theta}}^{(t+1)} | \tilde{\boldsymbol{s}}) + \sum_{q=1}^{Q} \frac{2}{\lambda_{q}} \mathsf{TV}(\boldsymbol{\vartheta}_{q}) \right]$$
(P.2)

with

$$\mathcal{L}_{\tilde{s}}(\vartheta, \theta | \tilde{s}) = \sum_{n=1}^{N} \mathcal{L}_{\tilde{s}_n}(\vartheta, \theta_n | \tilde{s}_n) = -\sum_{n=1}^{N} \log p(\tilde{s}_n | \vartheta, \theta_n)$$
(6.A.1)

where the summation runs over all N contrast-weighted low-resolution (LR) images \tilde{s}_n . Problems (P.1) and (P.2) are minimized using a trust-region Newton method (Coleman & Li, 1994). Such a gradient-based optimization algorithm benefits from having analytical expressions for the Jacobian and Hessian to avoid time-consuming finite difference computations. These analytical expressions are derived hereafter.

Nomenclature In what follows, we rewrite the forward operator sequence as $A_n = DBG_n$ and its adjoint sequence as $A_n^T = G_n^T B^T D^T$ to ease the notation, unless stated otherwise. As such, the forward model introduced in Eq. (6.2.1) of section 6.2.1, can be written more concisely as:

$$\boldsymbol{s}_n = |\boldsymbol{D}\boldsymbol{B}\boldsymbol{G}_n\boldsymbol{M}_{\boldsymbol{\theta}_n}\boldsymbol{r}_n| = |\boldsymbol{A}_n\boldsymbol{M}_{\boldsymbol{\theta}_n}\boldsymbol{r}_n|. \tag{6.A.2}$$

6.A.1 MAP estimation of motion parameters

Assuming no dependence of $\{\theta_n\}_{n=1}^N$ through index *n*, the rigid inter-image motion parameter optimization problem (P.1) can be decoupled into *N* parallel subproblems. In what follows, the optimization of a single rigid motion set θ_n corresponding with LR image \tilde{s}_n is considered.

The cost function of this estimation problem is given by

$$\mathcal{L}_{\tilde{s}_{n}}(\vartheta, \theta_{n} | \tilde{s}_{n}) = -\log p_{n}(\tilde{s}_{n}; \vartheta, \theta_{n})$$

$$= \sum_{l=1}^{N_{s}} \left[-\log \tilde{s}_{nl} + \log \sigma_{nl}^{2} + \frac{\tilde{s}_{nl}^{2}}{2\sigma_{nl}^{2}} + \frac{s_{nl}^{2}(\vartheta, \theta_{n})}{2\sigma_{nl}^{2}} - \log I_{0}\left(\frac{\tilde{s}_{nl}s_{nl}(\vartheta, \theta_{n})}{\sigma_{nl}^{2}}\right) \right].$$
(6.A.3)

Keeping only terms that are function of the unknown parameter vector θ_n , as only those are relevant for the minimization, Eq. (6.A.3) simplifies to

$$\mathcal{L}_{\tilde{\boldsymbol{s}}_{n}}(\boldsymbol{\vartheta},\boldsymbol{\theta}_{n}|\tilde{\boldsymbol{s}}_{n}) \sim \sum_{l=1}^{N_{\boldsymbol{s}}} \left[\frac{s_{nl}^{2}(\boldsymbol{\vartheta},\boldsymbol{\theta}_{n})}{2\sigma_{nl}^{2}} - \log I_{0} \left(\frac{\tilde{\boldsymbol{s}}_{nl}s_{nl}(\boldsymbol{\vartheta},\boldsymbol{\theta}_{n})}{\sigma_{nl}^{2}} \right) \right].$$
(6.A.4)

Assuming inter-rigid motion, the motion parameter vector $\theta_n \in \mathbb{R}^{6 \times 1}$ is defined as,

$$\boldsymbol{\theta}_{n} = \{\theta_{nk}\}_{k=1}^{6} = [t_{xn}, t_{yn}, t_{zn}, \alpha_{n}, \beta_{n}, \gamma_{n}]^{T}.$$
(6.A.5)

We then define the gradient w.r.t. the motion parameter θ_{nk} by taking the respective derivative of Eq. (6.A.4):

$$\nabla_{nk}^{\mathcal{L}} = \frac{\partial \mathcal{L}_{\tilde{\boldsymbol{s}}_n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_n | \tilde{\boldsymbol{s}}_n)}{\partial \theta_{nk}} = \boldsymbol{b}_n^{\mathsf{T}} \boldsymbol{c}_{nk}$$
(6.A.6)

where

$$\begin{cases} \boldsymbol{b}_{n} = \frac{\partial \mathcal{L}_{\tilde{\boldsymbol{s}}_{n}}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n} | \tilde{\boldsymbol{s}}_{n})}{\partial \boldsymbol{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})} = \left[\frac{\boldsymbol{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\boldsymbol{\sigma}_{n}^{2}} - \frac{\tilde{\boldsymbol{s}}_{n}}{\boldsymbol{\sigma}_{n}^{2}} \frac{l_{1}\left(\frac{\tilde{\boldsymbol{s}}_{n} \boldsymbol{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\boldsymbol{\sigma}_{n}^{2}}\right)}{l_{0}\left(\frac{\tilde{\boldsymbol{s}}_{n} \boldsymbol{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\boldsymbol{\sigma}_{n}^{2}}\right)} \right] \\ \boldsymbol{c}_{nk} = \frac{\partial \boldsymbol{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\partial \boldsymbol{\theta}_{nk}} = \frac{\partial |\boldsymbol{A}_{n} \boldsymbol{M}_{\boldsymbol{\theta}_{n}} \boldsymbol{r}_{n}|}{\partial \boldsymbol{\theta}_{nk}} = \operatorname{sgn}(\boldsymbol{A}_{n} \boldsymbol{M}_{\boldsymbol{\theta}_{n}} \boldsymbol{r}_{n}) \odot \left(\boldsymbol{A}_{n} \frac{\partial \boldsymbol{M}_{\boldsymbol{\theta}_{n}}}{\partial \boldsymbol{\theta}_{nk}} \boldsymbol{r}_{n}\right) \end{cases}$$
(6.A.7)

175

with $\boldsymbol{b}_n = \{b_{nl}\}_{l=1}^{N_s} \in \mathbb{R}^{N_s \times 1}$, $\boldsymbol{c}_{nk} \in \mathbb{R}^{N_s \times 1}$, and where \odot stands for point-wise multiplication. Finally, substitution of Eq. (6.A.7) in Eq. (6.A.6) results in the following expression for $\nabla_{nk}^{\mathcal{L}} \in \mathbb{R}$

$$\nabla_{nk}^{\mathcal{L}} = \underbrace{\mathbf{r}_{nk}^{\mathcal{T}}}_{\in \mathbb{R}^{1 \times N_{\mathbf{r}}}} \underbrace{\frac{\partial \mathbf{M}_{\boldsymbol{\theta}_{n}}^{\mathcal{T}} \mathbf{A}_{n}^{\mathcal{T}}}{\partial \boldsymbol{\theta}_{nk}} \mathbf{A}_{n}^{\mathcal{T}} \left[\underbrace{\underbrace{\mathrm{sgn}(\mathbf{A}_{n} \mathbf{M}_{\boldsymbol{\theta}_{n}} \mathbf{r}_{n})}_{\in \mathbb{R}^{N_{\mathbf{s}} \times 1}} \odot \underbrace{\left(\underbrace{\frac{\mathbf{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\boldsymbol{\sigma}_{n}^{2}} - \frac{\mathbf{\tilde{s}}_{n}}{\boldsymbol{\sigma}_{n}^{2}} \frac{l_{1}\left(\frac{\mathbf{\tilde{s}}_{n} \mathbf{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\boldsymbol{\sigma}_{n}^{2}}\right)}{l_{0}\left(\frac{\mathbf{\tilde{s}}_{n} \mathbf{s}_{n}(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n})}{\boldsymbol{\sigma}_{n}^{2}}\right)} \right)} \right]} \quad (6.A.8)$$

The exact implementation of $\frac{\partial M_{\theta_n}^T}{\partial \theta_{nk}}$ and the SRR forward model operators is further discussed in Section 6.B hereafter.

6.A.2 MAP estimation of tissue parameters

In contrast to problem (P.1), the tissue parameter estimation problem (P.2) is a large-scale minimization problem. The cost function of this estimation problem is given by

$$\mathcal{L}_{\tilde{\boldsymbol{s}}}(\boldsymbol{\vartheta},\boldsymbol{\theta}|\tilde{\boldsymbol{s}}) + \sum_{q=1}^{Q} \frac{2}{\lambda_{q}} \mathsf{TV}(\boldsymbol{\vartheta}_{q}) = -\sum_{n=1}^{N} \log p_{\tilde{\boldsymbol{s}}_{n}}(\tilde{\boldsymbol{s}}_{n};\boldsymbol{\vartheta},\boldsymbol{\theta}_{n}) + \sum_{q=1}^{Q} \frac{2}{\lambda_{q}} \mathsf{TV}(\boldsymbol{\vartheta}_{q}).$$
(6.A.9)

The tissue parameter maps to be inferred are $\boldsymbol{\vartheta} = \{\boldsymbol{\vartheta}_q\}_{q=1}^Q \in \mathbb{R}^{N_r \times Q}$, with $\boldsymbol{\vartheta}_q = \{\boldsymbol{\vartheta}_{qj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ the q^{th} tissue parameter map and $\boldsymbol{\vartheta}_{\bullet j} \in \mathbb{R}^{Q \times 1}$ all tissue parameters of the j^{th} voxel of $\boldsymbol{\vartheta}_q$. The gradient of the cost function w.r.t. the tissue parameter element $\boldsymbol{\vartheta}_{qj}$ can be written as:

$$\nabla_{\vartheta_{qj}}^{\mathcal{L}} = \frac{\partial \mathcal{L}_{\tilde{s}}(\vartheta, \theta | \tilde{s})}{\partial \vartheta_{qj}} = \sum_{n=1}^{N} \sum_{l=1}^{N_s} \frac{\partial \mathcal{L}_{\tilde{s}_{nl}}(\vartheta, \theta_n | \tilde{s}_n)}{\partial \vartheta_{qj}} = \sum_{n=1}^{N} \sum_{l=1}^{N_s} b_{nl} \frac{\partial s_{nl}(\vartheta, \theta_n)}{\partial \vartheta_{qj}} = \sum_{n=1}^{N} \sum_{l=1}^{N_s} b_{nl} J_{nl,qj}.$$
(6.A.10)

Here, $J_{nl,qj}$ denotes the elements of the Jacobian matrix, which can be further expressed by

$$J_{nl,qj} = \frac{\partial \boldsymbol{s}_{nl}(\boldsymbol{\vartheta},\boldsymbol{\theta}_n)}{\partial \vartheta_{qj}} = \frac{\partial \left| \sum_{j=1}^{N_r} \boldsymbol{A}_n \boldsymbol{M}_{\boldsymbol{\theta}_n} f_n\left(\boldsymbol{\vartheta}_{\bullet j}\right) \right|}{\partial \vartheta_{qj}} = \operatorname{sgn}\left(\varphi_n\right) \sum_{j=1}^{N_r} \boldsymbol{A}_n \boldsymbol{M}_{\boldsymbol{\theta}_n} \frac{\partial f_n\left(\boldsymbol{\vartheta}_{\bullet j}\right)}{\partial \vartheta_{qj}},$$
(6.A.11)

where we write $\varphi_n = \sum_{j=1}^{N_r} A_n M_{\theta_n} f_n(\vartheta_{\bullet j})$ to ease the notation in what follows.

Furthermore, the upwind Total Variation term $TV(\vartheta_q)$, as described in section 6.2.2.4, is given by:

$$\mathsf{TV}(\boldsymbol{\vartheta}_q) = \sum_{j} \left[\sqrt{\zeta_{qj}} - \epsilon \right]$$
(6.A.12)

with

$$\zeta_{qj} = \epsilon^2 + \sum_{m \in \{x, y, z\}} \left[\left(\Delta^{m, +}(\vartheta_{qj}) \right)^2 + \left(\Delta^{m, -}(\vartheta_{qj}) \right)^2 \right].$$
(6.A.13)

176

The derivative of Eq. (6.A.12) w.r.t. element ϑ_{qi} is then given by

$$\frac{\partial \mathsf{TV}(\boldsymbol{\vartheta}_q)}{\partial \boldsymbol{\vartheta}_{qj}} = \frac{1}{2} \sum_{j} \left(\zeta_{qj}\right)^{-1/2} \frac{\partial \zeta_{qj}}{\partial \boldsymbol{\vartheta}_{qj}}.$$
(6.A.14)

Note that a small offset $\epsilon > 0$ is introduced in Eq. (6.A.13) to avoid derivative singularities of TV when ϑ_q is locally constant.

The second order derivatives of cost function $\mathcal{L}_{\tilde{s}}(\vartheta, \theta | \tilde{s})$ w.r.t. the tissue parameter elements ϑ_{qj} can be calculated by taking the derivatives one order higher:

$$H_{\vartheta_{qj}\vartheta_{q'j'}}^{\mathcal{L}} = \frac{\partial}{\partial\vartheta_{qj}} \left(\nabla_{\vartheta_{q'j'}}^{\mathcal{L}} \right) = \sum_{n=1}^{N} \sum_{l=1}^{N_s} \frac{\partial}{\partial\vartheta_{qj}} \left(b_{nl} J_{nl,q'j'} \right) = \sum_{n=1}^{N} \sum_{l=1}^{N_s} \left(\nabla_{\vartheta_{qj}}^{b_{nl}} J_{nl,q'j'} + b_{nl} \nabla_{\vartheta_{qj}}^{J_{nl}} \right).$$

$$(6.A.15)$$

Using the shorthand notation $z_{nl} = \frac{\tilde{s}_{nl}s_{nl}(\boldsymbol{\vartheta},\boldsymbol{\theta}_n)}{\sigma_{nl}^2}$, the gradient terms $\nabla_{\boldsymbol{\vartheta}_{qj}}^{b_{nl}}$ and $\nabla_{\boldsymbol{\vartheta}_{qj}}^{J}$ are given by

$$\nabla_{\vartheta_{qj}}^{b_{nl}} = \frac{\partial b_{nl}}{\partial \vartheta_{qj}}
= \frac{\partial}{\partial \vartheta_{qj}} \left(\frac{\partial \mathcal{L}_{\tilde{s}_{nl}}(\vartheta, \theta_n | \tilde{s}_n)}{\partial s_{nl}(\vartheta, \theta_n)} \right)
= \frac{\partial^2 \mathcal{L}_{\tilde{s}_{nl}}(\vartheta, \theta_n | \tilde{s}_n)}{\partial s_{nl}^2(\vartheta, \theta_n)} \frac{\partial s_{nl}(\vartheta, \theta_n)}{\partial \vartheta_{qj}}
= \left[\frac{1}{\sigma_{nl}^2} - \frac{\tilde{s}_{nl}^2}{\sigma_{nl}^4} \left[1 - \frac{1}{z_{nl}} \frac{l_1(z_{nl})}{l_0(z_{nl})} - \frac{l_1^2(z_{nl})}{l_0^2(z_{nl})} \right] \right] J_{nl,qj}, \qquad (6.A.16)$$

$$\nabla_{\vartheta_{qj}}^{J_{nl}} = \frac{\partial J_{nl,q'j'}}{\partial \vartheta_{qj}}
= \frac{\partial}{\partial \vartheta_{qj}} \left(\frac{\partial s_{nl}(\vartheta, \theta_n)}{\partial \vartheta_{q'j'}} \right)
= \frac{\partial}{\partial \vartheta_{qj}} \left(\operatorname{sgn}(\varphi_n) \right) \sum_{j=1}^{N_r} A_n M_{\theta_n} \frac{\partial f_n(\vartheta_{\bullet j})}{\partial \vartheta_{q'j'}} + \operatorname{sgn}(\varphi_n) \sum_{j=1}^{N_r} A_n M_{\theta_n} \frac{\partial^2 f_n(\vartheta_{\bullet j})}{\partial \vartheta_{qj} \partial \vartheta_{q'j'}}
= \operatorname{sgn}(\varphi_n) \sum_{j=1}^{N_r} A_n M_{\theta_n} \frac{\partial^2 f_n(\vartheta_{\bullet j})}{\partial \vartheta_{qj} \partial \vartheta_{q'j'}},$$
(6.A.17)

where we have used that $\frac{d \operatorname{sgn}(x)}{dx} = 2\delta(x)$.

The partial derivatives $\frac{\partial f_n(\boldsymbol{\vartheta}_{\bullet,j})}{\partial \vartheta_{qj}}$ and $\frac{\partial^2 f_n(\boldsymbol{\vartheta}_{\bullet,j})}{\partial \vartheta_{qj}\partial \vartheta_{q'j'}}$ depend on the signal model of choice. In this work, a **T1-relaxometry signal model** was adopted as a showcase example (Barral et al., 2010):

$$f_n(\boldsymbol{\vartheta}_{\bullet j}) = \rho_j \left(1 - 2 \, e^{-\frac{\mathrm{Tl}_n}{T_{1,j}}} \right), \qquad (6.A.18)$$

with $\vartheta_{\bullet j} = [\rho_j, T_{1,j}]^T$ the tissue parameter vector at position x_j . A more extensive description of this signal model is given in section 6.2.1. The signal model considers Q = 2 tissue

parameter maps. Keeping track of the HR voxel index $j = 1, ..., N_r$, and tissue parameter index q = 1, ..., Q, the first and second order derivatives of $f_n(\vartheta_{\bullet j})$ w.r.t. the tissue parameters ϑ_{qj} are defined by Eq. (6.A.19) and Eq. (6.A.20), which are given as

$$\frac{\partial f_n(\boldsymbol{\vartheta}_{\bullet j})}{\partial \vartheta_{1j}} = 1 - 2e^{-\frac{\mathsf{TI}_n}{\mathsf{T}_{1j}}}, \qquad \frac{\partial f_n(\boldsymbol{\vartheta}_{\bullet j})}{\partial \vartheta_{2j}} = -2\rho_j e^{-\frac{\mathsf{TI}_n}{\mathsf{T}_{1j}}} \left(\frac{\mathsf{TI}_n}{(\mathsf{T}_{1j})^2}\right) \qquad (6.A.19)$$

$$\frac{\partial^2 f_n(\boldsymbol{\vartheta}_{\bullet j})}{\partial \vartheta_{1j}^2} = 0, \qquad \frac{\partial^2 f_n(\boldsymbol{\vartheta}_{\bullet j})}{\partial \vartheta_{2j}^2} = -2\rho_j e^{-\frac{\mathsf{TI}_n}{\mathsf{T}_{1j}}} \left(\frac{\mathsf{TI}_n}{(\mathsf{T}_{1j})^3}\right) \left(\frac{\mathsf{TI}_n}{\mathsf{T}_{1j}} - 2\right),$$

$$\frac{\partial^2 f_n(\boldsymbol{\vartheta}_{\bullet j})}{\partial \vartheta_{1j} \partial \vartheta_{2j'}} = -2e^{-\frac{\mathsf{TI}_n}{\mathsf{T}_{1j}}} \left(\frac{\mathsf{TI}_n}{(\mathsf{T}_{1j})^2}\right). \qquad (6.A.20)$$

Finally, we also give an expression for the second order derivative of the upwind Total Variation prior term in Eq. (6.A.12):

$$\frac{\partial}{\partial \vartheta_{qj}} \left(\frac{\partial \mathsf{TV}(\vartheta_q)}{\partial \vartheta_{q'j'}} \right) = \frac{\partial}{\partial \vartheta_{qj}} \left(\frac{1}{2} \sum_j \left(\zeta_{qj} \right)^{-1/2} \frac{\partial \zeta_{qj}}{\partial \vartheta_{q'j'}} \right) \\
= \frac{1}{2} \sum_j \left[\left(\zeta_{qj} \right)^{-1/2} \frac{\partial^2 \zeta_{qj}}{\partial \vartheta_{qj} \partial \vartheta_{q'j'}} - \frac{1}{2} \left(\zeta_{qj} \right)^{-3/2} \left(\frac{\partial \zeta_{qj}}{\partial \vartheta_{qj}} \right) \left(\frac{\partial \zeta_{qj}}{\partial \vartheta_{q'j'}} \right) \right].$$
(6.A.21)

Please note that for problem (P.2) the Hessian matrix was not explicitly stored in memory, but was implemented as a Hessian multiply function. This function gives the result of a Hessian-times-vector product without computing the Hessian directly, and thus avoids excessive memory usage.

6.B Model operators: implementation and derivatives

6.B.1 Warping operators and derivatives

The proposed SRR framework uses different warping operators in the forward model, described by Eq. (6.2.1) in section 6.2.1. The operator G_n describes the known geometric transformation, extracted from the LR image acquisition header information. This operator models the SRR acquisition, in which multiple LR contrast-weighted images at different orientations are acquired by rotation of the acquisition plane for each image around one fixed encoding axis. A second warping operator M_{θ_n} is introduced to model the effect of unintended rigid inter-image motion. Whereas the motion parameters for G_n are known from the acquisition, the motion parameters $\{\theta_n\}_{n=1}^N$ for M_{θ_n} are unknown, and have to be estimated from the data. The implementation of G_n is identical to that of M_{θ_n} , which will now be discussed.

In what follows, for ease of notation, the LR image index *n* is dropped. Furthermore, the elements of a single rigid motion parameter vector $\boldsymbol{\theta}$ are indexed numerically as $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6)$. In other words, $\theta_1, \theta_2, \theta_3$ correspond with the rigid translations, and $\theta_4, \theta_5, \theta_6$ with the Euler angles of the rigid rotations of Eq. (6.2.2). Similar to (Ramos-Llordén

et al., 2017; Cordero-Grande et al., 2016), the rigid motion is expressed as a series of linear phase modulations in *k*-space:

$$\begin{aligned} \boldsymbol{T}(\theta_{1},\theta_{2},\theta_{3}) &= \mathcal{F}^{H}\boldsymbol{U}(\theta_{1},\theta_{2},\theta_{3})\mathcal{F} \\ \boldsymbol{R}_{1}(\theta_{4}) &= \mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\tan}(\theta_{4})\mathcal{F}_{2}\mathcal{F}_{3}^{H}\boldsymbol{V}_{1}^{\sin}(\theta_{4})\mathcal{F}_{3}\mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\tan}(\theta_{4})\mathcal{F}_{2} \\ \boldsymbol{R}_{2}(\theta_{5}) &= \mathcal{F}_{3}^{H}\boldsymbol{V}_{2}^{\tan}(\theta_{5})\mathcal{F}_{3}\mathcal{F}_{1}^{H}\boldsymbol{V}_{2}^{\sin}(\theta_{5})\mathcal{F}_{1}\mathcal{F}_{3}^{H}\boldsymbol{V}_{2}^{\tan}(\theta_{5})\mathcal{F}_{3} \\ \boldsymbol{R}_{3}(\theta_{6}) &= \mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\tan}(\theta_{6})\mathcal{F}_{1}\mathcal{F}_{2}^{H}\boldsymbol{V}_{3}^{\sin}(\theta_{6})\mathcal{F}_{2}\mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\tan}(\theta_{6})\mathcal{F}_{1}, \end{aligned}$$
(6.B.1)

where \mathcal{F} represents the 3D DFT and \mathcal{F}_k corresponds with the DFT along dimension d, with $d = 1, \ldots, 3$, and where the superscript H denotes the Hermitian conjugate. Both transforms are implemented using MATLAB's built-in FFT functions. In addition, $U \in \mathbb{R}^{N_r \times N_r}$ and $V_d \in \mathbb{R}^{N_r \times N_r}$ are the diagonal matrices that describe, respectively, the applied translation and applied shear decomposed rotations along different axes (Larkin et al., 1997), and whose vectors u and v_d contain the diagonal elements, which are given by:

where k_d is the k-space coordinate vector of the spectral image voxels along dimension d, r_d is the spatial coordinate vector of the image voxels along dimension d, and \circ denotes the Hadamard product.

With this in mind, the rigid motion operator M_{θ} can then be rewritten as

$$\boldsymbol{M}_{\boldsymbol{\theta}} = \boldsymbol{T}(\theta_1, \theta_2, \theta_3) \boldsymbol{R}_1(\theta_4) \boldsymbol{R}_2(\theta_5) \boldsymbol{R}_3(\theta_6). \tag{6.B.3}$$

This helps in defining the partial derivatives of M_{θ} :

$$\frac{\partial \boldsymbol{M}_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}_{k}} = \begin{cases} \frac{\partial \boldsymbol{T}(\theta_{1}, \theta_{2}, \theta_{3})}{\partial \theta_{k}} \boldsymbol{R}_{1}(\theta_{4}) \boldsymbol{R}_{2}(\theta_{5}) \boldsymbol{R}_{3}(\theta_{6}), & 1 \leq k \leq 3\\ \boldsymbol{T}(\theta_{1}, \theta_{2}, \theta_{3}) \boldsymbol{R}_{1}'(\theta_{4}) \boldsymbol{R}_{2}(\theta_{5}) \boldsymbol{R}_{3}(\theta_{6}), & k = 4\\ \boldsymbol{T}(\theta_{1}, \theta_{2}, \theta_{3}) \boldsymbol{R}_{1}(\theta_{4}) \boldsymbol{R}_{2}'(\theta_{5}) \boldsymbol{R}_{3}(\theta_{6}), & k = 5\\ \boldsymbol{T}(\theta_{1}, \theta_{2}, \theta_{3}) \boldsymbol{R}_{1}(\theta_{4}) \boldsymbol{R}_{2}(\theta_{5}) \boldsymbol{R}_{3}'(\theta_{6}), & k = 6, \end{cases}$$
(6.B.4)

where

$$\frac{\partial \boldsymbol{T}(\theta_1, \theta_2, \theta_3)}{\partial \theta_k} = \mathcal{F}^H \frac{\partial \boldsymbol{U}(\theta_1, \theta_2, \theta_3)}{\partial \theta_k} \mathcal{F}, \quad 1 \le k \le 3.$$
(6.B.5)

$$\begin{aligned} \boldsymbol{R}_{1}^{\prime}(\theta_{4}) &= \mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\prime\,\text{tan}}(\theta_{4})\mathcal{F}_{2}\mathcal{F}_{3}^{H}\boldsymbol{V}_{1}^{\sin}(\theta_{4})\mathcal{F}_{3}\mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\tan}(\theta_{4})\mathcal{F}_{2} \\ &+ \mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\tan}(\theta_{4})\mathcal{F}_{2}\mathcal{F}_{3}^{H}\boldsymbol{V}_{1}^{\prime\,\sin}(\theta_{4})\mathcal{F}_{3}\mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\tan}(\theta_{4})\mathcal{F}_{2} \\ &+ \mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\tan}(\theta_{4})\mathcal{F}_{2}\mathcal{F}_{3}^{H}\boldsymbol{V}_{1}^{\sin}(\theta_{4})\mathcal{F}_{3}\mathcal{F}_{2}^{H}\boldsymbol{V}_{1}^{\prime\,\tan}(\theta_{4})\mathcal{F}_{2}, \end{aligned} \tag{6.B.6}$$

$$\begin{aligned} \boldsymbol{R}_{2}^{\prime}(\theta_{5}) &= \mathcal{F}_{3}^{H} \boldsymbol{V}_{2}^{\prime\,\text{tan}}(\theta_{5}) \mathcal{F}_{3} \mathcal{F}_{1}^{H} \boldsymbol{V}_{2}^{\sin}(\theta_{5}) \mathcal{F}_{1} \mathcal{F}_{3}^{H} \boldsymbol{V}_{2}^{\tan}(\theta_{5}) \mathcal{F}_{3} \\ &+ \mathcal{F}_{3}^{H} \boldsymbol{V}_{2}^{\tan}(\theta_{5}) \mathcal{F}_{3} \mathcal{F}_{1}^{H} \boldsymbol{V}_{2}^{\prime\,\text{sin}}(\theta_{5}) \mathcal{F}_{1} \mathcal{F}_{3}^{H} \boldsymbol{V}_{2}^{\tan}(\theta_{5}) \mathcal{F}_{3} \\ &+ \mathcal{F}_{3}^{H} \boldsymbol{V}_{2}^{\tan}(\theta_{5}) \mathcal{F}_{3} \mathcal{F}_{1}^{H} \boldsymbol{V}_{2}^{\sin}(\theta_{5}) \mathcal{F}_{1} \mathcal{F}_{3}^{H} \boldsymbol{V}_{2}^{\prime\,\text{tan}}(\theta_{5}) \mathcal{F}_{3}, \end{aligned} \tag{6.B.7}$$

$$\begin{aligned} \boldsymbol{R}_{3}^{\prime}(\theta_{6}) &= \mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\prime\,\text{tan}}(\theta_{6})\mathcal{F}_{1}\mathcal{F}_{2}^{H}\boldsymbol{V}_{3}^{\sin}(\theta_{6})\mathcal{F}_{2}\mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\tan}(\theta_{6})\mathcal{F}_{1} \\ &+ \mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\tan}(\theta_{6})\mathcal{F}_{1}\mathcal{F}_{2}^{H}\boldsymbol{V}_{3}^{\prime\,\text{sin}}(\theta_{6})\mathcal{F}_{2}\mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\tan}(\theta_{6})\mathcal{F}_{1} \\ &+ \mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\tan}(\theta_{6})\mathcal{F}_{1}\mathcal{F}_{2}^{H}\boldsymbol{V}_{3}^{\sin}(\theta_{6})\mathcal{F}_{2}\mathcal{F}_{1}^{H}\boldsymbol{V}_{3}^{\prime\,\text{tan}}(\theta_{6})\mathcal{F}_{1}. \end{aligned}$$
(6.B.8)

Finally, the derivatives of the diagonal elements of U and V_d can be summarized as

$$\frac{\partial \boldsymbol{u}(\theta_{1}, \theta_{2}, \theta_{3})}{\partial \theta_{d}} = -i\boldsymbol{k}_{d} \circ \boldsymbol{u}(\theta_{1}, \theta_{2}, \theta_{3})$$

$$\frac{\partial \boldsymbol{v}_{1}^{\tan}}{\partial \theta_{4}} = i\left(\frac{1+\tan^{2}(\theta_{4}/2)}{2}\right)\boldsymbol{k}_{2} \circ \boldsymbol{r}_{3} \circ \boldsymbol{v}_{1}^{\tan}$$

$$\frac{\partial \boldsymbol{v}_{2}^{\tan}}{\partial \theta_{5}} = i\left(\frac{1+\tan^{2}(\theta_{5}/2)}{2}\right)\boldsymbol{k}_{3} \circ \boldsymbol{r}_{1} \circ \boldsymbol{v}_{2}^{\tan}$$

$$\frac{\partial \boldsymbol{v}_{3}^{\tan}}{\partial \theta_{6}} = i\left(\frac{1+\tan^{2}(\theta_{6}/2)}{2}\right)\boldsymbol{k}_{1} \circ \boldsymbol{r}_{2} \circ \boldsymbol{v}_{3}^{\tan}$$

$$\frac{\partial \boldsymbol{v}_{1}^{\sin}}{\partial \theta_{4}} = -i\cos(\theta_{4})\boldsymbol{k}_{3} \circ \boldsymbol{r}_{2} \circ \boldsymbol{v}_{1}^{\sin}$$

$$\frac{\partial \boldsymbol{v}_{2}^{\sin}}{\partial \theta_{5}} = -i\cos(\theta_{5})\boldsymbol{k}_{1} \circ \boldsymbol{r}_{3} \circ \boldsymbol{v}_{2}^{\sin}$$

$$\frac{\partial \boldsymbol{v}_{3}^{\sin}}{\partial \theta_{6}} = -i\cos(\theta_{6})\boldsymbol{k}_{2} \circ \boldsymbol{r}_{1} \circ \boldsymbol{v}_{3}^{\sin}.$$
(6.B.9)

Note that this warping operator M_{θ_n} can be shown to be unitary (Ramos-Llordén et al., 2017), which means that its inverse is given by $M_{\theta_n}^H$. Hence, the motion operator M_{θ_n} is reversible, i.e. when applied to an image, this image can be retrieved by applying $M_{\theta_n}^H$ to the output of this operation.

6.B.1.1 Extended implementation details

It is important to discuss some additional implementation steps that were required to use FFT-based image warping with SRR. The following implementation steps were most prominent:

- Large angle extension Based on the original work of Larkin et al. (1997), in which the FFT is used to perform accurate rotations of sampled images, the warping operator has previously been used to jointly estimate small rigid 3D motion of patients in a T1 mapping experiment (Ramos-Llordén et al., 2017). In that work, the rotation angles that correspond with rigid patient motion in the MRI scanner are small and lie invariably in the range of $-45^{\circ} < \theta \le 45^{\circ}$. As discussed by Larkin et al. (1997), this is the theoretical angular range for which wraparound artefacts have little bearing of the overall rotated image. However, in SRR, one is also interested in larger rotation angles when the geometric acquisition of the low-resolution images with rotated slice-encoding direction is to be modelled. Therefore, as a first extension to the work of (Ramos-Llordén et al., 2017), the use of input angles outside the theoretical range was enabled based on geometric considerations, i.e. by applying appropriate dimensional permutations and flips to the image volume being warped, and by adjusting the input angles correspondingly. The large angle extensions for θ_4 , θ_5 , θ_6 are shown in Algorithms 2-4. Note that for each rotation angle the extension algorithm is applied prior to calling the corresponding rotation operator in Eq. (6.B.1).
- **Absence of imaginary component at Nyquist frequency** As described in Appendix A of Larkin et al. (1997), application of the discrete Fourier shift theorem on real images requires a real, band-limited interpolation function that produces only real components at the Nyquist frequency¹. Consequently, for even image dimensions, i.e. an even

¹The Nyquist frequency is the maximum frequency in a FFT of the Nyquist-sampled signal of length N.

number of sample points, it is necessary to multiply the Fourier transformed and shifted function by an additional phase factor that sets the imaginary component to zero at the Nyquist sampling frequency. If this phase factor is not implemented, an imaginary component is created at the Nyquist frequency, and the interpolation process is no longer fully reversible. For a more theoretical outline, the reader is referred to Larkin et al. (1997).

Algorithm 2:	Algorithm 3:	Algorithm 4:
Input: rotation angle $ heta_4$ and image r	Input: rotation angle $ heta_5$ and image $m{r}$	Input: rotation angle $ heta_6$ and image r
Output: adjusted $ heta_4^a$ and $m{r}^a$	Output: adjusted $ heta_5^a$ and r^a	Output: adjusted θ_6^a and r^a
Calculate $\zeta = \theta_4 \mod 360;$	Calculate $\zeta = \theta_5 \mod 360;$	Calculate $\zeta = \theta_6 \mod 360;$
if $0 \leq \zeta \leq 45$ then	if $0 \leq \zeta \leq 45$ then	if $0 \leq \zeta \leq 45$ then
$\theta_4^{\rm a} = \theta_4;$	$\theta_5^a = \theta_5;$	$\theta_6^a = \theta_6;$
$oldsymbol{r}^{a}=oldsymbol{r};$	$oldsymbol{r}^{a}=oldsymbol{r};$	$oldsymbol{r}^{a}=oldsymbol{r};$
else if $45 < \zeta \le 135$ then	else if $45 < \zeta \le 135$ then	else if $45 < \zeta \le 135$ then
$\theta_4^{\rm a} = \operatorname{sgn}(\theta_4) (\zeta - 90);$	$\theta_5^a = \operatorname{sgn}(\theta_5)(\zeta - 90);$	$\theta_6^a = \operatorname{sgn}(\theta_6) (\zeta - 90);$
if $\theta_4 \ge 0$ then	if $\theta_5 \ge 0$ then	if $\theta_6 \ge 0$ then
$ r^{a} = permute(flip(r,2),[1,3,2]);$	$ \boldsymbol{r}^{a} = \text{permute}(\text{flip}(\boldsymbol{r},1),[3,2,1]);$	$ r^{a} = flip(permute(r, [2, 1, 3]), 2);$
else	else	else
$r^{a} = flip(permute(r,[1,3,2]),2);$	$ \mathbf{r}^{a} = flip(permute(\mathbf{r},[3,2,1]),1);$	$r^{a} = permute(flip(r,2),[2,1,3]);$
end	end	end
else if $135 < \zeta \le 225$ then	else if $135 < \zeta \le 225$ then	else if $135 < \zeta \le 225$ then
$\theta_4^{\rm a} = \operatorname{sgn}(\theta_4) \left(\zeta - 180\right);$	$\theta_5^a = \operatorname{sgn}(\theta_5) (\zeta - 180);$	$\theta_6^a = \operatorname{sgn}(\theta_6) (\zeta - 180);$
$\boldsymbol{r}^{a} = flip(flip(\boldsymbol{r},2),3);$	$\vec{r}^{a} = flip(flip(r,1),3);$	$\vec{r}^{a} = flip(flip(r,1),2);$
else if $225 < \zeta \leq 315$ then	else if $225 < \zeta \leq 315$ then	else if $225 < \zeta \leq 315$ then
$\theta_4^{\rm a} = \operatorname{sgn}(\theta_4) (\zeta - 270);$	$\theta_5^a = \operatorname{sgn}(\theta_5) (\zeta - 270);$	$\theta_6^a = \operatorname{sgn}(\theta_6) (\zeta - 270);$
if $\theta_4 \ge 0$ then	if $\theta_5 \ge 0$ then	if $\theta_6 \ge 0$ then
$r^{a} = flip(permute(r,[1,3,2]),2);$	$r^{a} = flip(permute(r,[3,2,1]),1);$	$r^{a} = permute(flip(r,2),[2,1,3]);$
else	else	else
$r^{a} = \text{permute}(\text{flip}(r,2),[1,3,2]);$	$ \mathbf{r}^{a} = \text{permute}(\text{flip}(\mathbf{r},1),[3,2,1]);$	$ r^{a} = flip(permute(r, [2, 1, 3]), 2);$
end	end	end
else if $315 < \zeta \leq 360$ then	else if $315 < \zeta \leq 360$ then	else if $315 < \zeta \leq 360$ then
$\theta_4^a = \operatorname{sgn}(\theta_4) (\zeta - 360);$	$\theta_5^a = \operatorname{sgn}(\theta_5)(\zeta - 360);$	$\theta_6^a = \operatorname{sgn}(\theta_6) (\zeta - 360);$
end	end	end
return $ heta_4^a$, $m{r}^a$;	return $ heta_5^a$, $m{r}^a$;	return $ heta_6^a$, $m{r}^a$;

flip(x, n): flips 3D image volume x along dimension n.

permute(x, n): rearranges the dimensions of 3D image volume x so that they are in the order specified by the 3-element vector n.

Figure 6.B.1: Large angle extension algorithms.

6.B.1.2 Computational complexity analysis

To get an idea of the amount of computing resources that a particular algorithm consumes when it runs, one can perform a computational complexity analysis. This quantifies the order of floating point operations needed, and helps to discover particular run time bottlenecks.

For the FFT-based warping operator at hand, we can use the fact that a 1D FFT transformation has a computational complexity of $\mathcal{O}(N \log_2(N))$ (order of floating point operations), with N the dimension of the 1D vector. For 2D, and for a $M \times N$ image, this becomes $\mathcal{O}(MN \log_2(MN))$. Finally, for a 3D volume of dimensions $M \times N \times P$ the computational complexity can be characterized by $\mathcal{O}(MNP \log_2(MNP))$. If we use isotropic dimensions in 3D ($N \times N \times N$ volume), the computational complexity of a 3D FFT can be characterized by $\mathcal{O}(N^3 \log_2(N^3))$. The computational complexity of the inverse FFT follows similar conventions. As described by Eq. (6.B.1), a single rigid rotation using one of the Euler angles as input, can be implemented as a series of linear phase modulations in *k*-space, so called 'shearing' operations. For a single rotation in Eq. (6.B.1), we count six 1D FFT (or iFFT) operations, each time along different dimensions. In fact we have 3 similar 1D FFT (and iFFT) combinations. As such, the computational complexity of a single rotation can be approximated with a Big-O definition of $\mathcal{O}(3N_{pe} \cdot N_{fe} \cdot N_{se} \log_2(N_{fe} \cdot N_{se}))$ for an overall image size of $N_{pe} \times N_{fe} \times N_{se}$, with N_{pe} , N_{fe} , and N_{se} the volume dimensions along the phase-encoding, frequency-encoding, and slice-encoding dimension, respectively. For a simple $12 \times 12 \times 12$ cubic phantom, similar to what was used in the Monte Carlo simulation study of chapter 5, this would correspond to a floating point operation count in the range of 10 to 250 thousand depending on the amount of zeropadding that was used.

Next, considering all three rigid rotations, as well as the three rigid translations in Eq. (6.B.1), we recognize a total of 18 1D FFT (or iFFT) operations for the rotations, and one 3D FFT and one 3D iFFT combination to implement the translations. As such, the computational complexity of this full warping operator can be approximated with a Big-O definition of: $O\left(9N_{pe} \cdot N_{fe} \cdot N_{se}\log_2(N_{fe} \cdot N_{se}) + (N_{pe} \cdot N_{fe} \cdot N_{se})^2\log_2(N_{pe} \cdot N_{fe} \cdot N_{se})^2\right)$. Taking again the example of a simple 12 × 12 × 12 cubic phantom, typical floating point operation counts are in the range of 60 million to 14 billion.

As summarized in Eq. (6.B.4), the definition of the partial derivatives of M_{θ_n} w.r.t. the rigid motion parameters uses a large number of FFT operations. In fact, we can count a total of 108 1D FFT (or iFFT) operations for the rotations, and 8 3D FFT's (or iFFT's) for translational operations. As such, the computational complexity of the first order derivative of the operator can be approximated with a Big-O definition of:

$$\begin{cases} \frac{\partial M_{\theta}}{\partial \theta_{k}} \end{cases}_{k=1}^{3} \rightarrow \mathcal{O}\left(9N_{pe} \cdot N_{fe} \cdot N_{se} \log_{2}(N_{fe} \cdot N_{se}) + (N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \log_{2}(N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \right) \\ \frac{\partial M_{\theta}}{\partial \theta_{4}} \rightarrow \mathcal{O}\left(15N_{pe} \cdot N_{fe} \cdot N_{se} \log_{2}(N_{fe} \cdot N_{se}) + (N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \log_{2}(N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \right) \\ \frac{\partial M_{\theta}}{\partial \theta_{5}} \rightarrow \mathcal{O}\left(15N_{pe} \cdot N_{fe} \cdot N_{se} \log_{2}(N_{fe} \cdot N_{se}) + (N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \log_{2}(N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \right) \\ \frac{\partial M_{\theta}}{\partial \theta_{5}} \rightarrow \mathcal{O}\left(15N_{pe} \cdot N_{fe} \cdot N_{se} \log_{2}(N_{fe} \cdot N_{se}) + (N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \log_{2}(N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \right) \\ \frac{\partial M_{\theta}}{\partial \theta_{6}} \rightarrow \frac{\mathcal{O}\left(15N_{pe} \cdot N_{fe} \cdot N_{se} \log_{2}(N_{fe} \cdot N_{se}) + (N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \log_{2}(N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \right)}{\mathcal{O}\left(54N_{pe} \cdot N_{fe} \cdot N_{se} \log_{2}(N_{fe} \cdot N_{se}) + 4 \left(N_{pe} \cdot N_{fe} \cdot N_{se}\right)^{2} \log_{2}(N_{pe} \cdot N_{fe} \cdot N_{se})^{2} \right)} \end{cases}$$

For the $12 \times 12 \times 12$ cubic phantom in our simulations, floating point operation counts are in the range of 250 million to 60 billion depending on the amount of zeropadding that was used.

6.B.1.3 Alternative image warping implementations

Image warping operators are crucial to perform the affine transforms that describe the geometric acquisition of each low-resolution image. Besides the FFT-based warping operator with its unitary property, introduced in section 6.B.1, there also exist other image warping operator implementations which have been used successfully in SRR applications.

In the SRR work of Poot et al. (2010b), a warping operator was used that applies an affine transform T_f by splitting it into a (non unique) **S**et of 1D **SH**ear transforms (SSH). The shear transforms \tilde{T}_j , with $j \in \{1, ..., 2N\}$, satisfy $T_f = \prod_{i=1}^{2N} \tilde{T}_j$, with N the dimensions

of the problem. Each \tilde{T}_j differs from the identity matrix only in row d_j , so the ND image is only deformed along its d_j^{th} main axis. This allows for efficient interpolation with a 1D low pass filter. A general affine transform in N dimensions can thus be split in 2N 1D low pass filter steps. The complexity of this approach lies in the tuning of the filtering steps, in particular monitoring their combined effect on the generation of aliasing and geometric/spectral distortions.

Another type of image warping operators worth mentioning here, are the operators based on the use of multivariate spline interpolation to interpolate the regular grid data of the image that is warped. Such operators are linear maps, where each voxel in the warped image is a linear combination of voxels in the original image. A notable example is the implementation using tricubic interpolation, i.e. 3D or trivariate spline interpolation using cubic splines, by Renders et al. (2023). By use of symbolic computer algebra, a list of 64 polynomials can be generated that allows to compute a matrix representation of trivariate cubic image warping. By combining an on-the-fly computation of this matrix with a parallelized implementation of columnwise matrix multiplication, a CUDA accelerated, low memory implementation of the adjoint action of 3D cubic image warping can be obtained.

6.B.2 Blurring operator

The blurring operator **B** in Eq. (6.2.1) describes the point spread function (PSF) of the MRI signal acquisition process. For multislice acquisition methods that sample a rectangular part of *k*-space, the 3D PSF is separable and can be modeled as the product of three 1D PSFs that are applied in the orthogonal directions aligned with the MR image coordinate axis. The PSFs in the frequency and phase encoding direction are defined by the rectangular part of *k*-space that is regularly sampled. In this contribution, an in-plane 2D PSF is constructed as a convolution of two identical Gaussian functions, with a standard deviation set to $0.25 \times \Delta_{in-plane}$, with $\Delta_{in-plane}$ the in-plane resolution (Van Reeth et al., 2015). The remaining through-plane 1D PSF models the slice selection profile (SSP), as SRR relies on rotated SSP cross-talk to enhance the through-plane resolution while keeping the in-plane resolution fixed. In a multislice MRI acquisition, each slice is excited by incorporating a slice selective gradient which is often generated by applying either a (windowed) sinc or a Gaussian shaped RF pulse. In this contribution, the SSP in the slice-direction (i.e. the *z*-direction) corresponds to a windowed sinc slice excitation, and was modeled as a smoothed box function (Poot et al., 2010b):

$$SSP(z; \Delta S) = \begin{cases} 1 & \left|\frac{z}{\Delta S}\right| \le \frac{1}{3} \\ \frac{1}{2} - \frac{1}{2}\sin\left(3\pi\left(\left|\frac{z}{\Delta S}\right| - \frac{1}{2}\right)\right) & \frac{1}{3} < \left|\frac{z}{\Delta S}\right| < \frac{2}{3} \\ 0 & \frac{2}{3} \le \left|\frac{z}{\Delta S}\right| \end{cases}$$
(6.B.10)

where the full width at half maximum (FWHM) of the smoothed box equals the given slice thickness ΔS of the modeled LR images s_n . The spatially invariant blurring of the separable 3D PSF is performed using cyclic convolution, as described in (Hansen et al., 2006), where the blurring operator $\boldsymbol{B} \in \mathbb{R}^{N_r \times N_r}$ and its conjugate transpose $\boldsymbol{B}^H \in \mathbb{R}^{N_r \times N_r}$ are spectrally decomposed as:

$$\boldsymbol{B} = \mathcal{F}_{3}^{H} \boldsymbol{\Lambda}_{3} \mathcal{F}_{3} \mathcal{F}_{12}^{H} \boldsymbol{\Lambda}_{12} \mathcal{F}_{12}$$
(6.B.11)

$$\boldsymbol{B}^{H} = \mathcal{F}_{3}^{H} \boldsymbol{\Lambda}_{3}^{H} \mathcal{F}_{3} \mathcal{F}_{12}^{H} \boldsymbol{\Lambda}_{12}^{H} \mathcal{F}_{12}$$
(6.B.12)

with Λ_{12} the spectrum of a block-circulant-with-circulant-blocks matrix that describes the in-plane convolution, and Λ_3 a sparse diagonal matrix whose diagonal elements are the Fourier coefficients of the first column of a circulant blurring matrix created by circularly shifting the SSP array preceeding row forward. Furthermore, \mathcal{F}_{12} and \mathcal{F}_3 denote the 2D unitary DFT along the in-plane dimensions (d = 1 and d = 2) and the unitary 1D DFT along the through-plane dimension (d = 3), respectively.

6.B.3 Resampling operator

Downsampling along the through-plane direction is required to resample the HR image to a LR image with increased slice thickness. To allow for noninteger resampling, interpolation is required. The choice of interpolation paradigm should allow a straightforward transpose implementation for substitution in the analytical expressions of the Jacobian and Hessian of the gradient-based SRR optimization routine. Therefore, resampling was performed using cubic convolution-based interpolation, which was first introduced in (Keys, 1981). As the original proposition of this type of interpolation is put quite general and extensive, some extra choices are required regarding its computational implementation. To promote full reproducibility of our method, these choices will now be discussed.

Cubic convolution-based interpolation (CCI) (Keys, 1981; Meijering & Unser, 2003) of uniformly sampled data implies the use of an interpolation kernel $u : \mathbb{R} \to \mathbb{R}$, which determines the weights to be assigned to the samples $f_k = f(kT)$ of an original function $f : \mathbb{R} \to \mathbb{R}$ in computing the value of the interpolant g at any arbitrary $x \in \mathbb{R}$. In what follows, for ease of notation, but without loss of generality, we will use T = 1. CCI may then be described as

$$g(x) = \sum_{k \in \mathbb{Z}} f_k u(x - k).$$
(6.B.13)

As can readily be observed from Eq. (6.B.13), it is required that in order for g to be an interpolant, the kernel u must satisfy that u(0) = 1 and u(n) = 0 when n is any nonzero integer. A balanced trade-off between computational cost and accuracy is provided by the family of cubic convolution kernels that consist of piecewise third-degree polynomials and are once continuously differentiable. In this contribution, Keys' third-order cubic convolution kernel is used (Keys, 1981), which is defined as

$$u(x) = \begin{cases} \frac{3}{2}|x|^3 - \frac{5}{2}|x|^2 + 1 & \text{if } 0 \le |x| \le 1, \\ -\frac{1}{2}|x|^3 + \frac{5}{2}|x|^2 - 4|x| + 2 & \text{if } 1 \le |x| \le 2, \\ 0 & \text{if } 2 \le |x|. \end{cases}$$
(6.B.14)

This kernel has an approximation order of L = 3, which implies that the resulting interpolant converges to the original function as fast as the third power of the intersample distance. It also implies that the kernel is capable of reproducing polynomials up to second degree. Outside the interval (-2, 2), the interpolation kernel u(x) is zero. This means that only four data samples are used to evaluate the interpolant at some new position x. In practice, the original function f can only be observed on a finite interval. For values outside this interval, boundary conditions must be chosen. In our contribution, values outside the image matrix are assumed to have a zero weight contribution, i.e. $f_k = 0$, indicating that only three values are used to evaluate the interpolant at the outer background edges of the generated LR image.

The HR image r can be thought of as a function

$$\boldsymbol{r}:[n] \times [m] \times [o] \to \mathbb{R},$$
 (6.B.15)

where $n, m, o \in \mathbb{N}$ and $\forall k \in \mathbb{N} : [k] = \{1, \ldots, k\}$. For each pair of integer coordinates, it yields a HR voxel value. Following 3D volume considerations, downsampling along the third [o] through-plane dimension, i.e. the slice selection dimension of r, corresponds with $[n] \times [m]$ repeated one-dimensional CCI operations. A single CCI at a non-integer position a is given by

$$\mathbf{x}'(a) = c_1 \mathbf{x}(\mathbf{p}_1) + c_2 \mathbf{x}(\mathbf{p}_2) + c_3 \mathbf{x}(\mathbf{p}_3) + c_4 \mathbf{x}(\mathbf{p}_4),$$
 (6.B.16)

where p_1, \ldots, p_4 are the four integer valued points surrounding a, and c_1, \ldots, c_4 are the CCI coefficients obtained by substituting p_1, \ldots, p_4 in Eq. (6.B.14). The downsampling operator D transforms HR image r into a LR image s = Dr of which the (i, j, k)-th voxel value is obtained by

$$(Dr)(i, j, k) = x'((i, j, k)).$$
 (6.B.17)

Since, by Eq. (6.B.16), Eq. (6.B.17) is a linear combination of voxel values of r, we can interpret the action of D as a matrix vector product. Where the vectors are the $[n] \times [m]$ one-dimensional HR through-plane arrays. D can be represented by a matrix, with 4 non-zero coefficients on each row, namely the CCI coefficients of Eq. (6.B.16) at the corresponding voxel indices separated by the inter-slice distance. The adjoint operator D^T is then simply given by the matrix with the rows of D as its columns. The rows of D or equivalently, the columns of D^T can be computed on the fly, so there is no need to store these matrices explicitly. If we denote the *i*-th row of D, i.e. the *i*-th column of D^T by o_i , then the action of D^T on a vector $s \in \mathbb{R}^N$ can be implemented as follows:

$$\boldsymbol{D}^{\mathsf{T}}\boldsymbol{s} = \sum_{i=1}^{N} s_i \boldsymbol{o}_i. \tag{6.B.18}$$

With this approach we obtain an exact adjoint operator D^T that can be substituted in the analytical expressions for the Jacobian and Hessian of the gradient-based optimization routine.

6.B.3.1 Kernel extensions

The original work of Keys (1981) also describes variations of the cubic convolution algorithm. In particular, the construction of an interpolation kernel which has fourth-order accuracy (L = 4) and a higher order of convergence (i.e. a measure of how fast the approximation error goes to zero for decreasing sampling increments), that can be achieved with piecewise cubic polynomials:

$$u(x) = \begin{cases} \frac{4}{3}|x|^3 - \frac{7}{3}|x|^2 + 1 & \text{if } 0 \le |x| \le 1, \\ -\frac{7}{12}|x|^3 + 3|x|^2 - \frac{59}{12}|x| + \frac{15}{6} & \text{if } 1 \le |x| \le 2, \\ \frac{1}{12}|x|^3 - \frac{2}{3}|x|^2 + \frac{21}{12}|x| - \frac{3}{2} & \text{if } 2 \le |x| \le 3, \\ 0 & \text{if } 3 \le |x|. \end{cases}$$
(6.B.19)

Whereas the third-order kernel of Eq. (6.B.14) only requires 4 sample points, the fourthorder interpolation kernel in Eq. (6.B.19) requires 6 sample points for each query point of interpolation, increasing computational demands. Extensive testing of higher-order interpolation kernels in terms of accuracy and computational efficiency is part of ongoing research at the time of writing. For more detailed information, the reader is referred to (Keys, 1981).

6.B.4 Operator validation

In addition to identifying any mathematical or computational operations that could compromise an operator's computational efficiency, as discussed in section 6.B.1.2 for the FFT-based warping operator, it is also important to validate the adjoint (i.e. conjugate transpose) implementation of each forward operator. For this purpose, a set of operator tests can be performed, some of which are briefly explained below.

6.B.4.1 Matrix transpose test

Each linear operator function can be represented by a matrix, that acts (by matrix-vector multiplication) on an image, represented as a raveled vector. This matrix can also be extracted explicitly, by application of the operator function to each column of a unit matrix. In other words, by transforming an image with one voxel set to 1, while all other voxels are set to 0, we can obtain one column of the operator matrix.

By creating such a matrix explicitly for both the forward and transpose operator function, we end up with two matrices that should be each other's transpose, i.e. by interchanging row and column index for each element both matrices should be identical. It should be noted that this implementation test is computationally expensive, as the construction of each matrix representation requires one function operation per voxel. So typically this test is performed only for smaller phantom-like images.

As an example, consider a test for the resampling operator of section 6.B.3, where we resample an HR image r with dimensions $10 \times 10 \times 10$ and voxel size $2 \times 2 \times 2$ mm³ to an LR image s with dimensions $5 \times 10 \times 4$ and voxel size $2.2 \times 2 \times 5$ mm³. Taking into account the raveling of both images as column vectors, the resulting forward operator matrix D has dimensions 200×1000 , and dimensions 1000×200 for its transpose matrix D^{T} . Fig. 6.B.2 shows both constructed operator matrices visually using a colorbar to indicate the values of the matrix elements. The same figure also shows the matrix transpose of the operator by simply interchanging row and column index of each element, denoted as \tilde{D}^{T} , and the corresponding ratio based on element-wise division of D^{T} and \tilde{D}^{T} . Note that for a correct implementation of D^{T} , the ratio test should show values exactly equal to 1 for elements of D that are different from 0.

6.B.4.2 Dot-product test

Storing the linear operators explicitly as matrices results in extensive computer memory requirements. Therefore, it is more memory efficient to implement the operators as matrix-free programming functions. As such, for each operator D and its adjoint D^T there are two functions. The first amounts to the matrix multiplication Dr, while the adjoint routine computes D^Ts . The dot-product test allows to verify that the two routines are adjoint to each other.



Figure 6.B.2: Matrix transpose test for the resampling operator function of section 6.B.3, showing, respectively, the constructed forward operator matrix D, the constructed transpose operator matrix D^{T} , the transpose matrix \tilde{D}^{T} by interchanging row and column index of each element of D, and the ratio matrix created by element-wise division of D^{T} and \tilde{D}^{T} . Note that $\epsilon = 5e^{-14}$.

The associative property in linear algebra states that parentheses in a vector-matrix-vector multiplication are redundant, i.e. the grouping (or association) of the matrices does not change the result. The parentheses only determine the sequence of computation. Using symbolic notation:

$$\boldsymbol{s}^{\mathsf{T}} \left(\boldsymbol{D} \boldsymbol{r} \right) = \left(\boldsymbol{s}^{\mathsf{T}} \boldsymbol{D} \right) \boldsymbol{r} = \left(\boldsymbol{D}^{\mathsf{T}} \boldsymbol{s} \right)^{\mathsf{T}} \boldsymbol{r}$$
(6.B.20)

To perform the dot-product test, construct the random vectors r and s with appropriate dimensions. Using operator function D, compute the matrix-vector product $\tilde{s} = Dr$, and using the adjoint operator function D^{T} , compute $\tilde{r} = D^{T}s$. After substitution of \tilde{s} and \tilde{r} in

Eq. (6.B.20), it follows that:

$$s^{T} (Dr) - (D^{T}s)^{T} r = 0$$

$$\Leftrightarrow s^{T} \tilde{s} - \tilde{r}^{T} r = 0.$$
(6.B.21)

The difference of the scalar values in the left side of Eq. (6.B.21) should be equal to zero to validate that D^{T} is indeed the exact adjoint operator function of D. In practice, due to the computer's inability to represent real numbers with infinite precision, rounding errors can occur which can decrease the accuracy of the dot-product test. To avoid these errors, alternative dot-product algorithms have been investigated in the literature (Ogita et al., 2005).

6.C Applications in musculoskeletal imaging

In addition to neuroimaging, the proposed SRR method with joint motion estimation provides also a suitable and useful tool for other MRI applications. In particular, it was investigated how the SRR method can be applied in musculoskeletal MRI, which focuses specifically on joint structures, including hands, wrists, hips, knees, ankles and so on. A distinction can be made between the use of SRR in the application of *anatomical MRI*, where no signal model is used as part of the SR forward model and the aim is primarily to increase SNR and spatial resolution without increasing the total scan time, and *quantitative MRI*, where a signal model is embedded in the SR forward model to perform a quantitative MRI parameter mapping. Both types of SRR for musculoskeletal MRI are illustrated hereafter, using **knee MRI** as a carrying example. Knee MRI is the most frequently requested musculoskeletal exam, greatly affecting MRI workflow and patient throughput (Beker et al., 2017).

6.C.1 Anatomical knee MRI

We thank Céline Smekens and the radiology team of the Antwerp University Hospital (UZA) for their help in acquiring the *in-vivo* knee dataset that was reconstructed in this appendix section.

Current scanning protocols for knee MRI typically involve 2D intermediate-weighted (IW) and fat-suppressed T2-weighted turbo spin echo (TSE) sequences. These sequences offer excellent tissue contrast and high in-plane resolution, but they are commonly acquired with a large slice thickness, leading to partial volume averaging (Mugler III, 2014; Yao et al., 2007).

As an alternative to multiple 2D TSE acquisitions, all major MRI vendors have introduced 3D TSE sequences, including fast spin echo Cube (GE Healthcare), volume isotropic TSE acquisition (VISTA, Philips Medical Systems), and sampling perfection with application optimized contrast using different flip angle evolutions (SPACE, Siemens Healthcare) (Mugler III, 2014). The key advantage of these sequences lies in their capacity to provide a single-slab isotropic 3D volume covering the entire knee joint, thereby minimizing partial volume effects and eliminating inter-slice gaps (Yao et al., 2007). Moreover, the source data of these 3D sequences can be reformatted in any desired orientation, simplifying the depiction of oblique complex knee structures (e.g., meniscal roots) and obviating the need for multiplanar acquisitions (Mugler III, 2014; Yao et al., 2007; Garwood et al., 2017). Despite these advantages, current 3D TSE sequences still have limitations in terms of image quality due to the use of

long echo train lengths (Naraghi & White, 2012). Image blurring, caused by acquiring high frequencies at later echoes, can be problematic as it diminishes the visibility of low-contrast structures, such as the menisci, potentially hindering accurate diagnosis (Notohamiprodjo et al., 2009; Ristow et al., 2009; Subhas et al., 2011). Additionally, compared with 2D TSE sequences, conventional 3D TSE sequences are characterized by long acquisition times, increasing the likelihood of motion artifacts (Naraghi & White, 2012; Garwood et al., 2017). Efforts in knee MRI research have been directed towards developing strategies to achieve high-resolution isotropic 3D TSE MRI while reducing scan time. For instance, 3D SPACE with 2D controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) (Fritz et al., 2016a) and 3D TSE with compressed sensing (CS) (Fritz et al., 2016b; Lee et al., 2018), have been proposed. Yet, despite their potential, these 3D TSE techniques are not widely adopted in routine knee MRI.

As an alternative to the blur-sensitive and lengthy direct 3D MRI scans, model-based SRR could be employed to enhance the MRI trade-off between SNR, spatial resolution, and scan time (Van Dyck et al., 2020). As a proof-of-concept study and test case for anatomical knee MRI, the proposed SRR framework with joint motion estimation was applied to 2D TSE MRI to obtain high-resolution isotropic 3D knee MRI.

Methods A healthy volunteer was scanned on a 3T MR scanner (Magnetom PrismaFit, Siemens Healthcare, Erlangen, Germany) with a 15-channel knee coil (Quality Electrodynamics, Mayfield Village, OH). The right knee was imaged with a 2D TSE-based SRR protocol, for which the acquisition parameters are summarized in Table 6.C.1. The protocol consisted of 7 repetitions of a 2D TSE sequence with a low through-plane resolution (voxel size, $0.5 \times 0.5 \times 2.0 \text{ mm}^3$). Each acquisition was characterized by a specific rotation around the phase-encoding axis (i.e., 0°, 26°, 51°, 77°, 103°, 129°, 154°), as illustrated in Fig. 6.C.1. The scan time per anisotropic 2D slice stack was 1 minute 55 seconds, resulting in a total scan time of 13 minutes 25 seconds. Model-based super-resolution reconstruction was performed at an isotropic high-resolution grid with a voxel size of $0.5 \times 0.5 \times 0.5 \times 0.5 \text{ mm}^3$. The reconstruction was repeated with and without the use of motion estimation, each time using the same total variation regularization settings.

Sequence	2D Turbo Spin Echo (TSE)	Slice thickness/gap [mm]	2/0
Orientation	Sagittal	Voxel size [mm ³]	0.5 imes 0.5 imes 2.0
TR [ms]	3080	No. excitations	1
TE [ms]	36	Echo spacing [ms]	8.93
Acceleration factor	3 (GRAPPA)	Phase encoding direction	head to feet
Turbo factor	5	Phase sampling [%]	100
Receiver bandwidth [Hx/pixel]	256	Number of repetitions	7
Flip angle [°]	160	Angles of rotation [°]	0, 26, 51, 77, 103, 129, 154 (*)
Field of view [mm ²]	160×160	Total scan time	13 min 25 s
Matrix size	320 × 320		

Table 6.C.1: Acquisition parameters of the 2D TSE sequence	e used for anatomical SRF
--	---------------------------

*See Fig. 6.C.1 for a schematic representation of the SR acquisitions.

Results Fig. 6.C.1 and Fig. 6.C.2 display the results of both anatomical super-resolution reconstructions. Based on a qualitative inspection of these results, it can be observed how the use of joint inter-image motion correction results in improved sharpness and superior delineation of knee structures compared to the approach without motion estimation.



were acquired with high in-plane and low through-plane resolution, while rotating around the phase-encoding axis over angles of 0°, 26°, 51°, 77°, 103° 129°, and 154°. For comparison, the corresponding reconstruction result after model-based SRR with joint motion is shown on the bottom right corner



191

Conclusion It was demonstrated that SRR is technically feasible for 3D high-resolution isotropic IW knee MRI. Furthermore, the addition of joint inter-image motion estimation to the SRR framework provides a clear advantage over SRR without motion estimation. Finally, the reader is referred to the work of Van Dyck et al. (2020), which provides a more elaborate study on the use of SRR for anatomical knee MRI, while using the here proposed joint motion estimation strategy. Also, the comparison of SRR for 2D TSE images against a conventional 3D SPACE acquisition is thoroughly investigated using different performance measures.

6.C.2 Quantitative knee MRI

The work presented in this part of Appendix 6.C was performed in collaboration with Céline Smekens *et al.*, and was published as an ISMRM conference abstract, receiving a *Magna Cum Laude Merit Award*:

C. Smekens*, **Q. Beirinckx***, F. Vanhevel, P. Van Dyck, A. J. den Dekker, J. Sijbers, T. Janssens, and B. Jeurissen, "Super-resolution T_2^* mapping of the knee using UTE Spiral VIBE MRI", in *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM)*, Vol. 29, pp. 3920, 2021. (*Both authors contributed equally.)

Another application of knee MRI where SRR can prove its worth is the noninvasive imaging of the maturation process of the **anterior cruciate ligament (ACL)**, which can get injured during sports-related activities where jumping, pivoting and rapid change of direction occurs. Surgical ACL reconstruction using tendon graft is the standard to treat ACL injuries. However, little is known about the maturation process of human ACL graft and the role of adjacent structural abnormalities herein. As such, there currently exists a **high clinical need for improved noninvasive objective measures of ACL graft properties to help inform return to high-demand activities**. Next to anatomical MRI, qMRI techniques, such as T_2^* relaxometry and diffusion tensor imaging (DTI), have gained interest for musculoskeletal imaging, as they can provide objective measures of biophysical tissue properties that allow for monitoring of the tissue microstructure.

To perform T_2^* mapping and quantitative evaluation of highly organized collagen-rich knee structures with short mean transverse relaxation times it is necessary to use ultrashort echo time (UTE) MRI (de Mello et al., 2019). UTE T_2^* mapping can estimate ultrashort (< 1.0 ms) and short $(1-10 \text{ ms}) T_2^*$ relaxation times based on data sampled at multiple echo times (TEs), starting at 0.5 ms or shorter (de Mello et al., 2019; Chang et al., 2015). When the sampling is extended up to long TEs (> 10 ms), more comprehensive knee T_2^* maps, including tissues with long mean T_2^* , can be reconstructed (Williams et al., **2010)**. However, this wider T_2^* sensitivity comes at the expense of higher acquisition times. Consequently, in vivo validation studies commonly acquire UTE T_2^* -weighted images with lower (> 1 mm) through-plane resolution (Chu & Williams, 2019; Breda et al., 2020). Yet, accompanying partial volume effects may negatively affect the reliability of T_2^* measurements. There is thus a need for high-resolution T_2^* mapping methods that enable accurate and precise estimation of (ultra)short and long T_2^* values in a reasonable scan time. This is exactly where model-based SRR can help in providing a better trade-off between SNR, acquisition time, and spatial resolution compared to standard 3D relaxometry methods. To improve the aforementioned trade-off, the proposed SRR framework was applied to the knee as a proof-of-concept study.

Methods Three asymptomatic volunteers were scanned on a 3T MR scanner (Magnetom PrismaFit, Siemens Healthcare, Erlangen, Germany) with a 15-channel knee coil. Per volunteer, four T_2^* -weighted datasets were acquired with an accelerated prototypical 3D UTE Spiral VIBE sequence (Qian & Boada, 2008; Mugler et al., 2015). Acquisition parameters are listed in Table 6.C.2. As fat suppression leads to increased T_2^* estimates (Kim et al., 2019), acquisitions B, C, and super-resolution were acquired without fat saturation and TEs were chosen as close as possible to the in-phase TEs of water and fat. A TR = 11.40 ms was used for acquisitions A and B (Kim et al., 2019; Smekens et al., 2020), while a TR = 22.50 ms was chosen for acquisition C and super-resolution acquisitions to accommodate more and longer TEs. The super-resolution protocol consisted of five acquisitions with low through-plane resolution rotated around the frequency-encoding direction (see Fig. 6.C.3).

Table 6.C.2: Acquisition parameters of 4 UTE T_2^* mapping protocols based on accelerated 3D UTE Spiral VIBE MRI. A short TR was used for acquisitions A and B (4 TEs in 2 sets), while a longer TR was used for acquisition C and rotated super-resolution (SR) acquisitions with low through-plane resolution (10 TEs acquired in 2 (C) or 5 (SR) sets). SPIRiT: iterative parallel image reconstruction algorithm (Lustig & Pauly, 2010).

	Acquisition A	Acquisition B	Acquisition C	SR Acquisition
Slices per slab (#)	176	176	176	60
Acquisition matrix	224×224	224×224	224×224	224×224
Field of view [mm ³]	190×190×149.6	190×190×149.6	190×190×149.6	190×190×153.0
Voxel size [mm ³]	0.85×0.85×0.85	0.85×0.85×0.85	0.85×0.85×0.85	0.85×0.85×2.55
Fat suppression	Yes (Q-fat sat.)	No	No	No
Orientation	Sagittal	Sagittal	Sagittal	Sagittal (0°), 36°, 72°, 108°, 144° (*)
TR [ms]	11.40	11.40	22.50	22.50
TE [ms]	set 1: 0.06, 4.92 set 2: 0.50, 7.38	set 1: 0.06, 4.55 set 2: 0.10, 6.82	set 1: 0.04, 2.29, 6.82, 11.36, 15.90 set 2: 0.10, 4.55, 9.09, 13.64, 18.17	set 1: 0.04, 9.09 set 2: 0.10, 11.36 set 3: 2.27, 13.64 set 4: 4.55, 15.90 set 5: 6.82, 18.17
Flip angle [degrees]	6	6	6	6
Spiral interleaves (#)	212	212	212	184
Spiral duration [μ s]	1240	1240	1240	1360
Spiral iPAT factor	2	2	2	2
Reconstruction mode	SPIRiT	SPIRiT	SPIRiT	SPIRiT
Total scan time [min:s]	7:14	7:06	14:00	10:20

*See Fig. 6.C.3 for a schematic representation of the SR acquisitions.

For acquisitions A, B, and C, all images were rigidly registered to their respective first T_2^* -weighted image (TE₁) using the Advanced Mattes Mutual Information metric in Elastix (Klein et al., 2010; Wu et al., 2020). T_2^* relaxation times for these acquisitions were estimated voxel-wise using constrained non-linear least-squares fitting of a mono-exponential T_2^* relaxation model. Model-based super-resolution T_2^* mapping with joint inter-scan motion estimation was performed on the super-resolution data set using the proposed framework, using the following mono-exponential T_2^* relaxation model

$$f_n(\boldsymbol{\vartheta}_{\bullet j}) = \rho_j e^{-\frac{\mathsf{T} \mathbb{E}_n}{T_{2,j}^*}}, \qquad (6.C.1)$$


Figure 6.C.3: Schematic representation of the super-resolution (SR) T_2^* -weighted acquisitions. Five UTE Spiral VIBE data sets, consisting of two TEs each, were acquired with high in-plane and low through-plane resolution, while rotating around the frequency-encoding axis over angles of 0°, 36°, 72°, 108°, and 144°.



Figure 6.C.4: Representative T_2^* and proton density maps estimated from two short (A and B), and two long (C and SR) acquisition protocols (see Table 6.C.2).

with $\vartheta_{\bullet j} = [\rho_j, T_{2,j}^*]^T$ the tissue parameter vector at position x_j , representing the HR proton density and T_2^* relaxation value at that respective position. Regularization hyperparameters for the total variation terms were heuristically determined.

Results Fig. 6.C.4 displays representative T_2^* and PD maps from one volunteer corresponding to the 4 presented acquisitions (A, B, C and super-resolution). Acquisition A provides T_2^* maps with a noisy appearance and overall higher values than acquisition B, which displays the lowest T_2^* values overall. Acquisition C and super-resolution provide similar T_2^* maps.

Conclusion It could be demonstrated that the T_2^* maps obtained using the proposed modelbased SRR framework are comparable to maps generated with direct 3D UTE Spiral VIBE acquisitions, **while requiring approximately 25% less scan time**. SRR UTE T_2^* mapping thus shows great promise for high-resolution quantitative T_2^* mapping of knee structures within reasonable scan time.

References

- Aja-Fernández, S., Pieciak, T., & Sánchez-Ferrero, G. V. (2015). Spatially variant noise estimation in MRI: A homomorphic approach. *Med. Image Anal.*, 20(1), 184–197.
- Aja-Fernández, S., Pieciak, T., & Sánchez-Ferrero, G. V. (2016). Statistical Analysis of Noise in MRI: Modeling, Filtering and Estimation. Springer International Publishing, Cham, Switzerland, 1ed. ed.
- Aja-Fernández, S., Vegas-Sánchez-Ferrero, G., & Tristán-Vega, A. (2014). Noise estimation in parallel MRI: GRAPPA and SENSE. *Magn. Reson. Imaging*, 32(3), 281–290.
- Bano, W., Piredda, G. F., Davies, M., Marshall, I., Golbabaee, M., Meuli, R., Kober, T., Thiran, J. P., & Hilbert, T. (2020). Model-based super-resolution reconstruction of T2 maps. *Magn. Reson. Med.*, *83*(3), 906–919.
- Barral, J. K., Gudmundson, E., Stikov, N., Etezadi-Amoli, M., Stoica, P., & Nishimura, D. G. (2010). A robust methodology for in vivo T1 mapping. *Magn. Reson. Med.*, *64*(4), 1057–1067.
- Beck, A., & Tetruashvili, L. (2013). On the convergence of block coordinate descent type methods. *SIAM J. Optim.*, 23(4), 2037–2060.
- Beirinckx, Q., Ramos-Llordén, G., Jeurissen, B., Poot, D. H. J., Parizel, P. M., Verhoye, M., Sijbers, J., & den Dekker, A. J. (2020). Joint Maximum Likelihood Estimation of Motion and T1 Parameters from Magnetic Resonance Images in a Super-resolution Framework: a Simulation Study. *Fundam. Inform.*, 172, 105–128.
- Beker, K., Garces-Descovich, A., Mangosing, J., Cabral-Goncalves, I., Hallett, D., & Mortele, K. J. (2017). Optimizing MRI Logistics: Prospective Analysis of Performance, Efficiency, and Patient Throughput. Am. J. Roentgenol., 209(4), 836–844.
- Bladt, P., Beirinckx, Q., van der Plas, M. C. E., Schmid, S., Teeuwisse, W. M., Jeurissen, B., den Dekker, A. J., Sijbers, J., & van Osch, M. J. P. (2020). Super-resolution reconstruction of single-PLD pseudo-continuous ASL images. In *Proc. Intl. Soc. Mag. Reson. Med.*, vol. 28, (p. 3293).
- Bouhrara, M., Bonny, J.-M., Ashinsky, B. G., Maring, M. C., & Spencer, R. G. (2017). Noise estimation and reduction in magnetic resonance imaging using a new multispectral nonlocal maximum-likelihood filter. *IEEE Trans. Med. Imag.*, 36(1), 181–193.
- Breda, S. J., Poot, D. H. J., Papp, D., de Vries, B. A., Kotek, G., Krestin, G. P., Hernandez-Tamames, J. A., de Vos, R.-J., & Oei, E. H. G. (2020). Tissue-Specific T-2* Biomarkers in Patellar Tendinopathy by Subregional Quantification Using 3D Ultrashort Echo Time MRI. *Journal of Magnetic Resonance Imaging*, 52(2), 420–430.
- Carr, H. Y., & Purcell, E. M. (1954). Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Phys. Rev.*, 94, 630–638.
- Chambolle, A., Levine, S. E., & Lucier, B. J. (2011). An upwind finite-difference method for total variation-based image smoothing. *SIAM J. Imaging Sci.*, 4(1), 277–299.
- Chang, E. Y., Du, J., & Chung, C. B. (2015). UTE Imaging in the Musculoskeletal System. Journal of Magnetic Resonance Imaging, 41(4), 870–883.

- Christiaens, D., Cordero-Grande, L., Pietsch, M., Hutter, J., Price, A. N., Hughes, E. J., Vecchiato, K., Deprez, M., Edwards, A. D., Hajnal, J. V., & Tournier, J.-. D. (2021). Scattered slice SHARD reconstruction for motion correction in multi-shell diffusion MRI. *NeuroImage*, 225, 117437.
- Chu, C. R., & Williams, A. A. (2019). Quantitative MRI UTE-T2* and T2* Show Progressive and Continued Graft Maturation Over 2 Years in Human Patients After Anterior Cruciate Ligament Reconstruction. *Orthopaedic Journal of Sports Medicine*, 7(8).
- Coleman, T. F., & Li, Y. (1994). On the convergence of reflective Newton methods for large-scale minimization subject to bounds. *Math. Program.*, 67(2), 189–224.
- Constantinides, C. D., Atalar, E., & McVeigh, E. R. (1997). Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magn. Reson. Med.*, *38*(5), 852–857.
- Cordero-Grande, L., Teixeira, R. P. A. G., Hughes, E. J., Hutter, J., Price, A. N., & Hajnal, J. (2016). Sensitivity encoding for aligned multishot magnetic resonance reconstruction. *IEEE Trans. Comput. Imaging*, 2(3), 266–280.
- Coupé, P., Manjón, J. V., Gedamu, E., Arnold, D., Robles, M., & Collins, D. L. (2010). Robust Rician noise estimation for MR images. *Med. Image Anal.*, *14*(4), 483–493.
- de Mello, R., Ma, Y., Ji, Y., Du, J., & Chang, E. Y. (2019). Quantitative MRI musculoskeletal techniques: An update. AJR Am. J. Roentgenol., 213(3), 524–533.
- den Dekker, A. J., & Sijbers, J. (2014). Data distributions in magnetic resonance images: A review. *Physica Med.*, 30(7), 725–741.
- Ebner, M., Wang, G., Li, W., Aertsen, M., Patel, P. A., Aughwane, R., Melbourne, A., Doel, T., Dymarkowski, S., De Coppi, P., David, A. L., Deprest, J., Ourselin, S., & Vercauteren, T. (2020). An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI. *NeuroImage*, 206, 116324.
- Fessler, J. A., & Kim, D. (2011). Axial block coordinate descent (ABCD) algorithm for X-ray CT image reconstruction. In Proc. Intl. Mtg. on Fully 3D image Recon. in Rad. and Nuc. Med., (pp. 262–265).
- Fogtmann, M., Seshamani, S., Kim, K., Chapman, T., & Studholme, C. (2012). A unified approach for motion estimation and super resolution reconstruction from structural Magnetic Resonance imaging on moving subjects. In *MICCAI Workshop on Perinatal and Paediatric Imaging*, (pp. 9–26).
- Fogtmann, M., Seshamani, S., Kroenke, C., Cheng, X., Chapman, T., Wilm, J., Rousseau, F., & Studholme, C. (2014). A unified approach to diffusion direction sensitive slice registration and 3-D DTI reconstruction from moving fetal brain anatomy. *IEEE Trans. Med. Imag.*, *33*(2), 272–289.
- Fritz, J., Fritz, B., Thawait, G. G., Meyer, H., Gilson, W. D., & Raithel, E. (2016a). Three-Dimensional CAIPIRINHA SPACE TSE for 5-Minute High-Resolution MRI of the Knee. *Invest. Radiol.*, 51(10), 609–617.
- Fritz, J., Raithel, E., Thawait, G. K., Gilson, W., & Papp, D. F. (2016b). Six-Fold Acceleration of High-Spatial Resolution 3D SPACE MRI of the Knee Through Incoherent k-Space Undersampling and Iterative Reconstruction – First Experience. *Invest. Radiol.*, 51(6), 400–409.
- Gao, X., Hucker, P., Hennig, J., & Zaitsev, M. (2021). Strategies to improve intratrain prospective motion correction for turbo spin-echo sequences with constant flip angles. *Magn. Reson. Med.*, 86(2), 852–865.
- Garwood, E. R., Recht, M. P., & White, L. M. (2017). Advanced Imaging Techniques in the Knee: Benefits and Limitations of New Rapid Acquisition Strategies for Routine Knee MRI. *Am. J. Roentgenol.*, 209(3), 552–560.
- Gholipour, A., Estroff, J. A., & Warfield, S. K. (2010). Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. *IEEE Trans. Med. Imag.*, 29(10), 1739–1758.
- Greenspan, H., Oz, G., Kiryati, N., & Peled, S. (2002). MRI inter-slice reconstruction using superresolution. Magn. Reson. Imaging, 20(5), 437–446.
- Hansen, P. C., Nagy, J. G., & O'Leary, D. P. (2006). *Deblurring Images: Matrices, Spectra, and Filtering (Fundamentals of Algorithms 3) (Fundamentals of Algorithms)*. USA: Society for Industrial and Applied Mathematics.
- Hilbert, T., et al. (2018). Accelerated T2 mapping combining parallel MRI and model-based reconstruction: GRAPPATINI. J. Magn. Reson. Imaging, 48(2), 359–368.
- Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., & Hajnal, J. V. (2007). MRI of Moving Subjects Using Multislice Snapshot Images With Volume Reconstruction (SVR): Application to Fetal, Neonatal, and Adult Brain Studies. *IEEE Trans. Med. Imag.*, 26(7), 967–980.

- Kainz, B., Steinberger, M., Wein, W., Kuklisova-Murgasova, M., Malamateniou, C., Keraudren, K., Torsney-Weir, T., Rutherford, M., Aljabar, P., Hajnal, J. V., & Rueckert, D. (2015). Fast Volume Reconstruction From Motion Corrupted Stacks of 2D Slices. *IEEE Trans. Med. Imag.*, 34(9), 1901–1913.
- Keys, R. G. (1981). Cubic convolution interpolation for digital image processing. IEEE Trans. Acoust., Speech, Signal Process, (pp. 1153–1160).
- Kim, J., Lartey, R., Liu, K., Zhao, T., Benkert, T., Meyer, H., & Li, X. (2019). Comparison of radial and spiral UTE MRI and T2* quantification of the knee joint. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 27, (p. 1393).
- Klein, S., Staring, M., Murphy, K., Viergever, M. A., & Pluim, J. P. W. (2010). elastix: A Toolbox for Intensity-Based Medical Image Registration. *IEEE Transactions on Medical Imaging*, 29(1), 196–205.
- Lajous, H., Hilbert, T., Roy, C. W., Tourbier, S., de Dumast, P., Yu, T., Thiran, J.-P., J.-B., L., Piccini, D., Hagmann, P., Meuli, R., Kober, T., Stuber, M., van Heeswijk, R. B., & Bach Cuadraet, M. (2020). T2 mapping from super-resolution-reconstructed clinical fast spin echo magnetic resonance acquisitions. In A. L. Martel, et al. (Eds.) *Medical Image Computing and Computer-Assisted Intervention MICCAI 2020*, vol. 12262, (pp. 114–124). Springer, Cham.
- Larkin, K. G., Oldfield, M. A., & Klemm, H. (1997). Fast Fourier method for the accurate rotation of sampled images. *Opt. Commun.*, 139(1-3), 99–106.
- Lee, S. H., Lee, Y. H., & Suh, J.-S. (2018). Accelerating knee MR imaging: Compressed sensing in isotropic three-dimensional fast spin-echo sequence. *Magn. Reson. Imaging*, 46, 90–97.
- Lustig, M., & Pauly, J. M. (2010). SPIRIT: Iterative Self-consistent Parallel Imaging Reconstruction From Arbitrary k-Space. *Magnetic Resonance in Medicine*, 64(2), 457–471.
- Maclaren, J., Herbst, M., Speck, O., & Zaitsev, M. (2013). Prospective motion correction in brain imaging: A review. *Magn. Reson. Med.*, 69(3), 621–636.
- Maier, O., Schoormans, J., Schloegl, M., Strijkers, G. J., Lesch, A., Benkert, T., Block, T., Coolen, B. F., Bredies, K., & Stollberger, R. (2019). Rapid T1 quantification from high resolution 3D data with model-based reconstruction. *Magn. Reson. Med.*, *81*(3), 2072–2089.
- Maitra, R., & Faden, D. (2009). Noise estimation in magnitude MR datasets. *IEEE Trans. Med. Imag.*, 28(10), 1615–1622.
- Meijering, E., & Unser, M. (2003). A note on cubic convolution interpolation. IEEE Trans. Image Process., 12(4), 477–479.
- Morez, J., Szczepankiewicz, F., den Dekker, A. J., Vanhevel, F., Sijbers, J., & Jeurissen, B. (2023). Optimal experimental design and estimation for q-space trajectory imaging. *Human Brain Mapping*, 44(4), 1793–1809.
- Mugler, J. P., Fielden, S. W., Meyer, C. H., Altes, T. A., Wilson Miller, G., Stemmer, A., Pfeuffer, J., & Kiefer, B. (2015). Breath-hold ute lung imaging using a stack-of-spirals acquisition. In *Proceedings* of the International Society for Magnetic Resonance in Medicine, vol. 23, (p. 1476).
- Mugler III, J. P. (2014). Optimized three-dimensional fast-spin-echo MRI. *Journal of Magnetic Resonance Imaging*, 39(4), 745–767.
- Nachmani, A., Schurr, R., Joskowicz, L., & Mezer, A. A. (2019). The effect of motion correction interpolation on quantitative T1 mapping with MRI. *Med. Image Anal.*, *52*, 119–127.
- Naraghi, A., & White, L. M. (2012). Three-Dimensional MRI of the Musculoskeletal System. Am. J. Roentgenol., 199(3), W283–W293.
- Nicastro, M., Beirinckx, Q., Bladt, P., Jeurissen, B., Klein, S., Sijbers, J., Poot, D. H. J., & den Dekker, A. J. (2020). Optimal design of a T₁ super-resolution reconstruction experiment: a simulation study. In *Abstract Book of the 12th Annual Meeting of the ISMRM Benelux Chapter*, (pp. P–008).
- Notohamiprodjo, M., Horng, A., Pietschmann, M. F., Müller, P. E., Horger, W., Park, J., Crispin, A., Garcia del Olmo, J. R., Weckbach, S., Herrmann, K. A., Reiser, M. F., & Glaser, C. (2009). MRI of the Knee at 3T: First Clinical Results With an Isotropic PDfs-Weighted 3D-TSE-Sequence. *Invest. Radiol.*, 44(9), 585–597.
- Ogita, T., Rump, S. M., & Oishi, S. (2005). Accurate sum and dot product. SIAM Journal on Scientific Computing, 26(6), 1955–1988.
- Petrovic, A., Scheurer, E., & Stollberger, R. (2015). Closed-form solution for T2 mapping with nonideal refocusing of slice selective CPMG sequences. *Magn. Reson. Med.*, 73(2), 818–827.
- Pieciak, T., Aja-Fernández, S., & Sánchez-Ferrero, G. V. (2017). Non-stationary Rician noise estimation

in parallel MRI using a single image: A variance-stabilizing approach. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(10), 2015–2029.

- Plenge, E., Poot, D. H. J., Bernsen, M., Kotek, G., Houston, G., Wielopolski, P., van der Weerd, L., Niessen, W. J., & Meijering, E. (2012). Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magn. Reson. Med.*, 68(6), 1983–1993.
- Poot, D. H. J., den Dekker, A. J., Achten, E., Verhoye, M., & Sijbers, J. (2010a). Optimal Experimental Design for Diffusion Kurtosis Imaging. *IEEE Trans. Med. Imag.*, 29(3), 819–829.
- Poot, D. H. J., Jeurissen, B., Bastiaensen, Y., Veraart, J., Van Hecke, W., Parizel, P. M., & Sijbers, J. (2013). Super-resolution for multislice diffusion tensor imaging. *Magn. Reson. Med.*, 69(1), 103–113.
- Poot, D. H. J., Van Meir, V., & Sijbers, J. (2010b). General and efficient super-resolution method for multi-slice MRI. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, vol. 6361, (pp. 615–622).
- Qian, Y., & Boada, F. E. (2008). Acquisition-weighted stack of spirals for fast high-resolution threedimensional ultra-short echo time MR imaging. *Magnetic Resonance in Medicine*, 60(1), 135–145.
- Ramos-Llordén, G., den Dekker, A. J., Van Steenkiste, G., Jeurissen, B., Vanhevel, F., Van Audekerke, J., Verhoye, M., & Sijbers, J. (2017). A unified maximum likelihood framework for simultaneous motion and T₁ estimation in quantitative MR T₁ mapping. *IEEE Trans. Med. Imag.*, 36(2), 433–446.
- Renders, J., Jeurissen, B., Nguyen, A.-T., De Beenhouwer, J., & Sijbers, J. (2023). Imwip: Open-source image warping toolbox with adjoints and derivatives. *SoftwareX*, 24, 101524.
- Renders, J., Sijbers, J., & De Beenhouwer, J. (2021). Adjoint image warping using multivariate splines with application to four-dimensional computed tomography. *Med. Phys.*, 48(10), 6362–6374.
- Ristow, O., Steinbach, L., Sabo, G., et al. (2009). Isotropic 3D fast spin-echo imaging versus standard 2D imaging at 3.0 T of the knee image quality and diagnostic performance. *Eur. Radiol.*, *19*, 1263–1272.
- Rousseau, F., Glenn, O., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D., Barkovich, J., & Studholme, C. (2006). Registration-Based Approach for Reconstruction of High-Resolution In Utero Fetal MR Brain Images. *Acad. Radiol.*, *13*(9), 1072–1081.
- Rousseau, F., Kim, K., Studholme, C., Koob, M., & Dietemann, J. L. (2010). On Super-Resolution for Fetal Brain MRI. In *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2010*, *PT II*, vol. 6362 of *Lecture Notes in Computer Science*, (pp. 355–362).
- Scherrer, B., Gholipour, A., & Warfield, S. K. (2012). Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Med. Image Anal.*, 16(7), 1465–1476.
- Seiler, A., Nöth, U., Hok, P., Reiländer, A., Maiworm, M., Baudrexel, S., Meuth, S., Rosenow, F., Steinmetz, H., Wagner, M., Hattingen, E., Deichmann, R., & Gracien, R.-M. (2021). Multiparametric quantitative MRI in neurological diseases. *Front. Neurol.*, 12, 287.
- Shilling, R. Z., Robbie, T. Q., Bailloeul, T., Mewes, K., Mersereau, R. M., & Brummer, M. E. (2009). A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI. *IEEE Trans. Med. Imag.*, 28(5), 633–644.
- Sijbers, J., den Dekker, A. J., Scheunders, P., & Van Dyck, D. (1998). Maximum-likelihood estimation of Rician distribution parameters. *IEEE Trans. Med. Imag.*, *17*(3), 357–361.
- Smekens, C., Beirinckx, Q., Vanhevel, F., Van Dyck, P., den Dekker, A. J., Sijbers, J., Janssens, T., & Jeurissen, B. (2021). Super-resolution T2* mapping of the knee using UTE Spiral VIBE MRI. In *Proc. Intl. Soc. Mag. Reson. Med.*, vol. 29, (p. 3920).
- Smekens, C., Vanhevel, F., Jeurissen, B., Van Dyck, P., Sijbers, J., & Janssens, T. (2020). Short T2* quantification of knee structures based on accelerated UTE Spiral VIBE MRI with SPIRiT reconstruction. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, vol. 28, (p. 2687).
- Smith, S. M. (2002). Fast robust automated brain extraction. Hum. Brain Mapp., 17(3), 143–155.
- Subhas, N., Kao, A., Freire, M., Polster, J. M., Obuchowski, N. A., & Winalski, C. S. (2011). MRI of the Knee Ligaments and Menisci: Comparison of Isotropic-Resolution 3D and Conventional 2D Fast Spin-Echo Sequences at 3T. Am. J. Roentgenol., 197(2), 442–450.
- Sui, Y., Afacan, O., Gholipour, A., & Warfield, S. K. (2021). Fast and high-resolution neonatal brain MRI through super-resolution reconstruction from acquisitions with variable slice selection direction. *Front. Neurosci.*, 15, 709.

- van den Bos, A. (2007). Parameter Estimation for Scientists and Engineers. Hoboken, New Yersey, USA: John Wiley & Sons.
- Van Dyck, P., Smekens, C., Vanhevel, F., De Smet, E., Roelant, E., Sijbers, J., & Jeurissen, B. (2020). Super-Resolution Magnetic Resonance Imaging of the Knee Using 2-Dimensional Turbo Spin Echo Imaging. *Invest. Radiol.*, 55(8), 481–493.
- van Houdt, P. J., Yang, Y., & van der Heide, U. A. (2021). Quantitative magnetic resonance imaging for biological image-guided adaptive radiotherapy. *Front. in Oncol.*, *10*, 3190.
- Van Reeth, E., Tan, C. H., Tham, I. W., & Poh, C. L. (2015). Isotropic reconstruction of a 4-D MRI thoracic sequence using super-resolution. *Magn. Reson. Med.*, 73(2), 784–793.
- Van Reeth, E., Tham, I. W., Tan, C. H., & Poh, C. L. (2012). Super-resolution in magnetic resonance imaging: A review. *Concepts Magn. Reson.*, 40A(6), 306–325.
- Van Steenkiste, G., Jeurissen, B., Veraart, J., den Dekker, A. J., Parizel, P. M., Poot, D. H. J., & Sijbers, J. (2016). Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations. *Magn. Reson. Med.*, 75(1), 181–195.
- Van Steenkiste, G., Poot, D. H. J., Jeurissen, B., den Dekker, A. J., Vanhevel, F., Parizel, P. M., & Sijbers, J. (2017). Super-resolution T1 estimation: Quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. *Magn. Reson. Med.*, 77(5), 1818–1830.
- Walsh, D. O., Gmitro, A. F., & Marcellin, M. W. (2000). Adaptive reconstruction of phased array MR imagery. Magn. Reson. Med., 43(5), 682–690.
- Williams, A., Qian, Y., Bear, D., & Chu, C. R. (2010). Assessing degeneration of human articular cartilage with ultra-short echo time (UTE) T-2* mapping. Osteoarthritis and Cartilage, 18(4), 539–546.
- Wu, M., Zhao, W., Wan, L., Kakos, L., Li, L., Jerban, S., Jang, H., Chang, E. Y., Du, J., & Ma, Y.-J. (2020). Quantitative three-dimensional ultrashort echo time cones imaging of the knee joint with motion correction. *NMR in Biomedicine*, 33(1).
- Yao, L., Pitts, J. T., & Thomasson, D. (2007). Isotropic 3D Fast Spin-Echo with Proton-Density-Like Contrast: A Comprehensive Approach to Musculoskeletal MRI. Am. J. Roentgenol., 188(2), W199–W201.
- Zhang, T., Pauly, J. M., & Levesque, I. R. (2015). Accelerating parameter mapping with a locally low rank constraint. *Magn. Reson. Med.*, 73(2), 655–661.
- Zhao, B., Haldar, J. P., Christodoulou, A. G., & Liang, Z. (2012). Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints. *IEEE Trans. Med. Imag.*, 31(9), 1809–1820.
- Zhao, B., Haldar, J. P., Liao, C., Ma, D., Jiang, Y., Griswold, M. A., Setsompop, K., & Wald, L. L. (2019). Optimal Experiment Design for Magnetic Resonance Fingerprinting: Cramér-Rao Bound Meets Spin Dynamics. *IEEE Trans. Med. Imag.*, 38(3), 844–861.

7

A super-resolution reconstruction framework for quantitative brain perfusion mapping using pseudo-continuous Arterial Spin Labeling

Contents

7.1 Introduction	03
7.2 Theory	04
7.3 Methods	08
7.4 Results	18
7.5 Discussion	36
7.6 Conclusion	43
Appendices	43
7.A Slice-dependent PLD	43
7.B Background suppression	44
7.C Linear forward model	45
7.D Extra simulation experiment	47
7.E Motion parameter estimates	47
References	50

ABSTRACT

Arterial spin labeling (ASL) is a promising, non-invasive perfusion magnetic resonance imaging technique for quantifying cerebral blood flow (CBF). Unfortunately, ASL suffers from an inherently low signal-to-noise ratio (SNR) and spatial resolution, undermining its potential. Increasing spatial resolution without significantly sacrificing SNR or scan time represents a critical challenge towards routine clinical use. In this contribution, we propose a model-based super-resolution reconstruction (SRR) method with joint motion estimation that breaks the traditional SNR/resolution/scan-time trade-off. From a set of differently oriented 2D multi-slice pseudo-continuous ASL images with a low through-plane resolution, 3D-isotropic, high resolution, quantitative CBF maps are estimated using a Bayesian approach. Experiments on both synthetic whole brain phantom data, and on *in vivo* brain data, show that the proposed SRR Bayesian estimation framework outperforms state-of-the-art ASL quantification.

The work in this chapter has been published as:

Beirinckx, Q., Bladt, P., van der Plas, M. C. E., van Osch, M. J. P., Jeurissen, B., den Dekker, A. J., Sijbers, J., "Model-based super-resolution reconstruction for pseudocontinuous Arterial Spin Labeling", <u>NeuroImage</u>, Vol. 286, pp. 120606, 2024. doi: 10.1016/j.neuroimage.2024.120506

7.1 Introduction

Arterial spin labeling (ASL) is a magnetic resonance (MR) imaging technique to noninvasively measure cerebral blood flow (CBF), which is a biomarker for various brain disorders (Alsop et al., 2015; van Osch et al., 2018). ASL uses magnetically labeled arterial blood as an endogenous tracer, where the labeling is performed by inverting the inflowing blood magnetization in a plane proximal to the brain. After a specific period of time, called the post-labeling delay (PLD) time, during which labeled blood travels through the arterial vascular tree towards the brain tissue, a so-called label image is acquired. Additionally, a control image is acquired without prior labeling. The difference between the label and control image yields a perfusion weighted image that isolates the ASL signal. Next, a CBF map is computed from the perfusion weighted image using a perfusion model and a separately acquired calibration image (Alsop et al., 2015).

The consensus paper by Alsop et al. recommends pseudo-continuous ASL (pCASL), background suppression (BS), and segmented 3D readout for clinical implementation of ASL (Alsop et al., 2015). Segmented 3D acquisition schemes use a single excitation per TR, which is optimal for BS (Ye et al., 2000; Krüger & Glover, 2001; Garcia et al., 2005; Maleki et al., 2012; Paschoal et al., 2021). However, 3D readout sequences employ long echo trains, resulting in through-plane blurring due to T_2 decay along the echo train. Splitting the readout into more segments can reduce this blurring, but at the cost of a longer acquisition time and increased sensitivity to inter-shot motion and physiological fluctuations (Hernandez-Garcia et al., 2022). In addition, the long readout time of a 3D imaging sequence holds an increased risk of motion artefacts (Alsop et al., 2015).

As a viable alternative to 3D readout, single-shot 2D multi-slice (MS) readout methods based on echo-planar imaging (EPI) have been suggested (Alsop et al., 2015). 2D readout methods have several advantages over 3D readout methods. First, they have a much shorter readout and are hence less susceptible to motion during readout (Vidorreta et al., 2013, 2014; Alsop et al., 2015). Second, they are less susceptible to spatial blurring due to T_2 decay (Vidorreta et al., 2013). Third, they can be used at high field strengths where power deposition limits prohibit the use of multiple refocusing pulses (Hernandez-Garcia et al., 2022). Finally, 2D readout methods are readily available on all systems (Alsop et al., 2015).

However, 2D readout approaches also come with disadvantages. First, 2D MS imaging causes the PLD time to increase in subsequently acquired slices, which results in a significant degradation of the signal-to-noise ratio (SNR) in the last acquired slices due to longitudinal relaxation (van Osch et al., 2018). At the same time, however, the slice-wise increase of PLD can help accommodate unbiased CBF estimation in subjects with arterial transit time (ATT) values that increase from inferior to superior slices, which can be considered a consistent finding in most subjects. Second, the use of a separate excitation pulse for every slice complicates BS. In practice, BS can be optimal for only one slice and will be progressively less efficient for other slices (Alsop et al., 2015).

In this contribution, which is based on a preliminary study (Bladt et al., 2020), we propose an alternative 2D MS based image acquisition and parameter estimation method for single-PLD pCASL that alleviates the main disadvantages of traditional 2D MS imaging, while preserving its advantages. The method relies on MS super-resolution reconstruction (SRR). In this approach, a 3D isotropic high resolution (HR) image or parameter map is estimated from multiple, differently oriented, 2D MS images with a low through-plane resolution. SRR

has been shown to improve the inherent trade-off between spatial resolution, SNR, and acquisition time in MRI (Van Reeth et al., 2012; Plenge et al., 2012) and has previously been applied successfully in anatomical imaging (Poot et al., 2010; Van Dyck et al., 2020), diffusion MRI (Poot et al., 2013; Fogtmann et al., 2014; Van Steenkiste et al., 2016), and relaxometry (Van Steenkiste et al., 2017; Bano et al., 2020; Beirinckx et al., 2020, 2022). The current contribution introduces SRR in the field of ASL, proposing a model-based MS-SRR framework with joint motion estimation for direct whole brain CBF mapping from 2D MS single-PLD pCASL data. By choosing an SRR acquisition scheme in which low resolution (LR) label-control image pairs are acquired with varying slice-encoding directions, the negative effects of fading BS and increasing PLD values in subsequently acquired slices of the traditional 2D MS readout scheme for pCASL are made independent of location, i.e. averaged out. To explore its potential in ASL, our newly proposed method is evaluated on synthetic whole brain perfusion data. Moreover, our pCASL MS-SRR method is combined with multiband (MB) imaging, also known as simultaneous multi-slice (SMS) imaging, to accelerate image acquisition and hence provide a more constant and thus on average better BS as well as a more constant PLD across slices (van Osch et al., 2018). Finally, our method is validated on in vivo brain data, and compared to a conventional single-PLD pCASL experiment with 2D MS readout using a widely used Bayesian inference model (BASIL (Chappell et al., 2009; Groves et al., 2009)) for CBF quantification.

7.2 Theory

In what follows, the components of the proposed ASL SRR framework are discussed, namely the SRR forward model (Section 7.2.1), the single-PLD pCASL signal model, which encapsulates the CBF quantification formula (Section 7.2.2), and the joint Bayesian estimation framework, for direct CBF mapping with joint motion estimation from LR single-PLD pCASL data (Section 7.2.3).

7.2.1 Super-resolution reconstruction forward model

Let $s = \{s_n\}_{n=1}^{2N}$ be the set of N vectorized, noiseless, anisotropic LR 2D MS control (odd n) and N label (even n) magnitude images. Each image $s_n = \{s_{nl}\}_{l=1}^{N_s} \in \mathbb{R}^{N_s \times 1}$ is sampled at the LR grid points $y_n = \{y_{nl}\}_{l=1}^{N_s} \in \mathbb{R}^{3 \times N_s}$ with N_s the number of voxels per LR image, and can be modelled as:

$$s_n = DBG_n M_{\theta_n} r_n \tag{7.2.1}$$

where $\mathbf{r}_n = \{r_{nj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ represents the unknown, noiseless HR image with the same perfusion-weighting as \mathbf{s}_n and defined at the targeted isotropic HR grid points $\mathbf{x} = \{\mathbf{x}_j\}_{j=1}^{N_r} \in \mathbb{R}^{3 \times N_r}$, with N_r the number of voxels of the HR image. Furthermore, $\mathbf{M}_{\theta_n} \in \mathbb{R}^{N_r \times N_r}$, $\mathbf{G}_n \in \mathbb{R}^{N_r \times N_r}$, $\mathbf{B} \in \mathbb{R}^{N_r \times N_r}$, and $\mathbf{D} \in \mathbb{R}^{N_s \times N_r}$ are linear operators that describe unintended motion, a known geometric transformation that models the image acquisition with specific slice orientation, spatially invariant blurring, and down-sampling, respectively. The motion operator \mathbf{M}_{θ_n} is modeled as a parametric function of θ_n . Assuming rigid inter-image motion, the parameter vector $\theta_n \in \mathbb{R}^{6 \times 1}$ is given by $\theta_n = [t_{xn}, t_{yn}, t_{zn}, \alpha_n, \beta_n, \gamma_n]^T$, with t_{xn}, t_{yn}, t_{zn} the translation parameters and α_n , β_n , γ_n the Euler angles that describe rotation around the x, y, and z axis, respectively.

7.2.2 Single-PLD pCASL signal model

Let $\Delta \mathbf{r}_n = \{\Delta r_{nj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$, with *n* even, be the difference image $\mathbf{r}_{n-1} - \mathbf{r}_n$ and let $\vartheta_{rCBF} = \{\vartheta_{rCBF,j}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ represent the HR *relative* CBF parameter map to be estimated, expressed in arbitrary units (a.u.). According to the recommended quantification formula for single-PLD pCASL data (Alsop et al., 2015), Δr_{nj} is given by:

$$\Delta r_{nj}(\vartheta_{\mathsf{rCBF},j}) = \vartheta_{\mathsf{rCBF},j} \delta^{-1} \exp\left(-\frac{\mathsf{PLD}_{nj}}{\mathcal{T}_{1b}}\right) \quad , \tag{7.2.2}$$

with

$$\delta = 6000 \cdot \frac{\lambda}{2\alpha T_{1b} \left(1 - \exp\left(-\frac{\tau}{T_{1b}}\right)\right)}$$
(7.2.3)

a scalar constant that encapsulates the labeling efficiency α , the brain-blood partition coefficient λ , the labeling duration τ , and the longitudinal relaxation time of blood T_{1b} , which are all assumed to be known from experiment or fixed at their recommended population averages. In Eq. (7.2.2), PLD_{nj} is the PLD time that corresponds with the readout time of the corresponding slice within the label image s_n that contains the HR grid point x_j . Indeed, each slice of s_n is characterized by a unique PLD that depends on the slice acquisition order. If the MS acquisition proceeds in ascending slice order with a slice readout time t_{read} , the effective PLD in the M^{th} slice is given by: $(M - 1) \times t_{\text{read}} + \text{PLD}_{\text{base}}$, with PLD_{base} the time between the end of the labeling pulse train and the readout. In contrast to a conventional pCASL MS acquisition scheme, in a rotated SRR acquisition scheme, PLD_{nj} will depend on the slice direction (Fig. 7.2.1(d)). The mathematical model that describes the slice-dependent PLD_{nj} is provided in Appendix 7.A. Note that due to the PLD variations, the virtual HR label images (i.e., r_n with n is even) differ from each other, whereas the virtual HR control images (i.e., r_n can be modelled as:

$$r_{nj} = \begin{cases} r_{1,j}, & \text{if } n \text{ is odd} \\ r_{1,j} - \Delta r_{nj}(\vartheta_{\mathsf{rCBF},j}), & \text{if } n \text{ is even} . \end{cases}$$
(7.2.4)

Eq. (7.2.4) can be extended to include the effect of BS in 2D MS readout:

$$r_{nj} = \begin{cases} r_{1,j}b_{nj}, & \text{if } n \text{ is odd} \\ r_{1,j}b_{nj} - \Delta r_{nj}(\vartheta_{\mathsf{rCBF},j}), & \text{if } n \text{ is even}, \end{cases}$$
(7.2.5)

where $\mathbf{b}_n = \{b_{nj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ models inversion-recovery nulling for BS, under the assumption that BS is perfect for the first acquired slice and with T_{1t} the T_1 relaxation time of tissue t, i.e. $b_{nj} = 1 - 2 \cdot \exp(-\mathsf{TI}_{nj}/T_{1t,j})$, with TI_{nj} the optimal inversion time for perfect BS of the first slice. More details on how to model TI_{nj} for SRR are provided in Appendix 7.B.

Following the recommendations of the ASL white paper (Alsop et al., 2015), a calibration step is needed to translate CBF values in arbitrary units to absolute units of mL/100g/min of tissue, by voxel-wise dividing the relative CBF map ϑ_{rCBF} by a HR proton density weighted calibration image $\rho_{reg} = \{\rho_{reg,j}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$, registered to the HR reconstruction grid. In this contribution, we assume ρ_{reg} to be a known image, acquired from a separate acquisition. This calibration image is essentially a control image without background suppression acquired



Figure 7.2.1: MS SRR acquisition scheme: (a) acquisition coordinate system; (b) *k*-space coverage; (c) grid in image space; (d) slice dependent PLDs; (e) coronal LR MS images. The x-, y-, and z-direction represent the frequency-, phase-, and slice-encoding direction, respectively.

with the same readout as the ASL data (Clement et al., 2022). As such, the HR CBF parameter map is defined as $\vartheta_{CBF} = \{\vartheta_{CBF,j}\}_{j=1}^{N_r} = \{\vartheta_{rCBF,j}/\rho_{reg,j}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$. Note that by replacing ϑ_{rCBF} with ϑ_{CBF} in Eq. (7.2.4), a calibrated version of the single-PLD pCASL model for r_n can be obtained.

7.2.3 Joint Bayesian estimation framework

7.2.3.1 Maximum a posteriori estimator

Let $\tilde{s} = {\tilde{s}_n}_{n=1}^{2N} \in \mathbb{R}^{N_s \times 2N}$ denote the set of 2N measured, noisy LR MS images with $\tilde{s}_n = {\tilde{s}_n}_{l=1}^{N_s} \in \mathbb{R}^{N_s \times 1}$. Furthermore, let $\vartheta = [r_1^T \ \vartheta_{\mathsf{rCBF}}^T]^T \in \mathbb{R}^{2N_r \times 1}$ and $\theta = {\theta_n}_{n=1}^{2N} \in \mathbb{R}^{6 \times 2N}$ represent the perfusion parameters and the motion parameters to be estimated, respectively. Following a Bayesian approach, the data \tilde{s} and the parameters ${\vartheta, \theta}$ are modeled as random variables, where Bayes' theorem gives an expression for the *posterior* distribution of the parameters given the data:

$$p(\boldsymbol{\vartheta}, \boldsymbol{\theta} | \tilde{\boldsymbol{s}}) = \frac{p(\tilde{\boldsymbol{s}} | \boldsymbol{\vartheta}, \boldsymbol{\theta}) p(\boldsymbol{\vartheta}) p(\boldsymbol{\theta})}{p(\tilde{\boldsymbol{s}})} \quad , \tag{7.2.6}$$

with $p(\tilde{s}|\vartheta, \theta)$ the conditional probability distribution of \tilde{s} given the parameters $\{\vartheta, \theta\}$, $p(\vartheta)$ and $p(\theta)$ the *prior* distributions that encapsulate the prior knowledge about ϑ and θ , respectively, and with $p(\tilde{s})$ a normalization factor. When $p(\tilde{s}|\vartheta, \theta)$ is viewed as a function of the unknown parameters $\{\vartheta, \theta\}$ given the data \tilde{s} , it is called the *likelihood* function.

For a single coil acquisition, the noisy voxel intensities \tilde{s}_{nl} can be modeled as Rician distributed random variables, while for a multicoil acquisition, \tilde{s}_{nl} are governed by a non-central chi distribution (den Dekker & Sijbers, 2014). When the SNR is high enough (> 3), which is typically the case for the low resolution voxels \tilde{s}_{nl} , both distributions can be well approximated by a Gaussian distribution. If the voxel intensities are additionally assumed to be statistically independent and the standard deviation of the noise σ to be temporally and spatially invariant, the likelihood function $p(\tilde{s}|\vartheta, \theta)$ can be expressed as:

$$p(\tilde{\boldsymbol{s}}|\boldsymbol{\vartheta},\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2\sigma^2} \sum_{n=1}^{2N} \sum_{l=1}^{N_s} (\tilde{s}_{nl} - s_{nl}(\boldsymbol{\vartheta},\boldsymbol{\theta}_n))^2\right).$$
(7.2.7)

Furthermore, the prior distributions of the HR parameter maps r_1 and ϑ_{rCBF} , which are assumed to be smooth, are chosen as:

$$p(\mathbf{r}_1) \propto \exp\left(-\frac{\lambda_{\mathbf{r}_1}}{2} \|\Delta(\mathbf{r}_1)\|_2^2\right)$$
 and $p(\boldsymbol{\vartheta}_{\mathsf{rCBF}}) \propto \exp\left(-\frac{\lambda_{\boldsymbol{\vartheta}_{\mathsf{rCBF}}}}{2} \|\Delta(\boldsymbol{\vartheta}_{\mathsf{rCBF}})\|_2^2\right)$, (7.2.8)

respectively, where $\Delta(\cdot)$ denotes the 3D discrete Laplace operator (Poot et al., 2013), and $\lambda_{r_1} > 0$ and $\lambda_{\vartheta_{rCBF}} > 0$ are hyper-parameters that control the regularization strengths. For the motion parameters θ , a non-informative prior $p(\theta)$ is adopted, assuming $p(\theta)$ to be uniform over the range of values for which the likelihood function is non-negligible. The maximum a posteriori (MAP) estimator then maximizes $p(\vartheta, \theta|\tilde{s})$ w.r.t. the parameters $\{\vartheta, \theta\}$. Hence, by combining Eqs. (7.2.6)-(7.2.8), we obtain:

$$\{\hat{\boldsymbol{\vartheta}}, \hat{\boldsymbol{\theta}}\} = \arg\max_{\boldsymbol{\vartheta}, \boldsymbol{\theta}} p(\boldsymbol{\vartheta}, \boldsymbol{\theta} | \tilde{\boldsymbol{s}}) = \arg\min_{\boldsymbol{\vartheta}, \boldsymbol{\theta}} \left[-\ln p(\boldsymbol{\vartheta}, \boldsymbol{\theta} | \tilde{\boldsymbol{s}}) \right]$$
(7.2.9)
$$= \arg\min_{\boldsymbol{\vartheta}, \boldsymbol{\theta}} \left[\sum_{n=1}^{2N} \sum_{l=1}^{N_{s}} \left(\tilde{\boldsymbol{s}}_{nl} - \boldsymbol{s}_{nl} \left(\boldsymbol{\vartheta}, \boldsymbol{\theta}_{n} \right) \right)^{2} + \lambda_{\boldsymbol{r}_{1}}^{\prime} \| \Delta(\boldsymbol{r}_{1}) \|_{2}^{2} + \lambda_{\boldsymbol{\vartheta}_{rCBF}}^{\prime} \| \Delta(\boldsymbol{\vartheta}_{rCBF}) \|_{2}^{2} \right],$$
(7.2.10)

with $\lambda'_{r_1} = \sigma^2 \lambda_{r_1}$ and $\lambda'_{\vartheta_{rCBF}} = \sigma^2 \lambda_{\vartheta_{rCBF}}$ regularization parameters to be selected by the user. Note that σ does not have to be known or estimated in advance.

7.2.3.2 Optimization

The optimization problem in Eq. (7.2.9) is solved using the *alternating minimization* method, also known as the cyclic block-coordinate descent (cBCD) method (Fessler & Kim, 2011; Beck & Tetruashvili, 2013). In this method, the parameters $\{\vartheta, \theta\}$ are split into two blocks that contain the perfusion parameters ϑ , and the motion parameters θ , and the cost function is successively minimized with respect to each block in a cyclic order:

$$\hat{\boldsymbol{\vartheta}}^{(t+1)} = \arg\min_{\boldsymbol{\vartheta}} \left[\sum_{n=1}^{2N} \sum_{l=1}^{N_s} \left(\tilde{s}_{nl} - s_{nl}(\boldsymbol{\vartheta}, \hat{\boldsymbol{\theta}}_n^{(t)}) \right)^2 + \lambda_{\boldsymbol{r}_1}' \|\Delta(\boldsymbol{r}_1)\|_2^2 + \lambda_{\boldsymbol{\vartheta}_{\mathsf{rCBF}}}' \|\Delta(\boldsymbol{\vartheta}_{\mathsf{rCBF}})\|_2^2 \right]$$
(P.1)

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \arg\min_{\boldsymbol{\theta}} \sum_{n=1}^{2N} \sum_{l=1}^{N_s} \left(\tilde{s}_{nl} - s_{nl} (\hat{\boldsymbol{\vartheta}}^{(t+1)}, \boldsymbol{\theta}_n) \right)^2$$
(P.2)

with $\hat{\vartheta}^{(0)} = \vartheta_{\text{ini}}$ the initial values of the HR parameter maps ϑ , and with $\hat{\theta}^{(0)} = \theta_{\text{ini}}$ the initial values of the motion parameters θ , respectively. The procedure is terminated when

a maximum number of iterations is exceeded, or when a convergence tolerance on the relative difference of the tissue parameter estimates between consecutive iterations, defined as $\mathcal{E}^{(t)} = \|\hat{\vartheta}^{(t)} - \hat{\vartheta}^{(t-1)}\|_2 / \|\hat{\vartheta}^{(t)}\|_2$, is reached.

The alternating optimization routine requires suitable choices of the convergence tolerances and regularization weights, as well as choosing suitable solvers for model parameter optimization problem (P.1), and the motion parameter optimization problem (P.2). To efficiently solve the linear subproblem (P.1), the Conjugate Gradient Least Squares (CGLS) algorithm was used, in which parameter maps were initialized with zeros. The inter-image motion estimation problem (P.2), on the other hand, is nonlinear and adopts a particularly simple structure when the signal model parameters remain fixed. If the elements of θ are independent, problem (P.2) can be decoupled into 2*N* optimization problems that can be solved efficiently by parallel operations. Each of these decoupled problems is minimized using a trust-region Newton algorithm (Coleman & Li, 1994), with analytical expressions for the Jacobian to speed up convergence.

7.2.3.3 Implementation

The proposed method was written in MATLAB and partially in C++. Computations were performed on a computer with an Intel[®] CoreTM i7-6850K hexa-core CPU, with 32 GB of RAM, and a single NVIDIA GeForce GTX 1080 GPU. The computational complexity of the proposed algorithm is primarily defined by the Fast Fourier Transform (FFT)-based image warping operators M_{θ_n} and G_n in Eq. (7.2.1) (Beirinckx et al., 2022). To speed up reconstruction, the FFTs of these image warping operators were executed on the GPU. Furthermore, while the forward model given by Eq. (7.2.1) treats M_{θ_n} and G_n as separate operators, in our implementation we combined both operators to limit the number of FFTs and to maximize computational efficiency. Linear operators D and B followed the implementation of Beirinckx et al. (2022). MATLAB parallel computing tools were used to estimate θ_n for each value of n separately when solving problem (P.2) of the alternating minimization method. A single reconstruction took approximately 19 minutes for a simulated LR single-PLD pCASL dataset (without motion optimization), and 1 hour 10 minutes for the *in vivo* LR single-PLD pCASL dataset, respectively.

7.3 Methods

The proposed method, denoted as **SRR-pCASL**, was evaluated in simulation and *in-vivo* experiments, where its performance was compared to that of the following reference methods:

- **C-pCASL** Conventional acquisition of single-PLD pCASL data with 2D MS readout in which each control-label image pair is acquired multiple times at a 3D isotropic high resolution, with an inferior-superior slice-encoding direction, and with an ascending slice readout order. The reconstruction and direct CBF mapping are performed using the *same* joint Bayesian estimation framework as for SRR-pCASL.
- **BASIL** Conventional acquisition of single-PLD pCASL data with 2D MS readout, similar as for C-pCASL. From these data, CBF was quantified using the Bayesian Inference for Arterial Spin Labeling (BASIL) method (Chappell et al., 2009; Groves et al., 2009), which is part of the FSL toolbox (Smith et al., 2004; Woolrich et al., 2009).

Default settings were used to process single-PLD data, as described by the BASIL documentation guide and following the recommendations¹ of the consensus paper (Alsop et al., 2015). BASIL uses FSL's *mcflirt* (Jenkinson et al., 2002) to correct for motion between the ASL data and the calibration image. Note that this second reference method is primarily included to verify C-pCASL as a benchmark for optimal traditional CBF quantification w.r.t. SRR-pCASL. As such, a true one-to-one benchmarking between BASIL and the proposed MAP estimation framework is not the main objective, especially since the use of different prior information and motion correction strategies in both approaches complicates a fair comparison.

In addition to the above methods a multiband (MB) imaging version to **SRR-pCASL**, **C-pCASL** and **BASIL** was evaluated to partially prevent longitudinal relaxation effects due to increasing PLD values for ascending slices during acquisition. The corresponding MB augmented methods are denoted as **SRR-pCASL-MB**, **C-pCASL-MB**, and **BASIL-MB**, respectively.

7.3.1 Simulation experiments

Simulation experiments were set up to evaluate the proposed SRR method for single-PLD pCASL and compare its performance with that of the reference methods. First, to exclude a possible bias in the CBF estimation introduced by misregistration when comparing the different methods, a Monte Carlo simulation experiment was performed where the motion parameters θ_n were set to **0** and no motion correction was performed. Second, a Monte Carlo simulation experiment was performed in which the synthetic pCASL data were corrupted with unwanted inter-image motion to evaluate the estimation of both CBF and motion parameters. To guarantee realistic head movement, the inter-image motion parameters $\{\theta_n\}_{n=1}^{2N}$ were chosen equal to an estimated set of motion parameters obtained using a rigid registration routine² on the *in vivo* LR SRR data. The obtained true reference motion parameters for each of the 2N pCASL images in the simulation study are summarized in Figs. 7.E.1-7.E.2.

¹The compatibility with these recommendations was checked by setting BASIL's 'white paper mode' option to "ON", i.e. the arterial transit time was set to 0, both T_1 and T_{1b} were set to 1.65 seconds, the inversion efficiency was set equal to 0.9 for pCASL, and calibration with a provided proton density weighted image was performed voxel-wise. In addition, following the default recommendations, BASIL's adaptive non-local spatial smoothing prior was used (Groves et al., 2009). This spatial smoothing prior is used for CBF and is directly based on evidence in the data. It exploits the fact that neighboring voxels are likely to have similar CBF values, i.e. CBF variation in the brain is relatively smooth. It is also adaptive, so that in regions where the data does not support the use of smoothing the CBF image will not be smoothed. Motion correction, which uses FSL's mcflirt (Jenkinson et al., 2002) to estimate the motion between the ASL data (and the calibration image), was only turned on for processing of the real data. Finally, the arterial (macro-vascular) contribution flag was set to "OFF" in BASIL to facilitate comparability to the proposed method which currently implements the pCASL model omitting the local arterial contribution.

²The reference motion set θ_n was obtained from the *in vivo* LR SRR data using a procedure that involved three repetitions of: (i) upsampling of the LR SRR data by applying the adjoint operator $A_n^T = M_{\theta_n}^T G_n^T B^T D^T$ to each LR control and label image, (ii) calculation of the average HR control and rCBF maps from this upsampled data using the recommended quantification formula in Eq. (7.2.2) and averaging over the number of control-label pairs *N*, and (iii) motion estimation using subproblem (P.2) in which the HR control and rCBF map remained fixed. The motion parameters that resulted from this procedure were then used as reference motion component values.

7.3.1.1 Synthetic data generation

Both for the simulations without and with motion, four different synthetic datasets were generated, all having the same underlying HR ground truth parameter maps for CBF, PD and the relaxation time T_{1t} of tissue. These ground-truth parameter maps were generated starting from a 216 \times 180 \times 180 HR brain phantom with labeled tissue classes supplied by MRiLab (Liu et al., 2017) with a 1 mm³ isotropic resolution. Gray (GM) and white matter (WM) CBF values of 65 mL/100g/min and 20 mL/100g/min, respectively, reported for the healthy human brain, were assigned to the CBF map (Parkes et al., 2004; Zhang et al., 2014; Fan et al., 2016). To assess the identification of hyperintensities (Maier et al., 2021), we additionally simulated a hyperperfusion lesion of 113.75 mL/100g/min in GM and 50 mL/100g/min in WM, having a volume equal to 408 mm³ and 330 mm³, respectively. Subsequently, the CBF, PD, and T_{1t} maps were resampled onto a $72 \times 60 \times 60$ grid using cubic interpolation with a scale-variant kernel to prevent aliasing, matching a 3D isotropic resolution of 3 mm typical for HR ASL data. Finally, each HR ground-truth parameter map was zero padded to a $80 \times 80 \times 64$ grid, such that it corresponds to the dimensions of the reconstruction grid of the real data experiment (section 7.3.2). Starting from these ground truth parameter maps, the following noiseless datasets were generated:





- Dataset 1: LR 2D MS data (for SRR) Whole brain SRR single-PLD pCASL data was simulated assuming the rotational acquisition scheme depicted in Fig. 7.2.1. The acquisition settings, shown in Table 7.3.1, were chosen equal to those of the in vivo SRR experiment described in section 7.3.2. N = 24 control-label image pairs, each with a unique slice-encoding direction, were simulated by rotating the slice stack around the virtual phase encoding axis, aligned in the anterior-posterior direction, in increments of 180/N degrees. The acquisition settings, which include a labeling duration $\tau = 1.8$ s, a time between the end of labeling and the start of readout of the first slice $PLD_{base} = 1.8$ s, $N_{\text{slice}} = 16$ slices with a thickness of 12 mm, an in-plane isotropic resolution of 3 mm, and a readout time per slice t_{read} = 50 ms, correspond to a total scan time $T = 2N \cdot (\tau + \mathsf{PLD}_{\mathsf{base}} + N_{\mathsf{slice}} \cdot t_{\mathsf{read}}) \approx 210s.$ A schematic representation of the pCASL timing diagram is illustrated in Fig. 7.3.1. The LR control-label image pairs were simulated as follows. Starting from the $3 \times 3 \times 3$ mm³ HR ground truth CBF, PD, and T_{1t} maps described above, N = 24 HR whole-brain control-label image pairs were generated using Eq. (7.2.5), each with a unique PLD map. Next, for each HR control-label pair, a $3 \times 3 \times 12$ mm³ LR version was computed using the SRR forward model (Eq. (7.2.1)).
- **Dataset 2: HR 2D MS data** Whole brain single-PLD data was simulated assuming a 2D MS acquisition with an isotropic spatial resolution of $3 \times 3 \times 3$ mm³. The acquisition settings, which are tabulated in Table 7.3.1, were chosen identical to those of the *in*

vivo experiment described in section 7.3.2, except for the use of MB imaging, which was ignored in this dataset. Assuming $N_{\text{slice}} = 40$ slices with a thickness of 3 mm, $\tau = 1.8$ s, PLD_{base} = 1.8 s, and $t_{\text{read}} = 60$ ms, N = 22 HR control-label image pairs were simulated, leading to a longer total scan time than the *in vivo* experiment (see also dataset 4). Starting from the HR ground truth CBF, PD, and T_{1t} maps, Eq. (7.2.5) was used to generate a HR control-label image pair where the PLD increased along the inferior-superior slice-encoding direction, following the recommended ascending slice readout order (Alsop et al., 2015). Subsequently, the HR image pair was blurred by applying a spatially invariant 3D point spread function that corresponds with an MS acquisition at 3D isotropic spatial resolution.

- **Dataset 3: LR 2D MS MB data (for SRR)** To generate dataset 3, the procedure used to generate dataset 1 was repeated assuming MB imaging with an MB factor equal to 2. That is, an MB version of dataset 1 was generated assuming that slice *n* and slice $n + N_{\text{slice}}/2$, with $n = 1, 2, ..., N_{\text{slice}}/2$, were acquired at the same time, hence with the same PLD. The simultaneous acquisition of two slices with MB would reduce the scan time of dataset 3 compared to dataset 1, although the labeling duration and the PLD would still account for most of the scan time. The main advantage of MB will be a closer to optimal BS over the whole volume.
- **Dataset 4: HR 2D MS MB data** To generate dataset 4, the procedure used to generate dataset 2 was repeated assuming MB imaging with an MB factor equal to 2. As such, the acquisition settings were identical to those of the *in vivo* experiment described in Table 7.3.2, resulting in the same total scan time as dataset 1 and the *in vivo* scan to allow a fair comparison.

Finally, noise was added to the generated datasets. To facilitate an extensive Monte Carlo study, for each dataset, $N_{MC} = 100$ noise realisations were generated by adding zero-mean, Gaussian distributed noise with standard deviation $\sigma = \sqrt{\sigma_0^2 + \sigma_P^2}$, with σ_0 the standard deviation of the raw noise component, including thermal noise and scanner noise, and σ_P the standard deviation of the physiological noise component (Krüger & Glover, 2001). Unlike σ_0 , σ_P is proportional to the signal strength S, i.e., $\sigma_P = cS$, with c a scaling factor. Values for σ_0 and c in each dataset were chosen to match the temporal SNR (tSNR) values observed in the *in vivo* data. To this end, a voxel-wise tSNR map was calculated from the conventional HR 2D EPI data set (see Table 7.3.2), where the tSNR was defined per voxel as $\mu_{\Delta f_{ai}}/\sigma_{\Delta r_{ai}}$, with $\mu_{\Delta r_{nj}}$ the temporal voxel-wise mean and $\sigma_{\Delta r_{nj}}$ the temporal voxel-wise standard deviation of the difference images $\{\Delta r_n\}_{n=1}^N$, obtained from the N single-PLD pCASL label-control repetitions. Furthermore, an overall tSNR value was obtained by calculating the spatial mean inside a brain mask of the voxel-wise tSNR map. This procedure resulted in a tSNR ranging from approximately 0.2, in brain regions with almost no BS, to 3.4, in brain regions with perfect BS. Subsequently, values for σ_0 and c were tuned to match those tSNR values in the simulated datasets. Fig. 7.3.2 shows the voxel-wise tSNR map obtained from the conventional HR in vivo data alongside the tSNR maps used in the simulation experiment, as well as a comparison in overall tSNR value. Note that for the LR control-label images of dataset 1 and 3, the tSNR increased approximately 4-fold as a result of the increased slice thickness of those images when using the SRR forward model (Eq. (7.2.1)), as signal scales linearly with the imaged volume. The process of simulating (one noise realisation of) dataset 1 is summarized in a flowchart in Fig. 7.3.3 for the simulations with unwanted inter-image motion, and in Fig. 7.3.4 for the simulations without motion.



Figure 7.3.2: Coronal views and transverse slices of the voxel-wise temporal SNR maps of the simulated datasets 1-4, and as calculated from the conventional HR *in vivo* 2D EPI dataset (see Table 7.3.2). Dashed lines are used to indicate transitions between multibands. Note that for the HR Datasets 2 and 4, and the Conventional HR Dataset, only a selection of 24 out of the 40 transverse slices is shown. For each dataset, the overall tSNR measure is indicated, which was calculated by taking the spatial mean inside a whole-brain mask of each voxel-wise tSNR map.

7.3.1.2 Parameter estimation

SRR-pCASL and SRR-pCASL-MB were applied to all noise realizations of dataset 1 and 3, respectively, whereas C-pCASL & BASIL and C-pCASL-MB & BASIL-MB were applied to dataset 2 and 4, respectively. For the simulation experiments with motion, the parameter optimization routine used in C-pCASL, C-pCASL-MB, SRR-pCASL, and SRR-pCASL-MB alternated between (P.1) and (P.2). For the outer loop iterations combining both (P.1) and (P.2), a convergence tolerance on the relative difference of the tissue parameters between consecutive iterations $\mathcal{E}^{(t)}$ was set at $\mathcal{E}_{\min} = 10^{-4}$, with a maximum number of 10 iterations. Regarding the inner iterations, the convergence tolerance on the relative difference of the tissue parameters between consecutive iterations for (P.1) was also set at $\mathcal{E}_{min} = 10^{-4}$, with the maximum number of iterations set at 120. Each of the decoupled sub-problems of (P.2) was solved using a lower bound of $\mu = 10^{-6}$ on the step size as convergence tolerance, i.e. iterations end when $\|\boldsymbol{\theta}^{(t-1)} - \boldsymbol{\theta}^{(t)}\|_2 < \mu$. The regularization parameters in (P.1) were heuristically set to $\lambda'_{r_1} = 1.6 \cdot 10^{-3}$ and $\lambda'_{rCBF} = 2.0 \cdot 10^{-5}$, balancing the trade-off between the data consistency objective and the regularization objectives of the tissue parameter maps $r_{
m 1}$ and $artheta_{
m rCBF}$, respectively. To compare the estimation methods independent of the choice of regularization parameters, the same values for λ'_{r_1} and λ'_{rCBF} were used for C-pCASL, C-pCASL-MB, SRR-pCASL and SRR-pCASL-MB. For the simulations without motion, only (P.1) was solved in the parameter optimization routine of C-pCASL, C-pCASL-MB,



Figure 7.3.3: A flowchart of the data simulation process for single-PLD pCASL data using SRR, in correspondence with the procedure outlined in Section 7.3.1. Coronal slices are shown for four slice-encoding directions, illustrating the forward modelling of HR ground-truth parameter maps to LR MS images **with unwanted motion**. Signal intensities of the control and label images are shown in arbitrary units.



Figure 7.3.4: A flowchart of the data simulation process for single-PLD pCASL data using SRR, in correspondence with the procedure outlined in Section 7.3.1. Coronal slices are shown for four slice-encoding directions, illustrating the forward modelling of HR ground-truth parameter maps to LR MS images **without unwanted motion**. Signal intensities of the control and label images are shown in arbitrary units.

SRR-pCASL, and SRR-pCASL-MB, using the same tolerance settings and regularization parameters for (P.1) as for the simulations with motion. For BASIL and BASIL-MB, motion correction using FSL's *mcflirt* was only turned on in the simulations with motion. To facilitate voxel-wise division with the calibration image ρ_{reg} in the simulation experiments, the ρ_{reg} image and the first control image were assumed to be perfectly aligned and the first control image was used as a reference target image when estimating the motion of the other images, both for the proposed framework and for BASIL.

7.3.1.3 Performance analysis

The CBF estimates obtained by the individual methods were compared based on a voxel-wise analysis of the accuracy and precision of each method using the following performance measures (Beirinckx et al., 2020, 2022):

Absolute relative bias (arBias), which quantifies the *accuracy* of an estimator, calculated as $\left|(\bar{\hat{\vartheta}}_{CBF} - \vartheta_{CBF}) \oslash \vartheta_{CBF}\right|$, where $\bar{\hat{\vartheta}}_{CBF}$ and ϑ_{CBF} refer to the CBF maps which contain the element-wise sample mean of the N_{MC} estimates $\hat{\vartheta}_{CBF}$, and the true reference values, respectively, and where \oslash denotes the element-wise division operator.

Relative standard deviation (rSTD) , which quantifies the *precision* of an estimator, calculated as

 $\left(\frac{N_{\text{MC}}}{N_{\text{MC}-1}} \overline{(\hat{\boldsymbol{\vartheta}}_{\text{CBF}} - \bar{\hat{\boldsymbol{\vartheta}}}_{\text{CBF}}) \circ (\hat{\boldsymbol{\vartheta}}_{\text{CBF}} - \bar{\hat{\boldsymbol{\vartheta}}}_{\text{CBF}})}\right)^{\circ \frac{1}{2}} \oslash \boldsymbol{\vartheta}_{\text{CBF}}, \text{ where } \circ \text{ and the superscript } \circ \frac{1}{2} \text{ denote the Hadamard product and element-wise square-root operator, respectively.}$

Relative root-mean-squared error (rRMSE), which is a combined measure of accuracy and precision, calculated as $\left((\hat{\vartheta}_{\mathsf{CBF}} - \vartheta_{\mathsf{CBF}}) \circ (\hat{\vartheta}_{\mathsf{CBF}} - \vartheta_{\mathsf{CBF}})\right)^{\circ \frac{1}{2}} \oslash \vartheta_{\mathsf{CBF}}$.

In addition, the spatial mean of each of these performance measure maps was computed, yielding \overline{arBias} , \overline{rSTD} , and \overline{rRMSE} , respectively. To further assess image quality of the estimated CBF maps compared to the ground truth HR CBF map, average structural similarity index measure (SSIM) and peak SNR (PSNR) values were obtained for each method by calculating the sample mean of the SSIM and PSNR values obtained for each of the $N_{\rm MC}$ realisations.

To assess the ability of the different frameworks to estimate motion, the following performance measure was used:

Motion component root-(mean)-mean-squared-error (RMMSE) , which is defined as

 $\left(\frac{1}{2N}\sum_{n=1}^{2N}\overline{(\hat{\theta}_n-\theta_n)\circ(\hat{\theta}_n-\theta_n)}\right)^{\circ\frac{1}{2}}$, where θ_n refers to the true reference value and the operator $\overline{(\cdot)}$ denotes the element-wise sample mean over the $N_{\rm MC}$ estimates $\hat{\theta}_n$.

Next, to visually compare the estimated CBF values against the reference CBF values, a 2D scatter plot for each method was generated between $\hat{\vartheta}_{CBF}$ and ϑ_{CBF} . Additionally, following the definition of Delbany et al. (2019), the SNR gain map $\Gamma_{X,Y} \in \mathbb{R}^{N_r \times 1}$ between method X and method Y was calculated as $\Gamma_{X,Y} = \mathbf{SNR}_X \otimes \mathbf{SNR}_Y$. Here, \mathbf{SNR}_X represents the average SNR map of the reconstructed CBF maps in the simulation experiments for method X, which was calculated as the ratio of the element-wise sample mean and standard deviation of the N_{MC} estimates $\hat{\vartheta}_{CBF,X}$ for method X, i.e., $\begin{aligned} & \text{SNR}_X = \bar{\hat{\vartheta}}_{\text{CBF},X} \circ \left(\frac{N_{\text{MC}}}{N_{\text{MC}-1}} \overline{(\hat{\vartheta}_{\text{CBF},X} - \bar{\hat{\vartheta}}_{\text{CBF},X}) \circ (\hat{\vartheta}_{\text{CBF},X} - \bar{\hat{\vartheta}}_{\text{CBF},X})} \right)^{\circ \left(-\frac{1}{2}\right)}. \end{aligned} \\ & \text{ Note that } \Gamma_{X,Y} \\ & \text{ incorporates both the SNR gain due to the use of a different acquisition strategy as well as due to a different reconstruction algorithm being used between both methods. Finally, the spatial mean of each SNR gain map was computed, yielding overall SNR gain values $\overline{\Gamma}_{X,Y}$ between the different methods. } \end{aligned}$

Table 7.3.1: Acquisition settings for the synthetic data sets using 2D MS readout. A slice orientation angle of 0° corresponds with the slice-encoding axis directed from left to right, and with the phase-encoding axis perpendicularly directed from anterior to posterior. Each angle listed below is a rotation of the slice-encoding axis around the phase-encoding direction counterclockwise. Therefore, a 90° angle is consistent with an ascending slice order. These rotations are consistent with the rotations visualized in Fig. 7.2.1.

	Dataset 1 LR 2D MS	Dataset 2 HR 2D MS	Dataset 3 LR 2D MS MB	Dataset 4 HR 2D MS MB	
Slices per slab N_{slice} (#)	16	40	16	40	
Acquisition matrix	80×80	80 × 80	80×80	80 × 80	
FOV [mm ³]	$240\times240\times192$	$240\times240\times120$	$240\times240\times192$	$240 \times 240 \times 120$	
Voxel size [mm ³]	$3 \times 3 \times 12$	$3 \times 3 \times 3$	$3 \times 3 \times 12$	$3 \times 3 \times 3$	
Labeling duration $ au$ [ms]	1800	1800	1800	1800	
PLD _{base} [ms]	1800	1800	1800	1800	
PLD range ¹ [ms]	1800-2550	1800-3750	1800-2150	1800-2750	
Slice readout time t_{read} [ms]	50	60	50	60	
# control-label pairs N	24	22	24	22	
# slice encoding directions	24	1	24	1	
Slice orientation angles ² [°]	0, 7.5,, 172.5	90	0, 7.5,, 172.5	90	
Multiband factor ω	n.a.	n.a.	2	2	
Theor. scan time ³ T [min:s]	3:30	4:20	3:10	3:30	

¹ For a dataset with MB, the PLD range is given for a single band.

² For the LR datasets, the slice orientation angles were chosen by rotating the slice stack around the virtual phase encoding axis in increments of 180/N degrees. For N = 24, each rotational increment is equal to 7.5°.

³ Defined as $T = 2N \cdot (\tau + \mathsf{PLD}_{\mathsf{base}} + N_{\mathsf{slice}} \cdot t_{\mathsf{read}} / \omega)$

7.3.2 Real data experiment

The performance of the proposed SRR-pCASL method was also evaluated using *in vivo* brain MS single-PLD pCASL data from a healthy volunteer (adult, male, 29 years old), acquired using a 32-channel head coil on a 3 Tesla-scanner (Achieva, Philips Healthcare). Ethical approval from the local institutional review board was obtained and an informed consent was signed by the volunteer. The pCASL data was acquired using a single-shot 2D EPI readout method, as recommended by (Alsop et al., 2015). LR MS data for SRR as well as conventional MB MS data directly acquired at high resolution were collected using the acquisition settings tabulated in Table 7.3.2. Data sets were acquired without slice gap. A larger FOV for the LR data set compared to that of the conventional HR data set is needed because the entire brain has to be within the FOV for each rotation angle. Also note that, even though the readout time is significantly longer when acquiring 40 slices instead of 16, this results in only 2 control-label image pairs less for the conventional HR dataset compared to the LR dataset, given a fixed total scan time for both protocols (see Table 7.3.2). This

Table 7.3.2: Acquisition settings for SRR data and conventionally acquired data using 2D MS readout. A slice orientation angle of 0° corresponds with the slice-encoding axis directed from left to right, and with the phase-encoding axis perpendicularly directed from anterior to posterior. Each angle listed below is a rotation of the slice-encoding axis around the phase-encoding direction counterclockwise. Therefore, a 90° angle is consistent with an ascending slice order. These rotations are consistent with the rotations visualized in Fig. 7.2.1.

	LR 2D MS	HR 2D MS MB
Slices per slab (#)	16	40
Acquisition matrix	80 × 80	80 × 80
FOV [mm ³]	$240 \times 240 \times 192$	$240 \times 240 \times 120$
Voxel size [mm ³]	$3 \times 3 \times 12$	$3 \times 3 \times 3$
TR [ms]	4400	4800
Labeling duration [ms]	1800	1800
PLD _{base} [ms]	1800	1800
PLD range ¹ [ms]	1800-2550	1800-2750
Number of control-label pairs	24	22
Number of slice encoding directions	24	1
Slice orientation angles [°]	0, 7.5,, 172.5	90
SMS (multiband factor)	n.a.	yes (factor 2)
Total scan time [min:s]	3:30	3:30

¹ For the conventional HR data, the PLD range is given for a single band.

is a direct consequence of the fact that the labeling duration and the PLD take up most of the scan time. In addition to the pCASL data, a proton density weighted calibration image was acquired at isotropic high resolution for absolute CBF quantification. CBF map estimates were obtained from the LR MS data using the proposed SRR-pCASL method, and compared to the CBF maps estimated from the conventional MS MB data using BASIL-MB and C-pCASL-MB. For the acquisition of the conventional HR MS ASL data, an MB factor of 2 was used to limit ASL signal loss in the upper part of the brain. In contrast, MB was not used in the acquisition of the LR SRR-pCASL data because an MB acquisition required a mandatory calibration scan to be performed before the acquisition of each LR image pair with adjusted slice orientation. As a result, the condition of equal total scan time for the conventional HR dataset and the SRR-pCASL dataset would no longer apply.

The parameter optimization routine used in C-pCASL-MB and SRR-pCASL alternated between (P.1) and (P.2). For the outer loop iterations combining both (P.1) and (P.2), a convergence tolerance on the relative difference of the tissue parameters between consecutive iterations $\mathcal{E}^{(t)}$ was set at $\mathcal{E}_{\min} = 10^{-3}$, with a maximum number of 10 iterations. Regarding the inner iterations, the convergence tolerance on the relative difference of the tissue parameters between consecutive iterations for (P.1) was also set at $\mathcal{E}_{\min} = 10^{-3}$, with the maximum number of iterations set at 120. Each of the decoupled sub-problems of (P.2) was solved using a lower bound of $\mu = 10^{-3}$ on the step size as convergence tolerance, i.e. iterations end when $\|\boldsymbol{\theta}^{(t-1)} - \boldsymbol{\theta}^{(t)}\|_2 < \mu$. Regularization weights for the *in vivo* reconstructions were heuristically set at $\lambda'_{r_1} = 8 \cdot 10^{-3}$ and $\lambda'_{rCBF} = 1 \cdot 10^{-4}$, again similar for C-pCASL-MB and SRR-pCASL. The *in vivo* data for BASIL-MB was motion corrected using FSL's *mcflirt* (Jenkinson et al., 2002).

7.4 Results

7.4.1 Simulation experiments

7.4.1.1 Simulation experiments without motion

Table 7.4.1 summarizes the results of the whole brain simulation CBF mapping experiments in terms of the average SSIM, PSNR, arBias, rSTD, and rRMSE, for BASIL, C-pCASL, and SSR-pCASL (with and without MB). For each performance measure, the best performing framework is highlighted in green. It follows from Table 7.4.1 that SRR-pCASL consistently resulted in higher average SSIM and PSNR values compared to traditional CBF quantification with BASIL and C-pCASL. Of all methods studied, SRR-pCASL-MB outperformed the other approaches in terms of average SSIM and PSNR. C-pCASL-MB had the lowest arBias value of all methods, and performed best in terms of overall accuracy. SRR-pCASL and SRRpCASL-MB had the lowest \overline{rSTD} values, outperforming the other methods in terms of overall precision. In terms of overall RMSE, being a measure that incorporates both accuracy and precision, SRR-pCASL outperformed BASIL and C-pCASL, having an *rRMSE* value that is about 30% and 32% smaller than that of C-pCASL and BASIL, respectively. The addition of MB provided consistent improvement for each performance measure for each method, except for SRR-pCASL-MB where the arBias value slightly increased. Yet, combining SRR-pCASL with MB, resulting in SRR-pCASL-MB, provided a notable improvement in CBF estimation precision and RMSE. Indeed, rSTD and rRMSE for SRR-pCASL-MB decreased with about 14% and 10% compared to SRR-pCASL, respectively.

Table 7.4.1: Quantitative performance measures with standard error (SE) for the whole brain simulation experiment **without motion**, calculated over $N_{MC} = 100$ reconstruction results for CBF mapping, for each respective readout scheme and reconstruction framework. For each performance measure, the value of the best performing strategy is highlighted in green.

	BASIL		BASIL C-pCASL S		SRR-pC	SRR-pCASL		BASIL-MB		C-pCASL-MB		SRR-pCASL-MB	
	value	SE	value	SE	value	SE	value	SE	value	SE	value	SE	
SSIM	0.9833	1e-4	0.9846	1e-4	0.9927	1e-4	0.9894	1e-4	0.9905	1e-4	0.9940	1e-4	
PSNR [dB]	30.97	0.02	30.99	0.01	32.33	0.01	32.28	0.03	32.17	0.01	32.45	0.01	
arBias [%]	7.15	0.02	5.87	0.02	4.59	0.01	5.33	0.02	4.12	0.01	4.79	0.02	
rSTD [%]	17.15	0.03	17.29	0.03	11.71	0.02	13.87	0.02	13.81	0.02	10.07	0.02	
rRMSE [%]	19.27	0.04	18.76	0.03	13.07	0.02	15.42	0.02	14.81	0.02	11.68	0.02	

Fig. 7.4.1 shows coronal CBF maps estimated with each reconstruction framework as well as their absolute value of the rBias, rSTD, and rRMSE. Different aspects stand out. First, SRR-pCASL outperformed single-orientation pCASL in terms of CBF estimation accuracy, as illustrated by the coronal mid-slice of the arBias. For example, for BASIL a clear bias existed for gray matter estimates in the outer edges of the brain, while for C-pCASL there existed a significant bias in some voxels in the upper part of the brain. The latter may be attributed to the SNR of the ASL signal becoming critically low in these slices, which have the longest effective PLDs and the lowest degree of BS. For SRR-pCASL, the accuracy of the CBF estimation was more uniform across the brain, with no apparent differences between tissue types, or between top or lower parts of the brain. Second, the rSTD of the CBF estimates obtained using BASIL and C-pCASL increased from the lower parts of the brain towards the top parts of the brain (third row of Fig. 7.4.1). When using SRR-pCASL, on the other hand,



Figure 7.4.1: Coronal mid-slices of the CBF estimates and the corresponding quantitative performance measures for the whole brain simulation experiment **without motion**. The first row shows the numerical ground truth (left), followed by the estimated CBF maps for each method. Next, rows 2-4 show the absolute value of the relative bias, relative standard deviation, and relative RMSE, respectively, computed from the $N_{MC} = 100$ simulations.

the precision of CBF estimation was much more uniform, per tissue type, throughout the brain (Fig. 7.4.1). Furthermore, the addition of MB led to a reduction of the rSTD for each method. Whereas for BASIL-MB and C-pCASL-MB the precision improvement was limited to slices acquired in the second band only, for SRR-pCASL-MB these improvements were obtained across the whole brain. Third, in terms of rRMSE, SRR-pCASL clearly outperformed the other methods without MB, as indicated by the visibly darker rRMSE maps in the fourth row of Fig. 7.4.1. Here, the same trends as for the precision maps in Fig. 7.4.5 are visible, showing both an increase in the rRMSE for ascending slices during acquisition for BASIL and C-pCASL, and a more uniform rRMSE of the CBF estimation across brain regions for SRR-pCASL.

Fig. 7.4.2 (left) shows the locations of the transverse slices that were selected to visualize the variations in CBF estimation on a slice level. Each transverse slice is characterized by a unique PLD and degree of BS, depending on the acquisition settings of the processed dataset for each CBF estimation method. Ground truth values of CBF for these slice locations are given in the leftmost column of Fig. 7.4.3, including the GM and WM hyperperfusion lesions



Figure 7.4.2: Top left: locations of the transverse slices for the whole brain simulation experiment (shown in Fig. 7.4.3). Bottom left: locations of the transverse slices for the real data experiment (shown in Fig. 7.4.16). In addition, coronal mid-slice views of the slice-dependent effective PLD matrix are shown for each approach, reflecting the increase in PLD along the slice-encoding direction, starting from a PLD_{base} value of 1.8 seconds. Since for SRR-pCASL the effective PLD matrix rotates with the slice orientation of each LR image (see Fig. 7.2.1), only the effective PLD matrix corresponding with a slice orientation angle of 90° is shown for SRR-pCASL and SRR-pCASL-MB, consistent with an ascending slice order. Note that missing data in the bottom panel is due to these scans not being acquired in the real data experiment.

denoted by the white arrows in the fourth axial slice. Fig. 7.4.3 also shows the rRMSE maps for each slice location and method. Only slices D and E, which were positioned halfway and at the beginning of the second band, respectively, showed a lower rRMSE for BASIL-MB and C-pCASL-MB compared to BASIL and C-pCASL, respectively. Whereas for BASIL-MB and C-pCASL-MB the rRMSE improvement was limited to slices acquired in the second band only, for SRR-pCASL-MB these improvements were obtained across all slices. Axial slice views of the absolute value of the rBias and the rSTD maps, corresponding with the slice locations of Fig. 7.4.2, are shown in Figs. 7.4.4-7.4.5. Furthermore, to appreciate resolution enhancements and to ease a qualitative and visual comparison of the CBF map estimated by each method, Figs. 7.4.6-7.4.7, show orthogonal slice views of the simulated 2D control images for each simulated dataset in comparison with the estimated HR CBF map per framework, and zoomed close-ups of this CBF map in comparison with the ground truth CBF map as a reference, respectively.

Additionally, to demonstrate that the potential of SRR is not confined to a particular resolution, an extra simulation experiment was performed where a $2 \times 2 \times 2$ mm³ CBF map was super-resolution reconstructed from LR pCASL images with a resolution of $2 \times 2 \times 16$ mm³. Acquisition settings and results for this additional simulation experiment are summarized in Table 7.C.1 and Fig. 7.D.1 in Appendix 7.D.

Next, Fig. 7.4.8 shows the average estimated CBF values against the reference CBF values in a 2D scatter plot for BASIL (left), C-pCASL (mid), and SRR-pCASL (right) without (top) and with (bottom) MB. As indicated by the narrower distribution (i.e. better precision) and the higher number of voxels that match the ground truth values (i.e. better accuracy), SRR-pCASL and SRR-pCASL-MB outperformed the other methods.



Figure 7.4.3: Relative RMSE maps for CBF, calculated from the reconstruction results of the whole brain simulation experiment **without motion**. For each method, five transverse slices are shown, corresponding with the slice letter convention in Fig. 7.4.2. Overall relative RMSE values are summarized in Table 7.4.1. The numerical ground truth CBF map is shown in column 1. Both hyperperfusion lesions are indicated by white arrow marks in slice B of this ground truth CBF map.

Table 7.4.2: The overall SNR gain ($\overline{\Gamma}_{X,Y}$) with standard error (SE) between reconstruction methods as assessed by the whole brain simulation experiment **without motion**, calculated over the $N_{MC} = 100$ reconstruction results for CBF mapping. Note that row labels refer to method X, while column labels refer to method Y, in line with the definition of $\overline{\Gamma}_{X,Y}$ in Section 7.3.1.3.

$\overline{\Gamma}_{X,Y}$	BASIL		C-pCASL		SRR-pCASL		BASIL-MB		C-pCASL-MB		SRR-pCASL-MB	
	value	SE	value	SE	value	SE	value	SE	value	SE	value	SE
BASIL	1.000	0.000	1.047	0.001	0.709	0.001	0.845	0.001	0.859	0.001	0.623	0.001
C-pCASL	0.973	0.001	1.000	0.000	0.686	0.001	0.822	0.001	0.831	0.001	0.604	0.001
SRR-pCASL	1.446	0.001	1.505	0.001	1.000	0.000	1.200	0.001	1.217	0.001	0.875	0.001
BASIL-MB	1.252	0.001	1.310	0.001	0.870	0.001	1.000	0.000	1.017	0.001	0.754	0.001
C-pCASL-MB	1.251	0.001	1.302	0.002	0.867	0.001	0.998	0.001	1.000	0.000	0.750	0.001
SRR-pCASL-MB	1.674	0.001	1.748	0.002	1.153	0.001	1.371	0.001	1.389	0.001	1.000	0.000



Figure 7.4.4: Absolute value of relative bias maps for CBF, calculated from the reconstruction results of the synthetic whole brain simulations **without motion**. For each method, five transverse slices are shown, corresponding with the slice letter convention in Fig. 7.4.2.

Finally, Table 7.4.2 summarizes the overall SNR gains between the different methods. Here, SRR-pCASL outperformed the other methods in terms of SNR of the estimated CBF map, as illustrated by the overall SNR gains over BASIL and C-pCASL, even if these methods exploited MB during acquisition. The SNR gain was maximal when SRR-pCASL was combined with MB.



Figure 7.4.5: Relative standard deviation maps for CBF, calculated from the reconstruction results of the synthetic whole brain simulations **without motion**. For each method, five transverse slices are shown, corresponding with the slice letter convention in Fig. 7.4.2.



Figure 7.4.6: Orthogonal slice views of simulated HR 2D control images with high through-plane resolution $(3 \times 3 \times 3 \text{ mm}^3)$, without multiband (top panel, column 1) and with multiband (top panel, column 4), and simulated LR 2D control images with low through-plane resolution $(3 \times 3 \times 12 \text{ mm}^3)$, without multiband (bottom panel, columns 1-2) and with multiband (bottom panel, columns 4-5). For illustration purposes, only two slice orientation angles (0° and 90°) are shown for the 2D LR control images (see also Fig. 7.2.1). Simulated control images are compared to the corresponding HR CBF map estimates reconstructed with BASIL (top panel, column 2), C-pCASL (top panel, column 3), BASIL-MB (top panel, column 5), C-pCASL-MB (top panel, column 6), and the proposed SRR-pCASL (bottom panel, column 3) and SRR-pCASL-MB (bottom panel, column 6), respectively.



Figure 7.4.7: Orthogonal slice views with zoomed close-ups showing the high-resolution CBF map estimated with BASIL (column 3), C-pCASL (column 4), BASIL-MB (column 6), C-pCASL-MB (column 7), and the proposed SRR-pCASL (column 5) and SRR-pCASL-MB (column 8), compared to the ground truth CBF map as a reference (columns 1-2) for the whole brain simulation experiment **without motion**.



Figure 7.4.8: 2D histograms between reference values and estimated values for all methods, as assessed by the whole brain simulation experiment **without motion**. CBF values are given in mL/100g/min. The dashed line represents identity. Points below correspond to underestimation and points above to overestimation, compared to the reference value. For each method, values were averaged over the $N_{MC} = 100$ estimates.

7.4.1.2 Simulation experiments with motion

Table 7.4.3: Quantitative performance measures with standard error (SE) for the whole brain simulation experiment **with motion**, calculated over $N_{MC} = 100$ reconstruction results for CBF mapping, for each respective readout scheme and reconstruction framework. For each performance measure, the value of the best performing strategy is highlighted in green.

	BASIL		BASIL C-pCA		A SL	SL SRR-pCASL		BASIL	BASIL-MB		L-MB	SRR-pCASL-MB	
	value	SE	value	SE	value	SE	value	SE	value	SE	value	SE	
SSIM	0.9705	1e-4	0.9845	1e-4	0.9920	1e-4	0.9825	1e-4	0.9900	1e-4	0.9936	1e-4	
PSNR [dB]	28.56	0.02	30.79	0.01	31.54	0.01	30.41	0.03	31.86	0.01	32.29	0.01	
arBias [%]	15.86	0.05	6.27	0.02	5.51	0.02	12.92	0.04	4.61	0.02	5.37	0.02	
rSTD [%]	10.41	0.03	17.30	0.03	11.70	0.02	8.35	0.02	13.85	0.02	9.87	0.02	
rRMSE [%]	20.75	0.05	18.95	0.03	13.54	0.02	16.55	0.04	15.07	0.02	11.84	0.02	
RMMSE	value	SE	value	SE	value	SE	value	SE	value	SE	value	SE	
t_{x} [mm]	0.2	0.1	0.0016	3e-4	0.0020	3e-4	0.3	0.1	0.0021	3e-4	0.0028	4e-4	
<i>ty</i> [mm]	0.11	0.03	0.0034	4e-4	0.017	1e-3	0.18	0.03	0.0006	2e-4	0.0094	8e-4	
t _z [mm]	0.27	0.06	0.0015	2e-4	0.019	1e-3	0.56	0.09	0.0019	3e-4	0.042	2e-3	
α [deg]	0.8	0.2	0.0004	1e-4	0.0106	9e-4	0.8	0.2	0.0005	2e-4	0.021	1e-3	
β [deg]	0.20	0.05	0.0004	1e-4	0.0019	4e-4	0.3	0.1	0.0004	1e-4	0.0042	6e-4	
γ [deg]	0.5	0.2	0.0014	3e-4	0.0022	3e-4	0.8	0.2	0.0010	2e-4	0.0043	6e-4	

Table 7.4.3 summarizes the obtained quantitative performance measures for the whole brain simulation CBF mapping experiments with motion. It follows from Table 7.4.3 that the need to estimate unwanted motion during the CBF reconstruction degrades the average SSIM, PSNR, arBias, and rRMSE value for each method. This effect is most pronounced for BASIL and BASIL-MB, where motion between the pCASL images was corrected using a registration routine prior to CBF quantification. Without any exception, the addition of MB provided consistent improvement for each performance measure for each method. Similar to the simulations without motion, C-pCASL-MB had the lowest \overline{arBias} value of all methods, and performed best in terms of overall accuracy. For BASIL and BASIL-MB, the arBias value decreased by more than a factor of 2 compared to the simulations without motion, indicating a considerable drop in accuracy. In terms of overall precision, quantified by the rSTD value, C-pCASL and SRR-pCASL performed very similar compared to the simulations without motion (Table 7.4.1), whereas BASIL and BASIL-MB showed a noticeable improvement in rSTD. In terms of overall RMSE, SRR-pCASL clearly outperformed the other approaches without MB, having an *rRMSE* value that is about 34% and 28% smaller than that of BASIL and C-pCASL, respectively. A similar observation is true when MB was added, with SRR-pCASL-MB outperforming the other methods in terms of overall RMSE, having an *rRMSE* value that is about 24% and 21% smaller than that of BASIL-MB and C-pCASL-MB, respectively. For BASIL and BASIL-MB, which apply an adaptive spatial smoothing to the estimated perfusion image, the increased precision (reduced rSTD) somewhat compensates for the reduced accuracy (increased *arBias*), compared to the simulations without motion (Table 7.4.1). This compensating effect also follows from Fig. 7.4.9, which shows coronal mid-slices of the CBF estimates and the corresponding quantitative performance measures for the simulation experiment with motion. As indicated by the coronal mid-slice of the arBias (second row of Fig. 7.4.9), a clear bias existed for BASIL and BASIL-MB in all areas of the brain, whereas a reduced precision was observed of the rSTD map (third row of



Figure 7.4.9: Coronal mid-slices of the CBF estimates and the corresponding quantitative performance measures for the whole brain simulation experiment **with motion**. The first row shows the numerical ground truth (left), followed by the estimated CBF maps for each method. Next, rows 2-4 show the absolute value of the relative bias, relative standard deviation, and relative RMSE, respectively, computed from the $N_{MC} = 100$ simulations.

Fig. 7.4.9), most notably for gray matter voxels. Axial slice views of the absolute value of the rBias and the rSTD maps, corresponding with the slice locations of Fig. 7.4.2, are shown in Figs. 7.4.12-7.4.13.

In addition, Table 7.4.3 also summarizes the motion component RMMSE for each of the six rigid motion components. The proposed framework using joint motion estimation clearly outperformed the BASIL reference method in terms of the motion component RMMSE, with C-pCASL (without and with MB) performing best. Although RMMSE values for C-pCASL are lower than for SRR-pCASL, this does not result in lower *rRMSE* values, indicating that the benefits of a SRR acquisition with rotated slice-encoding and low through-plane resolution, i.e. more optimal BS and more constant PLD across slices, can outplay small inaccuracies/imprecision in motion estimation. The effect of an improved estimation of motion parameters is also visible from Fig. 7.4.11, where C-pCASL and SRR-pCASL, without and with the use of MB, result in a narrower distribution (i.e. better precision) and a higher number of voxels that match the ground truth values (i.e. better accuracy), compared to BASIL.



Figure 7.4.10: Relative RMSE maps for CBF, calculated from the reconstruction results of the whole brain simulation experiments **with motion**. For each method, five transverse slices are shown, corresponding with the slice letter convention in Fig. 7.4.2. Overall relative RMSE values are summarized in Table 7.4.3. The numerical ground truth CBF map is shown in column 1. Both hyperperfusion lesions are indicated by white arrow marks in slice B of this ground truth CBF map.

Finally, as can be seen from Table 7.4.4, which summarizes the overall SNR gains between the different methods for the simulation experiment with motion, SRR-pCASL outperformed C-pCASL and C-pCASL-MB. However, BASIL and BASIL-MB outperformed the proposed approach in terms of SNR of the estimated CBF map, with a maximal SNR gain for when BASIL was combined with MB. The increased SNR of BASIL may be attributed to the increase in spatial regularisation as a result of its adaptive smoothing prior, which also explains the increased precision (lower rSTD) and reduced accuracy (higher arBias) in Table 7.4.4. To further support this observation and indicate the increased smoothness in the reconstruction results of BASIL, particularly for gray matter, Fig. 7.4.14 shows orthogonal slice views with zoomed close-ups of the HR CBF maps estimated with the different approaches, compared to the ground truth CBF map as a reference.


Figure 7.4.11: 2D histograms between reference values and estimated values for all methods, as assessed by the whole brain simulation experiment **with motion**. CBF values are given in mL/100g/min. The dashed line represents identity. Points below correspond to underestimation and points above to overestimation, compared to the reference value. For each method, values were averaged over the $N_{MC} = 100$ estimates.

Table 7.4.4: The overall SNR gain ($\overline{\Gamma}_{X,Y}$) with standard error (SE) between reconstruction methods as assessed by the whole brain simulation experiment **with motion**, calculated over the $N_{MC} = 100$ reconstruction results for CBF mapping. Note that row labels refer to method *X*, while column labels refer to method *Y*, in line with the definition of $\overline{\Gamma}_{X,Y}$ in Section 7.3.1.3.

$\overline{\Gamma}_{X,Y}$	BASIL		C-pCASL		SRR-pCASL		BASIL-MB		C-pCASL-MB		SRR-pCASL-MB	
	value	SE	value	SE	value	SE	value	SE	value	SE	value	SE
BASIL	1.000	0.000	1.872	0.003	1.273	0.002	0.888	0.001	1.537	0.003	1.086	0.002
C-pCASL	0.589	0.001	1.000	0.000	0.694	0.001	0.497	0.001	0.830	0.001	0.592	0.001
SRR-pCASL	0.863	0.001	1.489	0.001	1.000	0.000	0.719	0.001	1.207	0.001	0.849	0.001
BASIL-MB	1.253	0.002	2.203	0.003	1.487	0.002	1.000	0.000	1.741	0.002	1.252	0.002
C-pCASL-MB	0.758	0.001	1.291	0.001	0.874	0.001	0.608	0.001	1.000	0.000	0.734	0.001
SRR-pCASL-MB	1.049	0.002	1.810	0.002	1.210	0.001	0.863	0.001	1.445	0.001	1.000	0.000



Figure 7.4.12: Absolute value of relative bias maps for CBF, calculated from the reconstruction results of the synthetic whole brain simulations **with motion**. For each method, five transverse slices are shown, corresponding with the slice letter convention in Fig. 7.4.2.



Figure 7.4.13: Relative standard deviation maps for CBF, calculated from the reconstruction results of the synthetic whole brain simulations **with motion**. For each method, five transverse slices are shown, corresponding with the slice letter convention in Fig. 7.4.2.



Figure 7.4.14: Orthogonal slice views with zoomed close-ups showing the high-resolution CBF map estimated with BASIL (column 3), C-pCASL (column 4), BASIL-MB (column 6), C-pCASL-MB (column 7), and the proposed SRR-pCASL (column 5) and SRR-pCASL-MB (column 8), compared to the ground truth CBF map as a reference (columns 1-2) for the whole brain simulation experiment **with motion**.

7.4.2 Real data experiment

Fig. 7.4.15 shows orthogonal mid-slice views of an HR-MB 2D EPI control image acquired with high through-plane resolution $(3 \times 3 \times 3 \text{ mm}^3)$ and MB, and an LR 2D EPI control image with low through-plane resolution $(3 \times 3 \times 12 \text{ mm}^3)$ acquired with a slice orientation angle of 0°, corresponding with the acquisition settings summarized in Table 7.3.2. Fig. 7.4.15 also shows the HR CBF map estimates obtained with BASIL-MB, C-pCASL-MB, and SRR-pCASL, respectively. Note that for the LR 2D EPI *in vivo* data set, SRR-pCASL successfully recovered the fine details from the set of LR images.



Figure 7.4.15: Orthogonal slice views for the real data experiment showing a HR-MB 2D EPI control image acquired with high through-plane resolution $(3 \times 3 \times 3 \text{ mm}^3)$ and multiband (first column), and a LR 2D EPI control image with low through-plane resolution $(3 \times 3 \times 12 \text{ mm}^3)$ corresponding with a slice orientation angle of 0° (column 4), compared with the HR CBF map estimates obtained with BASIL-MB (second column), C-pCASL-MB (third column), and SRR-pCASL (column 5), respectively.

In addition, a series of transverse slices at different locations in the brain of the estimated HR quantitative CBF maps is shown in Fig. 7.4.16 for BASIL-MB, C-pCASL-MB, and the proposed SRR-pCASL. The locations of these transverse slices are highlighted on a coronal view of the CBF map reconstructed with BASIL-MB in Fig. 7.4.2 (bottom left). As indicated in Fig. 7.4.2, slices A and B correspond with later acquired slices in the first MB segment for BASIL-MB and C-pCASL-MB, while the other four slices are acquired relatively early in the second MB segment. When comparing the different methods, two aspects stand out. First, as can be observed in Fig. 7.4.16, the HR CBF maps reconstructed using SRR-pCASL, C-pCASL-MB and BASIL-MB are comparable in terms of absolute values and visualized anatomical structures. This clearly demonstrates the feasibility of combining SRR with single-PLD pCASL. Second, the reconstructed slices shown in Fig. 7.4.16 for the proposed SRR-pCASL approach all have comparable CBF values, reflecting the relative uniformity in SNR throughout all regions in the brain. The CBF maps obtained from the conventional

HR MS data using C-pCASL-MB and BASIL-MB, on the other hand, clearly suffer from low SNR in the superior slices of the first MB segment (slices A and B in Fig. 7.4.16). For these slice locations, the proposed SRR-pCASL outperforms C-pCASL-MB and BASIL-MB in terms of reconstruction quality of the underlying anatomy.



Figure 7.4.16: Transverse slices at different locations in the brain of the estimated HR ($3 \times 3 \times 3$ mm³) CBF maps for the different real data approaches. In the first row the estimated CBF map is shown for the proposed SRR-pCASL on the LR dataset (SRR-pCASL), and in the following rows the estimated CBF map is shown for C-pCASL on the HR dataset with multiband factor 2 (C-pCASL-MB), and BASIL on the HR dataset with multiband factor 2 (BASIL-MB), respectively. The CBF maps estimated with SRR-pCASL all have comparable CBF values, which reflects the relative uniformity in average SNR throughout all regions in the brain as a consequence of acquiring the LR images with a rotational acquisition strategy. C-pCASL-MB and BASIL-MB, on the other hand, suffer from low SNR in the posterior slices of the first multiband segment (Slice A and B) due to longer effective PLDs and limited background suppression. Slice positions correspond with those given in Fig. 7.4.2.

As stated above, simulation experiments were performed for BASIL-MB and C-pCASL-MB, mimicking the same MB factor of 2 as in the real data HR pCASL experiment. Note that the stability of the CBF values across slices in SRR-pCASL (see Fig. 7.4.16) is consistent with the uniform RMSE of CBF estimation from LR MS data shown in the simulation experiment (see Fig. 7.4.3). Furthermore, the higher quality of the CBF map obtained using SRR-pCASL compared to that of BASIL-MB or C-pCASL-MB in regions of the brain that were imaged last within the MB segment (see the first two slices shown in Fig. 7.4.16) matches with the difference in RMSE of CBF estimation between both methods in those same regions

as predicted in the simulation experiment (see Fig. 7.4.3). While it is difficult to compare a qualitative assessment (real data) with a quantitative measurement of RMSE (simulation data), it is reasonable to assume both effects are correlated. It serves as an indication of the validity of the simulation experiment.

7.5 Discussion

In this contribution, we introduced a model-based SRR framework for single-PLD pCASL MRI. The framework, which integrates inter-image motion estimation, provides direct whole brain 3D isotropic high resolution CBF mapping from a set of 2D multi-slice control-label image pairs with a low through-plane resolution and slice orientations that are pair-wise rotated around a common phase encoding axis. Simulation and real data results show that this SRR acquisition strategy enables improved CBF mapping compared to the conventional 2D MS single-PLD pCASL acquisition scheme in which control-label image pairs are directly acquired at the target isotropic resolution with equal acquisition settings for each image pair, while using the same scan time. Our findings are discussed in more detail below.

7.5.1 Differences with existing techniques

We would like to point out that this contribution differs noticeably from another multi-image super-resolution resolution study for (pC)ASL, presented by Shou et al. (2021). In that work, the SLIce Dithered Enhanced Resolution (SLIDER) super-resolution technique proposed by Setsompop et al. (2015) is integrated with 2D SMS pCASL and a constrained slice-dependent background suppression (CSD-BS) scheme (Shao et al., 2018). Our approach improves upon the method presented by Shou et al. (2021) in various aspects. First, SLIDER relies on sub-voxel spatial shifts in the slice direction, whereas in our approach the slice orientations are rotated around the phase-encoding axis, which yields a more effective sampling of the k-space (Plenge et al., 2012). Previous studies that compared translational (i.e., sub-voxel shift) and rotational SRR schemes confirmed the superiority of the latter (Shilling et al., 2009; Nicastro et al., 2022). Furthermore, unlike the SRR-pCASL method proposed in this chapter, the SLIDER-SMS pCASL method proposed by Shou et al. (2021) does not estimate CBF directly from the LR images, nor does it integrate simultaneous motion estimation, which may introduce a bias due to error propagation. Moreover, the SLIDER-SMS pCASL method assumes a perfect slice profile, whereas our SRR-PCASL method models the slice profile as a more realistic smoothed box function (Poot et al., 2010). In addition, the long total scan time and lack of motion compensating steps of the SLIDER-SMS pCASL method increase the susceptibility to motion artifacts, which Shou et al. (2021) identify as a limitation of their study. As demonstrated in this contribution, the proposed SRR-pCASL framework integrates an inter-image motion model making it less susceptible to motion artifacts.

7.5.2 Improved CBF quantification from single-PLD pCASL data

Comparing the CBF estimation of the different approaches in the simulation experiment without motion, the proposed SRR-pCASL framework for LR MS pCASL data showed superior CBF parameter mapping RMSE compared to both C-pCASL and BASIL for processing of conventional HR MS pCASL data (see Fig. 7.4.1, Fig. 7.4.3 and Table 7.4.1). In addition, our results showed that SRR-pCASL consistently resulted in higher average SSIM and PSNR

values compared to C-pCASL and BASIL (see Table 7.4.1), indicating a better perceived image quality compared to the ground truth CBF map. Next, it was also demonstrated that the SRR-pCASL acquisition strategy, when combined with a MAP estimator, provides high resolution CBF maps with a more uniform and on average higher precision than CBF maps obtained with C-pCASL and BASIL (see Table 7.4.1, Fig. 7.4.1, and Fig. 7.4.5). This increased precision can be attributed to a two-fold gain in SNR, as provided by SRR-pCASL. First, SRR images are acquired with a low through-plane resolution, which increases the SNR as signal strength scales with the slice thickness (Delbany et al., 2019). Second, using a low through-plane resolution reduces the number of slices that need to be acquired to cover the same field-of-view (FOV) compared to higher through-plane resolution. As a result, for each LR image, the average effective PLD is shorter and the average level of BS improves compared to standard 2D MS acquisitions (see Fig. 7.5.1), which boosts the SNR throughout the entire volume. By augmenting each method with a MB factor of 2, which accelerates image acquisition and hence provides a more constant and on average better BS as well as a more constant PLD across slices, the SNR gain could be further maximized resulting in an additional improvement in CBF parameter mapping precision and RMSE for each method (see Fig. 7.4.5, Fig. 7.4.3, Tables 7.4.1 and 7.4.2). As calculated from our simulations without motion, this SNR gain was maximal when SRR-pCASL was combined with MB, relative to the other methods, with the non-MB version of SRR-pCASL even outperforming the MB versions of BASIL and C-pCASL in terms of SNR gain (see Table 7.4.2). Apart from the improved SNR, and although the labeling duration and the PLD inherently account for most of the time of the pCASL sequence, another advantage of lowering the spatial resolution in SRR-pCASL readout is a reduction of the scan time of an individual image, allowing to acquire more images within a certain amount of time compared to the acquisition of images with high spatial resolution.



Figure 7.5.1: A schematic representation of 2D MS readout with a high (left) and a low (right) through-plane resolution, both with an ascending acquisition order, as recommended for ASL (Alsop et al., 2015). Assuming the acquisition of an HR and an LR slice take up the same amount of scan time, the highlighted slices in green in both readout schemes will have the exact same effective PLD and level of BS. Regardless of the difference in SNR due to the difference in spatial resolution, the overall shorter effective PLD and the overall higher level of BS in the LR readout scheme will result in a higher SNR of the ASL signal on average throughout the brain.

7.5.3 Joint estimation of CBF and motion parameters

Besides the potential of SRR-pCASL to improve the traditional SNR/resolution/scan-time trade-off in ASL, the integration of a motion model and the single-PLD quantification model in the proposed SRR-pCASL reconstruction framework allows motion parameters and CBF parameters to be estimated directly and simultaneously. As such, a conventional two-step approach is avoided where HR perfusion-weighted images are reconstructed prior to voxel-wise quantification of CBF values. This benefit of the joint estimation of the motion and CBF parameters to avoid propagating errors originating from pre-registration routines (Beirinckx et al., 2022), was confirmed in the simulation experiment with motion, where it was shown that the proposed framework using joint motion estimation outperformed BASIL with FSL's mcflirt as a pre-registration routine in terms of motion component RMMSE (see Table 7.4.3). We also observed slightly better RMMSE values for C-pCASL than for SRR-pCASL, even when both approaches used the same Bayesian optimization framework with the same tolerance settings and regularization parameter selection. This may be the result of a combination of factors including cost function complexity, i.e. large-angle rotations are involved the SRR-pCASL forward model to compensate for the rotational acquisition, and sub-optimal hyper-parameter selection, but further research is required to investigate this discrepancy. Nevertheless, SRR-pCASL consistently resulted in lower CBF mapping RMSE values compared to C-pCASL in the simulations with motion (see Fig. 7.4.9, Fig. 7.4.10, and Table 7.4.3), indicating that the benefits of an SRR acquisition with rotated slice-encoding direction and low through-plane resolution, i.e. more optimal BS and more constant PLD across slices, provides a gain in precision that can compensate for small remaining inaccuracies in motion estimation compared to C-pCASL.

Although a thorough evaluation of BASIL and FSL's *mcflirt* was not the main scope of this contribution, it is worth emphasizing that default usage of these tools on motion corrupted pCASL data should be done with some precaution. As discussed in Appendix 7.E, for one control-label image pair of the simulated HR pCASL data sets corrupted with realistic motion, we noticed consistent outliers in the motion parameter component estimation that resulted from the pre-registration routine using *mcflirt*. This forced us to consider this image pair as an outlier and to discard it from the CBF quantification using BASIL. The obligatory use of outlier correction steps, and the added question on the basis of which criteria such correction steps should be evaluated, constitutes another unnecessary step towards accurate CBF quantification. The advantage of estimating CBF and motion parameters in a single integrated approach without any form of outlier correction, as proposed in this contribution, offers clear added value in that respect.

Further, we also observed an increased adaptive smoothing of BASIL in the simulations with motion, which resulted in an increased rSTD precision measure for BASIL (see Fig. 7.4.9, Fig. 7.4.13, and Table 7.4.3). As a result, BASIL and BASIL-MB outperformed the proposed approaches in terms of SNR gain (see Table 7.4.4). However, as indicated by our results, this over-smoothing of BASIL resulted in a significant reduction in accuracy, as confirmed by the arBias measure (see Fig. 7.4.9, Fig. 7.4.12, and Table 7.4.3), and a significant deterioration of BASIL in terms of RMSE compared to the simulation experiment without motion (see Fig. 7.4.10). In conclusion, similar to the simulations without motion, SRR-pCASL clearly outperformed BASIL and C-pCASL in terms of overall RMSE (see Fig. 7.4.9, Fig. 7.4.10) and Table 7.4.3).

Following the extensive simulation studies, the SRR-pCASL framework was also validated on in vivo brain data, and compared to a single-PLD pCASL experiment on conventional pCASL data acquired with MB directly at HR (see Table 7.3.2). When comparing SRRpCASL to the conventional HR pCASL experiment, two aspects stand out. First, the CBF maps reconstructed using SRR-pCASL (Fig. 7.4.16, top row), are comparable in terms of visualized anatomical structures to the CBF maps obtained from the conventional MS data using BASIL (Fig. 7.4.16, bottom row), and C-pCASL (Fig. 7.4.16, middle row), the last two approaches using data directly acquired at HR. Second, in certain slices (i.e. slices A-B of Fig. 7.4.16), SRR-pCASL appears to even outperform the conventional HR pCASL experiment in terms of reconstructing the underlying anatomy. This is a direct consequence of the benefit of acquiring LR data for SRR-pCASL in terms of SNR. Furthermore, the reconstructed slices of the CBF map shown in Fig. 7.4.16 for the SRR-pCASL experiment all have comparable signal intensities. This reflects the relative uniformity in average SNR throughout all regions in the brain related to the SNR benefits of acquiring LR pCASL images while using a rotational acquisition strategy. For the CBF map obtained from the conventional HR MS data, slices A-B shown in Fig. 7.4.16 clearly suffer from the low SNR due to long effective PLDs and limited BS. Similar as observed in the simulation experiment with motion, the adaptive regularization of BASIL-MB compensates for the low SNR by over-smoothing these slices A-B of the estimated CBF map, which were both acquired at the end of the first multiband (see Fig. 7.4.2). In comparison to BASIL-MB, such an over-smoothing effect was not observed in the CBF maps reconstructed with C-pCASL-MB (which uses a Laplacian prior) from the same conventional HR pCASL data set, where higher CBF intensities in the first multiband could be observed (see the coronal and sagittal views in Fig. 7.4.15). Moreover, the stability of the CBF values across slices for SRR-pCASL (see Fig. 7.4.16) is consistent with the uniform precision of CBF estimation from LR MS data using SRR-pCASL shown in the simulation experiment (see e.g. Fig. 7.4.9 and Fig. 7.4.13). Furthermore, the higher quality of the CBF map obtained from the SRR-pCASL experiment compared to that of the HR ASL experiment in regions of the brain that were imaged latest within the multiband segment (see the first two slices shown in Fig. 7.4.16) matches with the difference in RMSE of CBF estimation between both methods in those same regions as predicted in the simulation experiment (see Fig. 7.4.10). While it is difficult to compare a qualitative assessment (real data) with a quantitative measurement (simulation data), it is reasonable to assume both effects are correlated. It serves as an indication of the validity of the simulation experiment. Moreover, a qualitatively comparable CBF parameter mapping occurs for C-pCASL compared to BASIL, clearly indicative of the validity of using C-pCASL as a benchmark against SRR-pCASL in our in vivo experiment.

7.5.4 Model assumptions and prospective extensions.

It is worth highlighting that the proposed SRR framework is generic as other ASL models can be incorporated analogously. The current framework adopts the CBF quantification model prescribed by Alsop et al. (2015), which assumes that all labeled blood has arrived in the imaging voxel before the start of the readout (i.e., ATT from the labeling plane to the readout slice is assumed to be lower than the PLD), and has stayed intravascularly while decaying with the T1 relaxation time of blood. However, it has been shown that such model is sensitive to variations in the ATT of the labeled blood (Alsop & Detre, 1996), implying the need for sufficiently long PLD. If the PLD is shorter than the ATT, the risk of macrovascular artifacts (i.e., labeled blood in proximal arteries rather than the distal capillaries or tissue)

and ASL signal void (i.e., the labeled blood has not arrived yet in distal voxels resulting in ASL signal loss in these voxels) will increase. To address such issues, future work could focus on the extension of the proposed model-based SRR framework with a (nonlinear) multi-PLD pCASL model. Such a model allows to estimate the ATT parameter to improve the accuracy of CBF estimation (Buxton et al., 1998), and has already been applied for direct CBF mapping in combination with model-based reconstruction (Maier et al., 2021). Furthermore, a preliminary phantom simulation study in which the potential of SRR with multi-PLD pCASL was explored, has already been published as a conference proceeding (Bladt et al., 2017). In such an approach, each LR MS image is characterized by a unique slice orientation as well as a unique PLD time, allowing the estimation of both CBF and ATT. However, given the nonlinear nature of the multi-PLD pCASL model as well as the increased number of parameters to be estimated, the computational cost of the optimization is expected to increase. Note, also, that although such ASL model extensions may improve the accuracy of the estimated parameter maps, the addition of extra parameters to be estimated comes at the expense of a reduced precision. Furthermore, it is often assumed that the reduction in data averaging when using multiple time point protocols (required when acquiring the data in a matched scan time with a single-PLD protocol) leads to a reduction in the precision of the CBF estimates (Alsop et al., 2015; Dai et al., 2017; Teeuwisse et al., 2014), which could outweigh the benefits of correcting for ATT effects. At this point, future work could focus on investigating to what extent this reduction in CBF precision can be compensated for by the gain in precision associated with the use of SRR-pCASL. Overall, the effects of accuracy and precision should be carefully weighed against each other.

In this contribution, the calibration image ρ was acquired from a separate acquisition at the target resolution, i.e., the resolution of the reconstructed ϑ_{rCBF} map, followed by a multi-modal registration step to align $\vartheta_{\mathsf{rCBF}}$ and ρ , which then allowed absolute CBF quantification via voxel-wise division of ϑ_{rCBF} and ρ_{reg} to a HR CBF parameter map ϑ_{CBF} . An alternative approach could be to integrate the registration with the calibration image ρ as part of the cyclic block-coordinate descent optimization scheme in section 7.2.3.2, resulting in an additional set of motion parameters $heta_{
ho} \in \mathbb{R}^{6 imes 1}$ to be estimated simultaneously with $\{\vartheta, \theta\}$. This approach would avoid a separate calibration step, and could reduce a potential bias from propagating registration errors. However, similar to the addition of a multi-PLD model, the effects of accuracy and precision should be carefully considered as the estimation of additional motion parameters could come at the cost of a reduced precision. Furthermore, the benefit of simultaneous estimation of θ_{ρ} must be weighed against the extra computational cost associated with such estimation. In particular, modeling θ_{ρ} requires the introduction of an extra (computationally intensive) image warping operator in the numerator of the single-PLD pCASL signal model in Eq. (7.2.2). While image operators M_{θ_n} and G_n are currently combined in one warping operation to maximize computational efficiency, an additional image warping operator for $heta_
ho$ would demand a full operator call for every forward pass in the cost function. The latter would cause a significant increase in the computational cost of the framework. Therefore, in this contribution, the decision was made to consider the calibration step separately. Future work is encouraged to further investigate the impact of joint calibration on CBF quantification and computation time.

7.5.5 Limitations.

In our current implementation, the signal in the voxels of the LR images is assumed to be Gaussian distributed, which is a valid assumption for sufficiently high SNR. The MAP estimation then becomes a least-squares optimization, which can be solved efficiently as a linear optimization problem. However, other data distributions may apply when the SNR condition is not met, or when other coil acquisition setups, e.g. using parallel imaging, are used. In that case, the log likelihood can take a nonlinear form, resulting in a nonlinear optimization problem that requires more advanced computational solvers. Future work could focus on investigating the impact on CBF estimation when other data distributions, as well as spatial and temporal variations in the noise standard deviation maps of the LR images, are used.

The comparison between the considered CBF estimation frameworks depends on the imposed prior information on the unknown parameter maps, as well as the tuning of the associated regularization weights. While the SRR-pCASL framework uses a Laplacian prior (cfr. Section 7.2.3.1), and regularization weights can be chosen equally for SRR-pCASL and C-pCASL, BASIL uses an adaptive non-local spatial smoothing prior based on CBF variations in the brain, without a priori user-selected regularization weights. In order to compare the performance of the different methods in a reliable manner, ideally, the level of regularization should be the same in all reconstructions. At the same time, tuning regularization weights of a multiple regularization parameter selection problem as proposed, remains a difficult problem. In the current contribution, the hyper-parameters of the prior distributions in Eq. (7.2.8) were heuristically selected to be minimally intrusive in the reconstruction, balancing the trade-off between the data consistency objective and the regularization objectives of the tissue parameter maps. This approach may be sub-optimal.

A more fair comparison between SRR-pCASL, C-pCASL-MB and BASIL-MB in the real data experiments would be achieved when MB was (not) used in all experiments. One could argue that the current real data comparison was skewed in favour of the conventional HR ASL experiment, because MB was only used for HR ASL data acquisition. However, this choice was made for two reasons. On the one hand, a MB factor of 2 was used in the HR ASL experiment, as otherwise there would have been practically no ASL signal remaining in most of the upper part of the brain. As this was a proof-of-concept study, being able to verify whether the SRR-pCASL reconstructed CBF maps showed anatomical details comparable to those of BASIL-MB was more important than a true one-to-one comparison of the conventional HR ASL and SRR-pCASL experiment. On the other hand, MB was not used in the SRR-pCASL experiment, because MB acquisition required a calibration scan to be performed before acquisition of each LR image. This would have taken up too much of the available total scan time. Future work that investigates whether the SRR acquisition strategy can be combined with MB more efficiently, to allow for a more fair comparison between both strategies, is highly encouraged. As demonstrated in our simulation study without motion, the use of MB in SRR-pCASL-MB provided an additional improvement in relative RMSE of about 10% compared to SRR-pCASL, largely contributed to by the increased estimation precision.

Finally, the SRR-pCASL protocol needs to be validated on more subjects in order to demonstrate its intra- and inter-subject robustness. Ideally, data should be acquired repeatedly in individual subjects, in order to be able to determine sample standard deviations for CBF map estimates, similarly to the analysis done in simulations. This would allow to quantify the performance of SRR-pCASL and C-pCASL in more detail. In addition, CBF map estimates obtained using our SRR-pCASL method for 2D MS readout should be further compared with results obtained using recommended segmented 3D readout schemes. We anticipate that motion robustness and through-plane blurring are two effects where the proposed SRR-pCASL method with joint motion estimation offers potential improvements over segmented 3D readout.

7.5.6 Anticipated clinical impact and future perspectives.

Our research also highlights some advantages that could potentially have profound clinical impact. The FOV of the LR pCASL data for SRR-pCASL provides wider coverage in the slice direction (192mm) compared to the conventional HR data (120mm), acquired in the same total scan time. On the one hand, this wider coverage is required for SRR because the entire brain has to be within the FOV for each rotation angle. On the other hand, it offers a potential advantage to applications where such large coverage is required. An advantage, for example, is that the labeling plane can be scanned, which is not standard for conventional ASL acquisitions. Directly visualizing the labeling plane can offer valuable information for off-resonance correction schemes that investigate B0 field inhomogeneity distortions in the labeling plane (Berry et al., 2019). In addition, the larger coverage would potentially allow perfusion quantification in deep brain structures such as the brain stem, cerebellum, and even the spinal cord (Shou et al., 2021). These brain structures are important nodes of the structural and functional networks of the human brain. However, to date, very few perfusion measurements have been performed in these structures.

Further, we anticipate that the combination of a rotated SRR acquisition strategy with MB imaging (also known as SMS) offers great promise to overcome some of the limitations of ultra-high-field (UHF; 7T and higher) ASL techniques (Teeuwisse et al., 2010). The added value of such a combined acquisition strategy seems twofold. First, image SNR increases both with field strength (Gardener et al., 2009), and due to acquisition with low through-plane resolution and increased effectiveness of BS associated with SRR and MB, as clearly demonstrated in this contribution (see also Fig. 7.5.1). Second, the combination of SRR with MB offers an attractive approach to increase the currently limited spatial coverage at these higher field strengths (Ivanov et al., 2017), since both several spatially distributed imaging slices are excited and the through-plane resolution of each slice is increased.

It is worth discussing whether the proposed SRR approach can be extended to pCASL with 3D readout. Super-resolution reconstruction is conventionally defined as the recovery of high-frequency components corrupted by aliasing (Kang & Chaudhuri, 2003). In 2D multi-slice imaging, aliasing occurs in the through-plane direction, which facilitates SRR. However, there is consensus that super-resolution in MRI is not achievable in-plane (Greenspan et al., 2002; Scheffler, 2002; Plenge et al., 2012), nor in true 3D acquisitions, since the Fourier encoding scheme excludes aliasing in frequency and phase encoding directions. Notwithstanding the aliasing condition, the proposed estimation framework is fully compatible with 3D pCASL data (provided that the slice selection profile is turned off). In that case, the reconstruction will mainly benefit from the joint estimation of CBF and motion parameters, while the potential resolution gain is expected to be marginal. At the same time, 3D readout remains subject to several disadvantages, including spatial blurring due to T2 decay and a high sensitivity to (intra-scan) motion. These disadvantages could complicate resolution enhancement using LR

3D ASL scans. In particular, it is crucial that such effects are included in the forward model for (iterative) reconstruction which connects the ground truth CBF map to the observed data.

Finally, the combination of the proposed SRR acquisition strategy with alternative ASL labeling approaches, such as Hadamard time encoding (in the context of multi-PLD pCASL) (Teeuwisse et al., 2014), and velocity selective encoding (Qin et al., 2022), seems worth investigating and is suggested as a possible extension of this work.

7.6 Conclusion

This contribution has introduced a model-based super-resolution reconstruction framework for single-PLD pCASL MRI, building on a joint Bayesian estimation framework that aims to estimate motion-corrected 3D isotropic high-resolution quantitative CBF maps from a set of 2D multi-slice control-label image pairs acquired with low through-plane resolution and rotated slice-encoding direction. The framework has been validated in synthetic whole brain simulations and on *in vivo* human brain data, demonstrating successful CBF quantification while providing a more uniform distribution of PLD, improved SNR, and increased effectiveness of BS compared to conventional 2D MS readout with ascending slice order and isotropic resolution in the same scan time, even when multiband is applied in the latter. By improving upon existing disadvantages of 2D MS readout, the proposed framework provides a promising alternative to the recommended segmented 3D readout schemes, which to date remain sensitive to inter-shot motion and through-plane blurring due to T2 decay along the long echo trains.

Acknowledgments

The authors gratefully acknowledge support from the Research Foundation Flanders, Belgium (FWO) through project funding G084217N and 1S69918N. Furthermore, the authors would like to thank Sophie Schmid and Wouter M. Teeuwisse for their help during the acquisition of the *in-vivo* datasets.

Appendices

7.A Slice-dependent PLD

When using a conventional 2D multi-slice readout strategy for pCASL with subsequently acquired slices, distal slices will have longer effective PLDs than proximal slices. This slice-dependent effective PLD can be represented using a function that maps each ordered triplet of HR grid coordinates to a voxel value holding the effective PLD, i.e., **PLD**: $D \subset \mathbb{N}_0^3 \mapsto \mathbb{R}$, where $D = \{(i, j, k) \mid i = 1, ..., u; j = 1, ..., v; k = 1, ..., w; and <math>u, v, w \in \mathbb{N}_0\}$, with $N_r = u \times v \times w$. Assuming that the base PLD value, PLD_{base}, increases by a multiple of the slice readout time, t_{read} , for ascending slices in a conventional 2D multi-slice pCASL acquisition, the $(i, j, k)^{\text{th}}$ voxel value of **PLD** is defined as:

$$\mathsf{PLD}(i, j, k) = \mathsf{PLD}_{\mathsf{base}} + t_{\mathsf{read}} \cdot h(k) \quad , \tag{7.A.1}$$

243

with $h: \mathbb{N}_0 \to \mathbb{N}$, $h = \{h(k) \mid k = 1, ..., w \text{ and } w \in \mathbb{N}_0\}$, a function that defines the integer multiplication factor for the k^{th} slice of **PLD**:

$$h(k) = \frac{1}{\zeta} \left[(k-1) \mod \varrho - ((k-1) \mod \varrho) \mod \zeta \right], \qquad (7.A.2)$$

where ζ is equal to the anisotropy factor, defined as the ratio of the through-plane resolution to the in-plane resolution, and with ρ the number of HR slices per band, i.e. the ratio of the total number of HR slices to the multiband factor. The application of function h is illustrated in Fig. 7.A.1.



Figure 7.A.1: Illustration of Eq. (7.A.2) where function h operates on a 1-dimensional HR grid vector k = 1, ..., 12, oriented along the slice-encoding dimension, using an anisotropy factor $\zeta = 3$, with $\rho = 12$ (i.e., no SMS), and with $\rho = 6$ HR slices per band when SMS with multiband factor 2 is modelled.

When using a SRR acquisition strategy, 2D multi-slice images are acquired with anisotropic voxel size, where each LR image samples the HR scene in a distinct fashion to ensure that the acquired data contains complementary information about the HR image or HR parameter maps to be reconstructed. In this contribution, the LR images are acquired with varying slice-encoding directions (Fig. 7.2.1). Consequently, effective PLD values will vary according to the assumed slice-encoding direction of each LR image. Under the assumption that no labeling of cerebral blood is present at an infinitely long PLD, the effective PLD values of the n^{th} HR pCASL image r_n can be modelled using a function $PLD_n = \{PLD_{nj}\}_{i=1}^{N_r \times 1} \in \mathbb{R}^{N_r \times 1}$:

$$\mathbf{PLD}_n = \begin{cases} \infty, & \text{if } n \text{ is odd} \\ M_{\theta_n}^{-1} G_n^{-1} \mathbf{PLD}, & \text{if } n \text{ is even} \end{cases},$$
(7.A.3)

where $M_{\theta_n}^{-1} \in \mathbb{R}^{N_r \times N_r}$ and $G_n^{-1} \in \mathbb{R}^{N_r \times N_r}$ denote the exact inverse warping operators of operators M_{θ_n} and G_n , respectively, which are required to anticipate image warping in the super-resolution forward model (7.2.1). Note that in this work, similar to (Ramos-Llordén et al., 2017; Beirinckx et al., 2022), image warping is implemented very efficiently with Fast Fourier Transforms (FFT). With the FFT approach, M_{θ_n} (or G_n) can be shown to be unitary, which means that its inverse $M_{\theta_n}^{-1}$ is equal to $M_{\theta_n}^H \in \mathbb{R}^{N_r \times N_r}$, where the superscript H denotes the adjoint or Hermitian conjugate. Hence, the warping operator M_{θ_n} is easily reversible, i.e. when applied to an image, the image can be retrieved by applying $M_{\theta_n}^H$ to the output of this operation.

7.B Background suppression

Background suppression (BS) can be used to increase the SNR of the ASL signal by suppressing the physiological noise component that scales with the signal intensity in the

label and control images. BS can be achieved using a combination of a saturation pulse and a certain number of inversion pulses applied to the imaging volume (Garcia et al., 2005; Maleki et al., 2012). By timing these inversion pulses correctly with the readout excitation, the longitudinal magnetization of the static background tissue will pass through zero at the time of readout. For imaging methods that employ a single excitation per repetition time (TR), such as the segmented 3D approaches, BS can be highly effective, as the null point of the magnetization can be timed to coincide with the excitation pulse. However, in 2D multi-slice readout for pCASL, used for SRR, an excitation pulse is used for each individual slice and slices are acquired subsequently resulting in different slice acquisition times. As a result, BS can be optimal for one slice, but is progressively less efficient for other slices.

Under the assumption that BS is perfect for the first acquired slice and that signal in subsequent slices recovers towards equilibrium with T_1 of tissue, the slice-dependent variation of optimal inversion time points for BS in a conventional 2D multi-slice pCASL acquisition with ascending slice order, can be represented as a function TI: $D \subset \mathbb{N}_0^3 \mapsto \mathbb{R}$, where $D = \{(i, j, k) \mid i = 1, ..., u; j = 1, ..., v; k = 1, ..., w; and <math>u, v, w \in \mathbb{N}_0\}$, with $N_r = u \times v \times w$. Similar to the definition of slice-dependent PLD values in Eq. (7.A.1), it is assumed that inversion times increase by a multiple of the readout time per slice for ascending slice numbers. As such, the (i, j, k)-th voxel value of TI is defined as:

$$TI(i, j, k) = T_1(i, j, k) \cdot \ln(2) + t_{read} \cdot h(k) , \qquad (7.B.1)$$

with h following the same definition as in Eq. (7.A.2).

For a SRR acquisition, the optimal inversion times for perfect background suppression of each slice will depend on the corresponding slice-encoding direction of each separate acquisition. Therefore, for each HR pCASL image r_n , the corresponding $\Delta_n = \{\mathsf{TI}_{nj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ can be modelled as:

$$\mathsf{TI}_n = \boldsymbol{M}_{\boldsymbol{\theta}_n}^{-1} \boldsymbol{G}_n^{-1} \mathsf{TI} \quad . \tag{7.B.2}$$

Next, let $\mathbf{b}_n = \{b_{nj}\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ represent a vector that models the T_1 relaxation factor for inversion-recovery nulling for BS, assuming TR $\gg T_1$ and a perfect 180° RF inversion pulse (Barral et al., 2010), with

$$b_{nj} = 1 - 2 \cdot \exp\left(-\frac{\mathsf{TI}_{nj}}{\mathsf{T}_{1,j}}\right).$$
(7.B.3)

Then, Eq. (7.2.4) can be extended to include the effect of background suppression:

$$r_{nj} = \begin{cases} r_{1,j}b_{nj}, & \text{if } n \text{ is odd} \\ r_{1,j}b_{nj} - \Delta r_{nj}, & \text{if } n \text{ is even.} \end{cases}$$
(7.B.4)

7.C Linear forward model

It follows from Eq. (7.2.4) that the HR images \boldsymbol{r}_n can be modelled as a linear function of the parameter vector $\boldsymbol{\vartheta} = [\boldsymbol{r}_1^T \ \boldsymbol{\vartheta}_{rCBF}^T]^T \in \mathbb{R}^{2N_r \times 1}$:

$$\boldsymbol{r}_n(\boldsymbol{\vartheta}) = \boldsymbol{A}_n \boldsymbol{\vartheta} \tag{7.C.1}$$

where $A_n \in \mathbb{R}^{N_r \times 2N_r}$ represents the block matrix operator:

$$\boldsymbol{A}_{n} = \begin{cases} \begin{bmatrix} \boldsymbol{I}_{N_{r}} & \boldsymbol{0}_{N_{r}} \end{bmatrix}, & \text{if } n \text{ is odd} \\ \begin{bmatrix} \boldsymbol{I}_{N_{r}} & \text{diag}(\boldsymbol{v}_{n}) \end{bmatrix}, & \text{if } n \text{ is even }, \end{cases}$$
(7.C.2)

whose matrix elements are given by

$$\boldsymbol{A}_{n} = \begin{cases} \begin{bmatrix} 1 & 0 & \dots & 0 & 0 & 0 & \dots & 0 \\ 0 & 1 & \ddots & \vdots & 0 & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 & 0 \end{bmatrix}, & \text{if } n \text{ is odd} \\ \begin{bmatrix} 1 & 0 & \dots & 0 & v_{n,1} & 0 & \dots & 0 \\ 0 & 1 & \ddots & \vdots & 0 & v_{n,2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 & v_{n,N_{r}} \end{bmatrix}, & \text{if } n \text{ is even }, \end{cases}$$
(7.C.3)

with $\mathbf{I}_{N_r} \in \mathbb{R}^{N_r \times N_r}$ the identity matrix, $\mathbf{0}_{N_r} \in \mathbb{R}^{N_r \times N_r}$ the zero matrix, $\operatorname{diag}(\mathbf{v}_n) \in \mathbb{R}^{N_r \times N_r}$ a diagonal matrix with the elements of $\mathbf{v}_n = \{v_{nj}\}_{j=1}^{N_r} = \{-\Delta r_{nj}/\vartheta_{r\text{CBF},j}\}_{j=1}^{N_r} = \{-\delta^{-1} \exp\left(-\operatorname{PLD}_{n,j}/T_{1b}\right)\}_{j=1}^{N_r} \in \mathbb{R}^{N_r \times 1}$ on its diagonal. Consequently, when combining Eq. (7.C.1) with the linear SRR forward model operators in Eq. (7.2.1), the overall forward model in SRR-pCASL remains linear, which allows for efficient solving of (P.1) using linear optimization routines.

Table 7.C.1: Acquisition settings for the synthetic data set using 2D MS readout, that was used for the additional simulation experiment with adjusted resolution. A slice orientation angle of 0° corresponds with the slice-encoding axis directed from left to right, and with the phase-encoding axis perpendicularly directed from anterior to posterior. Each angle listed below is a rotation of the slice-encoding axis around the phase-encoding direction counterclockwise. Therefore, a 90° angle is consistent with an ascending slice order. These rotations are consistent with the rotations visualized in Fig. 7.2.1.

	Extra Dataset	
	LR 2D MS	
Number of slices per slab N _{slice}	12	
Acquisition matrix	120×120	
FOV [mm ³]	$240 \times 240 \times 192$	
Voxel size [mm ³]	$2 \times 2 \times 16$	
Labeling duration $ au$ [ms]	1800	
PLD _{base} [ms]	1800	
PLD range [ms]	1800-2350	
Number of control-label pairs N	24	
Number of slice encoding directions	24	
Slice orientation angles [°]	0, 7.5,, 172.5	
Multiband factor ω	n.a.	

7.D Extra simulation experiment

In order to demonstrate that the potential of SRR is not confined to a particular resolution, an additional simulation experiment was performed where a $2 \times 2 \times 2$ mm³ CBF map was super-resolution reconstructed from LR images with a resolution of $2 \times 2 \times 16$ mm³. Acquisition settings for this addition simulation experiment are summarized in Table 7.C.1. Fig. 7.D.1 shows the result of this extra simulation experiment side-by-side with the original simulation experiment where a $3 \times 3 \times 3$ mm³ CBF was reconstructed. No unwanted patient motion was simulated for this experiment.



Figure 7.D.1: Orthogonal slice views of simulated LR 2D control images with low through-plane resolution compared to the HR CBF map estimates reconstructed with SRR-pCASL, for the original acquisition protocol using LR images with a resolution of $3 \times 3 \times 12$ mm³ as input (left), and the new acquisition protocol using LR images with a resolution of $2 \times 2 \times 16$ mm³ as input (right), as summarized in Table 7.C.1. For illustration purposes, only the slice orientation angles corresponding with a slice orientation of 0° and 90° are shown for the 2D LR control images (see also Fig. 7.2.1).

7.E Motion parameter estimates

The motion parameter estimates for the simulation experiment with motion, both without and with the use of multiband, are summarized in Figs. 7.E.1-7.E.2. Each figure shows the true reference motion component values $\theta_n = {\{\theta_{nk}\}}_{k=1}^6$ being used to corrupt each image number, $n = 1 \dots 2N$, in the whole brain simulations with added motion. Next, Fig. 7.E.1 and Fig. 7.E.2 also show the mean motion parameter component $\overline{\theta}_{nk}$, where the mean was calculated over the N_{MC} estimates for each image number *n*. In addition, for each component and framework combination, the RMSE was plotted using a barplot, where the RMSE value

was calculated per motion component as $\text{RMSE}(\theta_{nk}) = \left(\overline{(\theta_{nk} - \overline{\theta}_{nk})^2}\right)^{\frac{1}{2}}$, where $\overline{\theta}_{nk}$ denotes the sample mean of the N_{MC} estimates of the motion component k for image number n, and where $\overline{(\cdot)}$ denotes the element-wise sample mean operator over the N_{MC} estimates. Note that, in line with the definitions in Section 7.2.2, pCASL images r_n were pairwise ranked in alternating order as control-label-control-label-.... As highlighted by the red bars in Fig. 7.E.1 and Fig. 7.E.2, the control-label image pair corresponding with image numbers 39 and 40 was discarded from the CBF quantification routine in BASIL and BASIL-MB due to consistent outliers of the estimated motion parameter components that resulted from the pre-registration routine using FSL's *mcflirt*.



Figure 7.E.1: Graphs of the mean motion component estimate and associated RMSE, for each image number and respective framework **without the use of multiband**, calculated over the $N_{MC} = 100$ results for the whole brain simulation experiment with motion. The true reference values for each motion component are shown in column 1. The outlier discarded control-label image pair for BASIL is annotated in red.



Figure 7.E.2: Graphs of the mean motion component estimate and associated RMSE, for each image number and respective framework **using multiband**, calculated over the $N_{MC} = 100$ results for the whole brain simulation experiment with motion. The true reference values for each motion component are shown in column 1. The outlier discarded control-label image pair for BASIL-MB is annotated in red.

References

- Alsop, D. C., & Detre, J. A. (1996). Reduced Transit-Time Sensitivity in Noninvasive Magnetic Resonance Imaging of Human Cerebral Blood Flow. J. Cereb. Blood Flow Metab., 16(6), 1236–1249.
- Alsop, D. C., et al. (2015). Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. *Magn. Reson. Med.*, 73(1), 102–116.
- Bano, W., Piredda, G. F., Davies, M., Marshall, I., Golbabaee, M., Meuli, R., Kober, T., Thiran, J. P., & Hilbert, T. (2020). Model-based super-resolution reconstruction of T2 maps. *Magn. Reson. Med.*, *83*(3), 906–919.
- Barral, J. K., Gudmundson, E., Stikov, N., Etezadi-Amoli, M., Stoica, P., & Nishimura, D. G. (2010). A robust methodology for in vivo T1 mapping. *Magn. Reson. Med.*, *64*(4), 1057–1067.
- Beck, A., & Tetruashvili, L. (2013). On the convergence of block coordinate descent type methods. *SIAM J. Optim.*, 23(4), 2037–2060.
- Beirinckx, Q., Jeurissen, B., Nicastro, M., Poot, D. H. J., Verhoye, M., den Dekker, A. J., & Sijbers, J. (2022). Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping. *Comput. Med. Imaging. Graph.*, 100, 102071.
- Beirinckx, Q., Ramos-Llordén, G., Jeurissen, B., Poot, D. H. J., Parizel, P. M., Verhoye, M., Sijbers, J., & den Dekker, A. J. (2020). Joint Maximum Likelihood Estimation of Motion and T1 Parameters from Magnetic Resonance Images in a Super-resolution Framework: a Simulation Study. *Fundam. Inform.*, 172, 105–128.
- Berry, E. S., Jezzard, P., & Okell, T. W. (2019). Off-resonance correction for pseudo-continuous arterial spin labeling using the optimized encoding scheme. *NeuroImage*, *199*, 304–312.
- Bladt, P., Beirinckx, Q., van der Plas, M. C. E., Schmid, S., Teeuwisse, W. M., Jeurissen, B., den Dekker, A. J., Sijbers, J., & van Osch, M. J. P. (2020). Super-resolution reconstruction of single-PLD pseudo-continuous ASL images. In *Proc. Intl. Soc. Mag. Reson. Med.*, vol. 28, (p. 3293).
- Bladt, P., Beirinckx, Q., Van Steenkiste, G., Jeurissen, B., Achten, E., den Dekker, A. J., & Sijbers, J. (2017). Super-resolution multi-PLD PCASL: a simulation study. *Magn. Reson. Mater. Phy.*, 30 (Suppl. 1)(S396).
- Buxton, R. B., Frank, L. R., Wong, E. C., Siewert, B., Warach, S., & Edelman, R. R. (1998). A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magn. Reson. Med.*, 40(3), 383–396.
- Chappell, M. A., Groves, A. R., Whitcher, B., & Woolrich, M. W. (2009). Variational bayesian inference for a nonlinear forward model. *IEEE Trans. Signal Process.*, 57(1), 223–236.
- Clement, P., Petr, J., Dijsselhof, M. B. J., Padrela, B., Pasternak, M., Dolui, S., Jarutyte, L., Pinter, N., Hernandez-Garcia, L., Jahn, A., Kuijer, J. P. A., Barkhof, F., Mutsaerts, H. J. M. M., & Keil, V. C. (2022). A Beginner's Guide to Arterial Spin Labeling (ASL) Image Processing. *Front. Radio.*, *2*.
- Coleman, T. F., & Li, Y. (1994). On the convergence of reflective Newton methods for large-scale minimization subject to bounds. *Math. Program.*, 67(2), 189–224.
- Dai, W., Fong, T., Jones, R. N., Marcantonio, E., Schmitt, E., Inouye, S. K., & Alsop, D. C. (2017). Effects of arterial transit delay on cerebral blood flow quantification using arterial spin labeling in an elderly cohort. J. Magn. Reson. Imaging, 45(2), 472–481.
- Delbany, M., Bustin, A., Poujol, J., Thomassin-Naggara, I., Felblinger, J., Vuissoz, P.-A., & Odille, F. (2019). One-millimeter isotropic breast diffusion-weighted imaging: Evaluation of a superresolution strategy in terms of signal-to-noise ratio, sharpness and apparent diffusion coefficient. *Magn. Reson. Med.*, 81(4), 2588–2599.
- den Dekker, A. J., & Sijbers, J. (2014). Data distributions in magnetic resonance images: A review. *Physica Med.*, 30(7), 725–741.
- Fan, A. P., Jahanian, H., Holdsworth, S. J., & Zaharchuk, G. (2016). Comparison of cerebral blood flow measurement with [¹⁵O]-water positron emission tomography and arterial spin labeling magnetic resonance imaging: A systematic review. J. Cereb. Blood Flow Metab., 36(5), 842–861.
- Fessler, J. A., & Kim, D. (2011). Axial block coordinate descent (ABCD) algorithm for X-ray CT image reconstruction. In Proc. Intl. Mtg. on Fully 3D image Recon. in Rad. and Nuc. Med., (pp. 262–265).
- Fogtmann, M., Seshamani, S., Kroenke, C., Cheng, X., Chapman, T., Wilm, J., Rousseau, F., & Studholme, C. (2014). A unified approach to diffusion direction sensitive slice registration and 3-D DTI reconstruction from moving fetal brain anatomy. *IEEE Trans. Med. Imag.*, 33(2), 272–289.

- Garcia, D. M., Duhamel, G., & Alsop, D. C. (2005). Efficiency of inversion pulses for background suppressed arterial spin labeling. *Magn. Reson. Med.*, 54(2), 366–372.
- Gardener, A. G., Gowland, P. A., & Francis, S. T. (2009). Implementation of quantitative perfusion imaging using pulsed arterial spin labeling at ultra-high field. *Magn. Reson. Med.*, 61(4), 874–882.
- Greenspan, H., Oz, G., Kiryati, N., & Peled, S. (2002). MRI inter-slice reconstruction using superresolution. Magn. Reson. Imaging, 20(5), 437–446.
- Groves, A. R., Chappell, M. A., & Woolrich, M. W. (2009). Combined spatial and non-spatial prior for inference on MRI time-series. *NeuroImage*, 45(3), 795–809.
- Hernandez-Garcia, L., Aramendía-Vidaurreta, V., Bolar, D. S., Dai, W., Fernández-Seara, M. A., Guo, J., Madhuranthakam, A. J., Mutsaerts, H., Petr, J., Qin, Q., Schollenberger, J., Suzuki, Y., Taso, M., Thomas, D. L., van Osch, M. J. P., Woods, J., Zhao, M. Y., Yan, L., Wang, Z., Zhao, L., & Okell, T. W. (2022). Recent Technical Developments in ASL: A Review of the State of the Art. *Magn. Reson. Med.*, *88*(5), 2021–2042.
- Ivanov, D., Poser, B. A., Huber, L., Pfeuffer, J., & Uludağ, K. (2017). Optimization of simultaneous multislice EPI for concurrent functional perfusion and BOLD signal measurements at 7T. *Magn. Reson. Med.*, 78(1), 121–129.
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2), 825–841.
- Kang, M. G., & Chaudhuri, S. (2003). Super-resolution image reconstruction. IEEE Signal Process. Mag., 20(3), 19–20.
- Krüger, G., & Glover, G. H. (2001). Physiological noise in oxygenation-sensitive magnetic resonance imaging. Magn. Reson. Med., 46(4), 631–637.
- Liu, F., Velikina, J. V., Block, W. F., Kijowski, R., & Samsonov, A. A. (2017). Fast Realistic MRI Simulations Based on Generalized Multi-Pool Exchange Tissue Model. *IEEE Trans. Med. Imag.*, 36(2), 527–537.
- Maier, O., Spann, S. M., Pinter, D., Gattringer, T., Hinteregger, N., Thallinger, G. G., Enzinger, C., Pfeuffer, J., Bredies, K., & Stollberger, R. (2021). Non-linear fitting with joint spatial regularization in arterial spin labeling. *Med. Image Anal.*, 71, 102067.
- Maleki, N., Dai, W., & Alsop, D. C. (2012). Optimization of background suppression for arterial spin labeling perfusion imaging. *MAGMA*, 25(2), 127–133.
- Nicastro, M., Jeurissen, B., Beirinckx, Q., Smekens, C., Poot, D. H. J., Sijbers, J., & den Dekker, A. J. (2022). To shift or to rotate? comparison of acquisition strategies for multi-slice super-resolution magnetic resonance imaging. *Front. Neurosci.*, 16(1044510).
- Parkes, L. M., Rashid, W., Chard, D. T., & Tofts, P. S. (2004). Normal cerebral perfusion measurements using arterial spin labeling: Reproducibility, stability, and age and gender effects. *Magn. Reson. Med.*, 51(4), 736–743.
- Paschoal, A. M., Leoni, R. F., Pastorello, B. F., & van Osch, M. J. P. (2021). Three-dimensional gradient and spin-echo readout for time-encoded pseudo-continuous arterial spin labeling: Influence of segmentation factor and flow compensation. *Magn. Reson. Med.*, 86(3), 1454–1462.
- Plenge, E., Poot, D. H. J., Bernsen, M., Kotek, G., Houston, G., Wielopolski, P., van der Weerd, L., Niessen, W. J., & Meijering, E. (2012). Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? *Magn. Reson. Med.*, 68(6), 1983–1993.
- Poot, D. H. J., Jeurissen, B., Bastiaensen, Y., Veraart, J., Van Hecke, W., Parizel, P. M., & Sijbers, J. (2013). Super-resolution for multislice diffusion tensor imaging. *Magn. Reson. Med.*, 69(1), 103–113.
- Poot, D. H. J., Van Meir, V., & Sijbers, J. (2010). General and efficient super-resolution method for multi-slice MRI. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, vol. 6361, (pp. 615–622).
- Qin, Q., Alsop, D. C., Bolar, D. S., Hernandez-Garcia, L., Meakin, J., Liu, D., Nayak, K. S., Schmid, S., van Osch, M. J. P., Wong, E. C., Woods, J. G., Zaharchuk, G., Zhao, M. Y., Zun, Z., Guo, J., & the ISMRMPerfusion Study Group (2022). Velocity-selective arterial spin labeling perfusion mri: A review of the state of the art and recommendations for clinical implementation. *Magn. Reson. Med.*, 88(4), 1528–1547.
- Ramos-Llordén, G., den Dekker, A. J., Van Steenkiste, G., Jeurissen, B., Vanhevel, F., Van Audekerke, J., Verhoye, M., & Sijbers, J. (2017). A unified maximum likelihood framework for simultaneous motion and T₁ estimation in quantitative MR T₁ mapping. *IEEE Trans. Med. Imag.*, 36(2), 433–446.

Scheffler, K. (2002). Superresolution in MRI? Magn. Reson. Med., 48(2), 408-408.

- Setsompop, K., Bilgic, B., Nummenmaa, A., Fan, Q., Cauley, S. F., Huang, S., Chatnuntawech, I., Rathi, Y., Witzel, T., & Wald, L. L. (2015). SLIce Dithered Enhanced Resolution Simultaneous MultiSlice (SLIDER-SMS) for high resolution (700 um) diffusion imaging of the human brain. In *Proc. Intl. Soc. Mag. Reson. Med.*, vol. 23, (p. 339).
- Shao, X., Wang, Y., Moeller, S., & Wang, D. J. (2018). A constrained slice-dependent background suppression scheme for simultaneous multislice pseudo-continuous arterial spin labeling. *Magn. Reson. Med.*, 79(1), 394–400.
- Shilling, R. Z., Robbie, T. Q., Bailloeul, T., Mewes, K., Mersereau, R. M., & Brummer, M. E. (2009). A super-resolution framework for 3-D high-resolution and high-contrast imaging using 2-D multislice MRI. *IEEE Trans. Med. Imag.*, 28(5), 633–644.
- Shou, Q., Shao, X., & Wang, D. J. J. (2021). Super-resolution arterial spin labeling using slice-dithered enhanced resolution and simultaneous multi-slice acquisition. *Front. Neurosci.*, 15(737525), 1–10.
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., De Luca, M., Drobnjak, I., Flitney, D. E., Niazy, R. K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J. M., & Matthews, P. M. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23, S208–S219.
- Teeuwisse, W. M., Schmid, S., Ghariq, E., Veer, I. M., & van Osch, M. J. (2014). Time-encoded pseudocontinuous arterial spin labeling: Basic properties and timing strategies for human applications. *Magn. Reson. Med.*, 72(6), 1712–1722.
- Teeuwisse, W. M., Webb, A. G., & van Osch, M. J. P. (2010). Arterial spin labeling at ultra-high field: All that glitters is not gold. *Int. J. Imaging. Syst. Technol.*, 20(1), 62–70.
- Van Dyck, P., Smekens, C., Vanhevel, F., De Smet, E., Roelant, E., Sijbers, J., & Jeurissen, B. (2020). Super-resolution magnetic resonance imaging of the knee using 2-dimensional turbo spin echo imaging. *Investig. Radiol.*, 55(8), 481–493.
- van Osch, M. J. P., Teeuwisse, W. M., Chen, Z., Suzuki, Y., Helle, M., & Schmid, S. (2018). Advances in arterial spin labelling MRI methods for measuring perfusion and collateral flow. *J. Cereb. Blood. Flow. Metab.*, 38(9), 1461–1480.
- Van Reeth, E., Tham, I. W., Tan, C. H., & Poh, C. L. (2012). Super-resolution in magnetic resonance imaging: A review. *Concepts Magn. Reson.*, 40A(6), 306–325.
- Van Steenkiste, G., Jeurissen, B., Veraart, J., den Dekker, A. J., Parizel, P. M., Poot, D. H. J., & Sijbers, J. (2016). Super-resolution reconstruction of diffusion parameters from diffusion-weighted images with different slice orientations. *Magn. Reson. Med.*, 75(1), 181–195.
- Van Steenkiste, G., Poot, D. H. J., Jeurissen, B., den Dekker, A. J., Vanhevel, F., Parizel, P. M., & Sijbers, J. (2017). Super-resolution T1 estimation: Quantitative high resolution T1 mapping from a set of low resolution T1-weighted images with different slice orientations. *Magn. Reson. Med.*, 77(5), 1818–1830.
- Vidorreta, M., Balteau, E., Wang, Z., De Vita, E., Pastor, M. A., Thomas, D. L., Detre, J. A., & Fernández-Seara, M. A. (2014). Evaluation of segmented 3D acquisition schemes for whole-brain high-resolution arterial spin labeling at 3T. *NMR Biomed.*, *27*(11), 1387–1396.
- Vidorreta, M., Wang, Z., Rodríguez, I., Pastor, M. A., Detre, J. A., & Fernández-Seara, M. A. (2013). Comparison of 2d and 3d single-shot asl perfusion fmri sequences. *NeuroImage*, *66*, 662–671.
- Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M. A., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M., & Smith, S. M. (2009). Bayesian analysis of neuroimaging data in FSL. *NeuroImage*, 45(1, Supplement 1), S173–S186.
- Ye, F. Q., Frank, J. A., Weinberger, D. R., & McLaughlin, A. C. (2000). Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). *Magn. Reson. Med.*, 44(1), 92–100.
- Zhang, K., Herzog, H., Mauler, J., Filss, C., Okell, T. W., Kops, E. R., Tellmann, L., Fischer, T., Brocke, B., Sturm, W., Coenen, H. H., & Shah, N. J. (2014). Comparison of Cerebral Blood Flow Acquired by Simultaneous [¹⁵O]Water Positron Emission Tomography and Arterial Spin Labeling Magnetic Resonance Imaging. J. Cereb. Blood Flow Metab., 34(8), 1373–1380.

8

Conclusions and Future Perspectives

The overall goal of this dissertation was to develop improved methods for **motion-robust quantitative magnetic resonance imaging (qMRI) of the brain**, specifically focusing on the use of model-based super-resolution reconstruction (SRR) as a technology to optimize the trade-off between spatial resolution, signal-to-noise ratio, and scan time in qMRI.

As outlined in the Prologue, our society is experiencing rapid growth and aging, resulting in an increased prevalence of neurodegenerative disorders and age-related diseases, underscoring the urgent need for timely disease detection to impede or delay their progression. This demand emphasizes the **need for reliable and easily accessible quantitative biomarkers** capable of identifying diseases before clinical symptoms manifest. MRI, renowned for its excellent soft tissue contrast and inherent patient safety, is a preferred biomedical imaging tool. However, its widespread use as a biomarker detection tool encounters one particular challenge. Conventional MRI relies on qualitative image contrast evaluation, complicating the quantitative comparison of tissue properties within and between scans or subjects. Transitioning to quantitative MRI (qMRI), as emphasized in Chapter 3, is crucial to overcome these limitations, enabling absolute quantification of tissue characteristics independent of experimental design, thereby enhancing diagnostics.

Unfortunately, the dissemination of qMRI faces challenges such as **low spatial resolution**, **low signal-to-noise ratio (SNR)**, and long scan times. These long scan times, required to compensate for the low SNR and spatial resolution, can impact patient comfort and compliance, increase the risk of motion artifacts, and reduce patient throughput. To address the need for rapid MRI techniques, without compromising the spatial resolution or SNR of the MR scans, this thesis employs model-based super-resolution reconstruction as an advanced image reconstruction tool. As discussed in Chapter 4, SRR addresses certain deficiencies (blur, SNR loss, etc.) evident in existing MRI resolution enhancement techniques. Additionally, the presence of aliased frequency content along the throughplane dimension of 2D multi-slice MRI acquisitions serves as a foundational prerequisite for the application and efficacy of SRR. Furthermore, as also highlighted in Chapter 4 and throughout the contributions of this thesis, the utilization of SRR in MRI mandates a meticulous selection of acquisition strategy coupled with an accurate physical modeling of the imaging process. Regarding the latter, this thesis provides a comprehensive guidebook of the essential components and considerations required for the mathematical formulation and construction of a model-based SRR framework for quantitative MRI, including the utilization of image coordinate transformations for image warping, point spread function blurring, volumetric resampling, and the implementation of realistic noise models for magnitude MR data. Also, a clear distinction exists in the MR image reconstruction formulation, differentiating between anatomical SRR aimed at estimating a single high-resolution MR image and the more intricate model-based SRR for qMRI, which integrates a biophysical MR signal model into the forward model to estimate multiple high-resolution quantitative parameter maps.

In an attempt to gradually elaborate the expansion and complexity of the proposed method, Chapter 5 first described the proof-of-concept of joint motion and qMRI parameter estimation on small-image checkerboard-like phantoms, using T_1 mapping as the MR relaxometry model of choice. By means of an **extensive Monte Carlo simulation study with benchmarking against three alternative SRR approaches** – one without motion estimation, and two with motion correction as a preprocessing step – the potential of augmenting model-based SRR for quantitative T_1 mapping with **joint inter-image motion estimation** was explored. In particular, employing model-based super-resolution reconstruction using a maximum likelihood estimation framework, leveraging prior data knowledge to estimate motion and model parameters, demonstrated a substantial reduction in potential bias in the estimated T_1 map caused by motion. This improvement was observed when compared to a previously reported SRR-based T_1 mapping approach, in which motion registration was applied as a preprocessing step prior to T_1 mapping.

However, to further extend and validate the use of the model-based SRR with joint motion estimation on real data, a more advanced physical forward modeling of the MR acquisition of low-resolution contrast-weighted scans from the underlying high-resolution parameter maps was required, including, among others, an updated blurring model that takes into account the slice selection profile of a 2D multi-slice acquisition and a realistic Rician noise model for magnitude data. Moreover, the method needed to be generally applicable for other biophysical signal models besides T_1 relaxometry. Therefore, in Chapter 6, a rigorous unified framework for model-based super-resolution reconstruction with joint patient motion estimation using a Bayesian maximum a posteriori estimator was proposed. The framework allows the joint estimation of 3D isotropic high-resolution tissue parameter maps and inter-image motion parameters from a set of multi-slice contrast-weighted magnitude images with a low through-plane resolution. Additionally, to facilitate the framework's applicability to other imaging modalities beyond MR relaxometry, we have designed it to be modular with respect to both the quantitative signal model and the assumed distribution of the MR data. To underline the framework's strength and importance for qMRI, we have validated its use in both synthetic whole brain simulations and by using two in vivo human brain data sets, for T_1 and T_2 relaxometry parameter mapping, respectively. It has been demonstrated that the proposed SRR framework provides a more detailed delineation of brain structures and

shows **superior motion parameter estimation and improved quantitative tissue parameter mapping root-mean-square error** compared to state-of-the-art SRR approaches.

To emphasize the flexibility of the proposed model-based SRR framework with joint motion estimation, we extended its application to the more advanced scenario of Arterial Spin Labeling MRI for brain perfusion quantification. This involved establishing a research collaboration with colleagues at the Leiden University Medical Center in the Netherlands, leveraging their expertise and scanner hardware for acquiring pCASL data. As extensively detailed in Chapter 7, we introduced a new SRR framework for brain perfusion quantification using pseudo-continuous Arterial Spin Labeling (pCASL), which is capable of estimating 3D isotropic high-resolution quantitative cerebral blood flow (CBF) maps from a series of single-post-labeling-delay (single-PLD) pCASL control-label image pairs. These images were acquired with low through-plane resolution and rotated slice-encoding direction in a 2D multi-slice readout scheme. Building upon the SRR framework from Chapter 6, we jointly estimated motion between control and label images within a Bayesian estimation framework. This enabled accurate and precise CBF quantification without propagating pre-registration errors, while effectively leveraging prior knowledge of tissue properties and noise statistics. The rotation of the slice-encoding direction for each control-label image pair, along with the lower through-plane resolution, ensured a more uniform PLD distribution throughout the brain and enhanced background suppression efficacy. Consequently, this approach significantly improved SNR compared to conventional 2D multi-slice readout strategies, where CBF quantification is hindered by perfusion SNR slice dependence. Validation was performed both qualitatively and quantitatively through synthetic whole brain simulations and in vivo human brain data. Results demonstrated superior cerebral blood flow estimation in terms of root-mean-square error compared to a state-of-the-art approach using conventional 2D multi-slice readout strategies, even with additional hardware acceleration techniques like multiband applied in the latter. This SRR-pCASL framework addresses existing limitations of 2D multi-slice readout for pCASL and presents a promising alternative to the currently recommended segmented 3D readout schemes, which are sensitive to inter-shot motion and through-plane blurring due to T_2 decay along long echo trains. Future research may include direct comparison of 3D readout and 2D SRR readout schemes for pCASL imaging within the same time frame. Additionally, continued investigation into the optimal experimental design of the SRR acquisition strategy with multiband, which has shown significant gains in CBF estimation precision in simulations, is warranted.

Multiple ways forward from the presented research can be identified, including challenges inherent in joint optimization reconstruction frameworks involving different sets of parameters, such as **tuning multiple regularization hyperparameters and handling nonlinear parameter coupling**. Investigating whether these hyperparameters can be (implicitly) learned from training data, potentially through the use of recurrent inference machines (RIMs) (Putzky & Welling, 2017; Lønning et al., 2019), variational networks (Hammernik et al., 2018), or alternative learning-based methods, represents a promising research line. Ideally, such methods should be transferable across various modalities and MRI types, addressing also concerns of data scarcity of high-resolution isotropic quantitative parameter maps. Concurrently, **advancements in MRI acquisition sequences** are vital to support (super-resolution) reconstruction frameworks, with a focus on delivering data and contrast-weighted images within clinically acceptable scan times. For instance, the research discussed in Chapter 6 on T_2 parameter mapping optimized the MESE sequence for acquiring low-resolution T_2 -weighted data. Although such sequence suits the acquisition of multiple images at different TEs, it

remains a time-consuming readout scheme. Recent developments in sequence protocols, such as the GRAPPATINI sequence for T_2 -weighted imaging (Hilbert et al., 2018), may further facilitate the integration of SRR for quantitative T_2 parameter mapping into clinical practice. However, it's important to recognize the inherent limitations in MRI dictated by the relaxation timescales of biological tissues, which cannot be easily circumvented. Therefore, it remains imperative to underline the necessity for parallel advancements in both acquisition and reconstruction techniques.

Moreover, beyond the biophysical signal models explored in this thesis for MR relaxometry and perfusion MRI, ongoing research is delving into more advanced gMRI models with additional parameters to address existing limitations in physical modeling. For instance, the implementation of a nonlinear multi-regime multi-PLD pCASL model with model-based SRR has been briefly explored in checkerboard-like simulations (Bladt et al., 2017), and could be extended to real data validation experiments in line with the promising results for single-PLD pCASL in Chapter 7. This would allow for simultaneous high-resolution parameter mapping of both the cerebral blood flow and the arterial transit time. While such advanced models offer potentially more accurate estimations, the incorporation of multiple additional parameters carries the risk of a reduced precision. As such, parameter estimation should carefully monitor the impact on accuracy and precision, and the statistical performance of different signal models with alternative parameterizations needs to be thoroughly evaluated and compared. **Optimal experimental design** studies, where experiments are designed so as to minimize the variance of unbiased estimators towards the theoretically predicted Cramér-Rao lower bound, are essential to ensure that collected data is informative and precise to guarantee robust quantitative imaging. These studies should consider acquisition parameters such as the number of contrast-weighted images, the choice of TE, TI, and PLD values, parallel imaging, the impact of post-processing steps, and the number of slice encoding orientations, for SRR acquisitions in particular. Ideally, the optimal experiment should be repeatable and generally applicable across different MRI exams (brain, joints, etc.) and patient demographics, despite the inherent challenges posed by patient-specific variations.

The effectiveness of Al-based methods in radiology hinges on the quality of the training datasets and the machine learning models. These algorithms are trained using large datasets of images. The diversity and size of these datasets are crucial in developing robust AI models that can generalise well to new, unseen images. For some Al-based methods targeted at resolution enhancement in MRI it is crucial that high-resolution isotropic data is available. In this context, **SRR can play a pivotal role in unrolling Al-based methods**, by generating high-resolution datasets retrospectively from low-resolution scans, that can in turn be used as training data for Al or deep learning methods. In a recent study, the method developed in this thesis was used to construct high-resolution FLAIR images from low-resolution orthogonal FLAIR scans, demonstrating potential applications in lesion segmentation of MR scans of people with Multiple Sclerosis (MS) (Giraldo et al., 2023).

Finally, MRI accessibility remains low and extremely inhomogeneous around the world. According to the 2020 Organisation for Economic Co-operation and Development (OECD) statistics (OECD, 2021), there are approximately 65,000 installations of MRI scanners worldwide (\sim 7 per million inhabitants) compared to \sim 200,000 for CT and \sim 1,500,000 for ultrasound scanners. The distribution of MRI scanners is mainly concentrated within high income countries, with scarce availability in low and middle income countries and in rural

areas. Consequently, approximately 70% of the world's population have little to no access to MRI. The cost-prohibitive nature of high-field superconducting MRI scanners present itself as a major stumbling block. In particular the need for specialized hardware, dedicated RF-shielded hospital spaces, high maintenance costs for helium refill/re-liquification, and highly-trained personnel (Liu et al., 2021). The applicability of MRI is further limited by various contraindications for use (e.g., metal implants or claustrophobia) and limited accessibility. Additionally, patient age complicates access to MRI, as older patients tend to be more immobile, frailer, or more likely to have metallic implants. Ultimately, these factors present a major roadblock to MRI accessibility in healthcare. Recently, therefore, MRI research has been pivoting towards more accessible, low-cost MRI technologies, operating at ultra-low-field strengths (Arnold et al., 2023). Ultra-low-field MRI (ULF MRI) holds clear potential in creating a new class of low-cost MRI technologies for accessible healthcare, with scanners that are simple to onboard, maintain, and operate. ULF MRI holds several inherent attractions when compared to conventional high-field MRI, including the use of open-magnet configurations to reduce claustrophobia, low acoustic noise levels during scans, low sensitivity to metallic implants, and low RF specific absorption rate (Liu et al., 2021). In essence, ULF MRI can be operated routinely in a doctor's office, an elderly care facility, or at someone's own doorstep - thereby enhancing accessibility for elderly individuals with limited mobility, a need for caregiver assistance, claustrophobia, or metal implants. Consequently, **ULF MRI could revolutionize large-scale brain health screening**, crucial for addressing the growing number of brain diseases cases in our rapidly aging society and preventing healthcare systems from strain. Unfortunately, similar to qMRI, clinical adoption of ULF MRI is lagging behind due to its long scan time requirements. For ULF MRI, long scan times are currently necessary to compensate for the low SNR and low spatial resolution, inherent to current scanner designs and imaging at reduced magnetic field strengths. As explained in Chapter 3, scanners that operate in the low magnetic field regime typically suffer from reduced image quality that stems from the low Boltzmann polarization at those magnetic field strengths, resulting in weak nuclear magnetic resonance signals. Accordingly, images obtained at low-field suffer from low SNR, which can be mitigated in part by increased acquisition times (Marques et al., 2019). Alas, prolonged scan times are disadvantageous as they increase the risk of patient motion artifacts while also reduce throughput and patient comfort. Given this existing need for optimizing the trade-off between SNR, spatial resolution, and scan time, similar as for gMRI, the research in this dissertation could therefore also be viewed in the context of this ongoing effort to develop ULF technologies. In fact, the **synergistic** combination of model-based SRR with ULF MRI presents an opportunity to optimize the existing trade-off between SNR, spatial resolution, and scan time. Specifically, SRR could facilitate the direct estimation of high-resolution images or parameter maps from sets of low-resolution ULF MRI scans, enabling accurate and precise biomarker detection within limited scan durations. As such, the proposed research could further promote the development of a performant ULF MRI technology, enabling patient-centric and site-agnostic MRI scanners to fulfill unmet clinical needs across various global healthcare sites, while also opening the door to democratizing MRI for low and middle-income countries.

References

Arnold, T. C., Freeman, C. W., Litt, B., & Stein, J. M. (2023). Low-field MRI: Clinical promise and challenges. *J. Magn. Reson. Imaging*, *57*(1), 25–44.

Bladt, P., Beirinckx, Q., Van Steenkiste, G., Jeurissen, B., Achten, E., den Dekker, A. J., & Sijbers,

J. (2017). Super-resolution multi-PLD PCASL: a simulation study. *Magn. Reson. Mater. Phy., 30 (Suppl. 1)*(S396).

- Giraldo, D. L., Beirinckx, Q., den Dekker, A. J., Jeurissen, B., & Sijbers, J. (2023). Super-Resolution Reconstruction of Multi-Slice T2-W FLAIR MRI Improves Multiple Sclerosis Lesion Segmentation. *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2023, 1–4.
- Hammernik, K., Klatzer, T., Kobler, E., Recht, M. P., Sodickson, D. K., Pock, T., & Knoll, F. (2018). Learning a variational network for reconstruction of accelerated MRI data. *Magn. Reson. Med.*, 79(6), 3055–3071.
- Hilbert, T., Sumpf, T. J., Weiland, E., Frahm, J., Thiran, J.-P., Meuli, R., Kober, T., & Krueger, G. (2018). Accelerated T2 mapping combining parallel MRI and model-based reconstruction: GRAPPATINI. J. Magn. Reson. Imaging, 48(2), 359–368.
- Liu, Y., et al. (2021). A low-cost and shielding-free ultra-low-field brain MRI scanner. *Nature Communications*, *12*(1).
- Lønning, K., Putzky, P., Sonke, J.-J., Reneman, L., Caan, M. W., & Welling, M. (2019). Recurrent inference machines for reconstructing heterogeneous MRI data. *Med. Image Anal.*, *53*, 64–78.
- Marques, J. P., Simonis, F. F., & Webb, A. G. (2019). Low-field MRI: An MR physics perspective. J. Magn. Reson. Imaging, 49(6), 1528–1542.
- OECD (2021). *Magnetic resonance imaging (MRI) units (indicator)*. The Organization for Economic Cooperation and Development (OECD). doi: 10.1787/1a72e7d1-en.
- Putzky, P., & Welling, M. (2017). Recurrent Inference Machines for Solving Inverse Problems. arXiv, abs/1706.04008.

List of Abbreviations

1D	one-dimensional
2D	two-dimensional
3D	three-dimensional
ACL	anterior cruciate ligament
AD	Alzheimer's disease
ADC	apparent diffusion coefficient
AF	anisotropy factor
ΔΙ	artificial intelligence
arBias	absolute relative bias
	absolute relative blas
AJL	arterial transit time
	Arterial transit time
BASIL	Bayesian interence for Arterial Spin Labeling
DELSFU	Delgiali Science Policy
BEI	Brain Extraction Tool
BIDS	Brain Imaging Data structure
BOLD	Blood oxygenation level dependent
BS	background suppression
BW	bandwidth
CAIPIRINHA	Controlled Aliasing in Parallel Imaging Results in Higher Acceleration
CASL	continuous arterial spin labeling
cBCD	cyclic block-coordinate descent
CBF	cerebral blood flow
CCI	cubic convolution-based interpolation
CGLS	conjugate gradient least squares
CI	confidence interval
C-pCASL	conventional pseudo-continuous arterial spin labeling
CS	compressed sensing
CSD-BS	constrained slice-dependent background suppression
CSF	cerebrospinal fluid
СТ	computed tomography
CUDA	Compute Unified Device Architecture
CV	curriculum vitae
DICOM	Digital Imaging and Communications in Medicine
DFT	discrete Eourier transform
DSC	dynamic suscentibility contrast
	diffusion tensor imaging
FIRALI	European Imaging Biomarkers Alliance
EM	electromagnetic
	ocho planar imaging
	European Society of Padialogy
	ache train langth
	The angle
FDA	Food and Drug Administration (US)
	frequency encoding
FFI	fast Fourier transform
FID	free induction decay
tMRI	functional magnetic resonance imaging
FOV	tield of view
FSE	tast spin echo
FSL	FMRIB Software Library
FWHM	full with at half maximum

FWO	Fonds voor Wetenschappelijk Onderzoek
GB	gigabyte
GCV	generalized cross validation
GE	gradient echo
GM	gray matter
GRAPPATINI	GRAPPA + MARTINI
GRAPPA	generalized auto-calibrating partially parallel acquisition
GRASE	gradient-and-spin-echo
GRE	gradient recalled echo
GI	ground truth
HPC	High Performance Computing
HR	high-resolution
HS	
IBP	iterative back projection
IDF I	Inverse discrete Fourier transform
IES	inter echo spacing
	inverse fast Fourier transform
IR	inversion recovery
ISMRM	International Society for Magnetic Resonance in Medicine
	intravoxel incoherent motion
JMLE	joint maximum likelihood estimator
LL	Look-Locker
LLS	linear least squares
LM	Levenberg-Marquardt
LR	low-resolution
LS	least squares
MAP	Maximum a Posteriori
MARTINI	Model-based Accelerated Relaxometry by Iterative Nonlinear Inversion
MATLAB	matrix laboratory
MB	multiband
MC	Monte Carlo
MD	mean diffusivity
MESE	multi-echo spin echo
MI	mutual information
ML	maximum likelihood
MLE	maximum likelihood estimator
MPRAGE	magnetization prepared rapid acquisition gradient echoes
MR	magnetic resonance
MRF	magnetic resonance fingerprinting
MRI	magnetic resonance imaging
MS	multi-slice
MS	multiple sclerosis
MSE	mean squared error
MSK	musculoskeletal
NDAC	neurodegenerative diseases and cancer
NDAD	neurodegenerative disorders and age-related diseases
NIFTI	Neuroimaging Informatics Technology Initiative
NIST	National Institute of Standards and Technology
NLLS	nonlinear least squares
NMR	nuclear magnetic resonance
UECD	Organisation for Economic Co-operation and Development
PASL	pulsed arterial spin labeling
pCASL	pseudo-Continuous Arterial Spin Labeling
PD	proton density
PDF	probability density function

PE	phase encoding
PET	positron emission tomography
PI	parallel imaging
PLD	post-labeling delay
POCS	projection onto convex sets
PSF	point spread function
PSNR	peak signal-to-noise ratio
PVE	partial volume effects
QIBA	Quantitative Imaging Biomarkers Alliance
aMRI	quantitative MRI
RAM	random access memory
rCBF	relative cerebral blood flow
RF	radio frequency
RMSE	root-mean-square-error
RMMSE	root-(mean)-mean-square-error
ROI	region of interest
rRMSE	relative root-mean-square-error
rSTD	relative standard deviation
SA	signal averages
SAD	signal averages
SAN	spin ocho
	spin echo
	scandaru error
	Selisitivity encoding
SLIDER	simultaneous multi slice
SIVIS	simultaneous multi-silce
SPP	super-resolution
	single shot echo planar imaging
	structural similarity index massure
SOL	slice selection profile
551 STD	standard deviation
STOPM	Sumar resolution Tomographic Deconstruction for MDI
	T1 weighted
	T2 weighted
	T mApping with Partial Inversion Recovery
	acha tima
	turba factor
	total generalized variation
	inversion time
TP	ropotition time
TCE	turbo spin ocho
+SNP	tomporal signal to poise ratio
	total variation
INTE	ultra-high_field
	ultra-low-field
	ultrashort acha tima
	unbiased predictive rick estimator
	variable flip ande
	volume of interest
V/SVCI	velocity-selective arterial chin labeling
V) C	weighted linear least squares
	white matter
* * ! * !	Winte matter

Academic overview

Journal papers

- Beirinckx, Q., Bladt, P., van der Plas, M. C. E., van Osch, M. J. P., Jeurissen, B., den Dekker, A. J., and Sijbers, J., "Model-based super-resolution reconstruction for single-PLD pseudo-continuous Arterial Spin Labeling", *NeuroImage*, Vol. 286, pp. 120606, 2024. doi: 10.1016/j.neuroimage.2024.120506
- Van Brempt, N., Sgammato, R., Beirinckx, Q., Hammerschmid, D., Sobott, F., Dewilde, S., Moens, L., Herrebout, W. A., Johannessen, C., and Van Doorslaer, S., "The effect of pH and nitrite on the haem pocket of GLB-33, a globin-coupled neuronal transmembrane receptor of Caenorhabditis elegans", *Biochimica et Biophysica Acta (BBA) - Proteins & Proteomics*, Vol. 1871, no. 4, pp. 140913, 2023. doi: 10.1016/j.bbapap.2023.140913
- Nicastro, M., Jeurissen, B., Beirinckx, Q., Smekens, C., Poot, D. H. J., Sijbers, J., and den Dekker, A. J., "To shift or to rotate? Comparison of acquisition strategies for multi-slice super-resolution magnetic resonance imaging", *Frontiers in Neuroscience*, Vol. 16, pp. 1044510, 2022. doi: 10.3389/fnins.2022.1044510
- Beirinckx, Q., Jeurissen, B., Nicastro, M., Poot, D. H. J., Verhoye, M., den Dekker, A. J., and Sijbers, J., "Model-based super-resolution reconstruction with joint motion estimation for improved quantitative MRI parameter mapping", *Computerized Medical Imaging & Graphics*, Vol. 100, pp. 102071, 2022. doi: 10.1016/j.compmedimag.2022.102071
- Beirinckx, Q., Ramos-Llordén, G., Jeurissen, B., Poot, D. H. J., Parizel, P. M., Verhoye, M., Sijbers, J., and den Dekker, A. J., "Joint Maximum Likelihood estimation of motion and T1 parameters from magnetic resonance images in a super-resolution framework: a simulation study", *Fundamenta Informaticae*, Vol. 172, no. 2, pp. 105-128, 2020. doi: 10.3233/FI-2020-1896
- Van Doorslaer, S., Beirinckx, Q., Nys, K., Mangiameli, M. F., Cuypers, B., Callens, F., Vrielinck, H., and González, J. C., "EPR and DFT analysis of biologically relevant chromium(V) complexes with d-glucitol and d-glucose", *Journal of Inorganic Biochemistry*, Vol. 162, pp. 216-226, 2016. doi: 10.1016/j.jinorgbio.2016.07.012

Conference papers

 Giraldo, D. L., Q. Beirinckx, A. J. den Dekker, B. Jeurissen, and J. Sijbers, "Super-Resolution Reconstruction of Multi-Slice T2-W FLAIR MRI Improves Multiple Sclerosis Lesion Segmentation", 2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), July 24-27, 2023, Sydney, Australia, pp. 1-4, 2023. doi: 10.1109/EMBC40787.2023.10341047

 Nicastro, M., B. Jeurissen, Q. Beirinckx, C. Smekens, D. H. J. Poot, J. Sijbers, and A. J. den Dekker, "Comparison of MR acquisition strategies for super-resolution reconstruction using the Bayesian Mean Squared Error", *Proceedings of the 16th Virtual International Meeting on Fully 3D Image Reconstruction in Radiology and Nuclear Medicine (Fully 3D), 19-23 July, 2021, Leuven, Belgium*, arXiv:2110.04143, pp. 435-439, 2021.

Conference abstracts

- Khan, H., H. C. Woodruff, D. L. Giraldo, Q. Beirinckx, J. Sijbers, L. M. Peeters, and P. Lambin, "Predicting Disability Progression in People with Multiple Sclerosis: Combining Volumetric and Radiomics Features from T2-Weighted FLAIR MR Images", Conference: European Conference of Radiology (ECR), February 28 - March 3, 2024, Vienna, Austria. Proceedings: ECR 2024 Book of Abstracts, *Insights Imaging, Vol. 15* (*Suppl.*).
- Khan, H., D. L. Giraldo, Q. Beirinckx, J. Sijbers, P. Lambin, H. C. Woodruff, and L. M. Peeters, "Exploring the Correlation between Disability Status and Brain Volumetric Measures Using Real-World Retrospective Magnetic Resonance Images in People with Multiple Sclerosis", Conference: 9th Joint European Committee for Treatment and Research in Multiple Sclerosis (ECTRIMS) Americas Committee for Treatment and Research in Multiple Sclerosis (ACTRIMS) Meeting, October 11-13, 2023, Milan, Italy. Proceedings: Multiple Sclerosis Journal, Vol. 29, No. 3 suppl., pp. 841-842, 2023. doi: 10.1177/13524585231196195
- Smekens, C., F. Vanhevel, Q. Beirinckx, T. Janssens, and P. Van Dyck, "Accelerated Dual-contrast Three-dimensional Knee Magnetic Resonance Imaging Using Superresolution Reconstructed Deep Learning-enhanced Two-dimensional Dixon Turbo Spinecho Imaging", Conference: 30th Annual Scientific Meeting of the European Society of Musculoskeletal Radiology (ESSR), June 22-24, 2023, Bilbao, Spain. Proceedings: Seminars in Musculoskeletal Radiology, Vol. 27 (Suppl. 01), pp. S1-S24, 2023. doi: 10.1055/s-0043-1770025
- Smekens*, C., Q. Beirinckx*, F. Vanhevel, P. Van Dyck, A. J. den Dekker, J. Sijbers, T. Janssens, and B. Jeurissen, "High-resolution T₂* mapping of the knee based on UTE Spiral VIBE MRI", Conference: 38th Annual Scientific Meeting of the European Society for Magnetic Resonance in Medicine & Biology (ESMRMB), October 7-9. Proceedings: Magn Reson Mater Phy, Vol.34 (Suppl. 1), C1.O3, pp. 553-554, 2021. doi: 10.1007/s10334-021-00947-8. (*Both authors contributed equally.)
- Nicastro, M., B. Jeurissen, Q. Beirinckx, C. Smekens, D. H. J. Poot, J. Sijbers, and A. J. den Dekker, "Rotated or shifted sets of multi-slice MR images for super-resolution reconstruction? A Bayesian answer", Conference: *38th Annual Scientific Meeting of the European Society for Magnetic Resonance in Medicine & Biology (ESMRMB)*, October 7-9. Proceedings: Magn Reson Mater Phy, Vol.34 (Suppl. 1), C1.O7, pp. S56-S57, 2021. doi: 10.1007/s10334-021-00947-8.

- Smekens*, C., Q. Beirinckx*, F. Vanhevel, P. Van Dyck, A. J. den Dekker, J. Sijbers, T. Janssens, and B. Jeurissen, "Super-resolution T2* mapping of the knee using UTE Spiral VIBE MRI", in *Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM)*, Vol. 29, pp. 3920, 2021. (*Both authors contributed equally.)
- Bladt, P., Q. Beirinckx, M. C. E. van der Plas, S. Schmid, W. M. Teeuwisse, B. Jeurissen, A. J. den Dekker, J. Sijbers, and M. J. P. van Osch, "Super-resolution reconstruction of single-PLD pseudo continuous ASL images", Conference: 28th Annual Meeting of ISMRM & SMRT, Virtual conference and exhibition, 2020.
- 8. Nicastro, M., **Q. Beirinckx**, P. Bladt, B. Jeurissen, S. Klein, J. Sijbers, D. H. J. Poot, and A. J. den Dekker, "Optimal design of a T1 super-resolution reconstruction experiment: a simulation study", Conference: 12th Annual Meeting of the ISMRM Benelux Chapter, Arnhem, The Netherlands, 2020.
- Beirinckx, Q., P. Bladt, M. C. E. van der Plas, S. Schmid, W. M. Teeuwisse, B. Jeurissen, A. J. den Dekker, J. Sijbers, and M. J. P. van Osch, "Super-resolution strategies for single-PLD pseudo continuous ASL", Conference: *12th Annual Meeting of the ISMRM Benelux Chapter*, Arnhem, The Netherlands, 2020.
- Beirinckx, Q., B. Jeurissen, M. Verhoye, A. J. den Dekker, and J. Sijbers, "Superresolution T1 mapping with integrated motion compensation in a joint maximum likelihood framework", Conference: *36th Annual Scientific Meeting of the European Society for Magnetic Resonance in Medicine & Biology (ESMRMB)*, Rotterdam, The Netherlands, October 3-5. Proceedings: Magn Reson Mater Phy. 2019 Oct, Vol.32 (Suppl. 1): S14.05. doi: 10.1007/s10334-019-00754-2.
- Ribeiro Sabidussi, E., M. Nicastro, S. Bazrafkan, Q. Beirinckx, B. Jeurissen, A. J. den Dekker, S. Klein, and D. H. J. Poot, "A deep learning approach to T1 mapping in quantitative MRI", Conference: *36th Annual Scientific Meeting of the European Society for Magnetic Resonance in Medicine & Biology (ESMRMB)*, Rotterdam, The Netherlands, October 3-5. Proceedings: Magn Reson Mater Phy. 2019 Oct, Vol.32 (Suppl. 1): S09.05. doi: 10.1007/s10334-019-00754-2.
- 12. Ramos-Llordén, G., **Q. Beirinckx**, A. J. den Dekker, and J. Sijbers, "An educational presentation on accurate and precise MRI relaxometry: the often disregarded but critical role of statistical parameter estimation", *10th Annual Meeting of the ISMRM Benelux Chapter*, Antwerp, Belgium, 2018.
- 13. Ramos-Llordén, G., **Q. Beirinckx**, A. J. den Dekker, and J. Sijbers, "Accurate and precise MRI relaxometry: the often disregarded but critical role of statistical parameter estimation", *26th Annual Meeting of the ISMRM*, Paris, France, 2018.
- Bladt*, P., Q. Beirinckx*, G. Van Steenkiste, B. Jeurissen, E. Achten, A. J. den Dekker, and J. Sijbers, "Super-resolution multi-PLD PCASL: a simulation study", Conference: 34th Annual Scientific Meeting of the European Society for Magnetic Resonance in Medicine & Biology (ESMRMB), Barcelona, Spain, October 19-21. Proceedings: Magn Reson Mater Phy 2017 Oct, Vol. 30 (Suppl. 1): S396. doi: 10.1007/s10334-017-0634-z. (*Both authors contributed equally.)
Awards

 Magna Cum Laude Merit Award (ISMRM 2021) for the work: Smekens*, C., Q. Beirinckx*, F. Vanhevel, P. Van Dyck, A. J. den Dekker, J. Sijbers, T. Janssens, and B. Jeurissen, "Super-resolution T2* mapping of the knee using UTE Spiral VIBE MRI", in *Proceedings of the International Society for Magnetic Resonance in Medicine* (ISMRM), Vol. 29, pp. 3920, 2021. (*Both authors contributed equally.)

Invited talks & co-organiser

1. B-Q MINDED Summer School: an intensive, multidisciplinary course in Quantitative MRI (Q-MRI), Universiteit Antwerpen, Antwerp, Belgium, August 27 - September 4, 2018. Guest speaker on the topic of 'model-based super-resolution reconstruction', and co-organiser of Workshop Unishell/Git/Python in collaboration with the Netherlands eScience Center.

Relevant courses and workshops

- 1. ESMRMB Lectures on Magnetic Resonance: MRI simulation for sequence development, protocol optimization, and education, Eindhoven University of Technology, De Zwarte Doos, Eindhoven, The Netherlands, June 28-30, 2017.
- 2. Annual meeting of the FWO Scientific Research Network (WOG) "Turning images into value through statistical parameter estimation": Model selection, Irish College, Leuven, Belgium, June 11, 2018.
- B-Q MINDED Summer School: an intensive, multidisciplinary course in Quantitative MRI (Q-MRI), Universiteit Antwerpen, Antwerp, Belgium, August 27 - September 4, 2018.
- 4. B-Q MINDED Summer School: an intensive, multidisciplinary course in Quantitative MRI (Q-MRI), Universiteit Antwerpen, Antwerp, Belgium, August 26-30, 2019.
- 5. Annual meeting of the FWO Scientific Research Network (WOG) "Turning images into value through statistical parameter estimation": Image reconstruction, Zaal Nick Ervinck, NV Zebrastraat, Ghent, Belgium, September 20, 2019.
- 6. Annual meeting of the FWO Scientific Research Network (WOG) "Turning images into value through statistical parameter estimation": Deep Learning, September 15, 2020.

Teaching and supervision

- 2017-2021: Tutor exercises 'Fysica m.i.v. wiskunde' (1st year B. Sc. Pharmacy and 1st year B. Sc. Biomedical Sciences, ~ 150 students), supervising lecturer: Prof. Dr. Jan Sijbers.
- 2. **2017-2021**: Tutor exercises 'Fysica voor biomedisch onderzoek' (1st year B. Sc. Biomedical Sciences), supervising lecturer: Prof. Dr. Jan Sijbers.

3. **2021-2023**: Co-supervisor of Mateo Sierens' research projects in B. Sc. and M. Sc. in Computer Sciences, University of Antwerp, on "The development of a Python-based super-resolution reconstruction framework using Recurrent Inference Machines as iterative solvers". (Supervisor: Prof. Dr. Jan Sijbers)

Technical training and research support

1. Trained operator of the FlexCT micro-CT scanner system of the imec-Vision Lab, University of Antwerp. Providing scanning and reconstruction support for internal and external clients, e.g. 3D CT scans of human or animal explant lungs in collaboration with the Antwerp Surgical Training, Anatomy and Research Centre (ASTARC), or scans of stained java finches for beak movement analysis in collaboration with the Functional Morphology (FunMorph) research group at the University of Antwerp.