Three-Dimensional Hybrid Radial Cartesian Echo Planar Imaging for Functional MRI



Nadine N. Grädel
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2016

Preface

List of Publications Arising from this Thesis

Journal Articles

1. **N.N. Graedel**, J.A. McNab, M. Chiew, K.L. Miller. "Motion correction for functional MRI with three-dimensional hybrid radial-Cartesian EPI.", *Magnetic Resonance in Medicine*. In press.

Conference Proceedings

1. **N. N. Graedel**, M. Chiew, and K. L. Miller. "Motion correction for functional MRI with hybrid radial-Cartesian 3D EPI". *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 940

2. **N. N. Graedel**, M. Chiew, J. A. McNab, and K. L. Miller. "FMRI using a 3D radial-Cartesian trajectory: spatio-temporal tunability and artifact correction". *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, Canada, 2015, p. 360.

3. **N. N. Graedel**, M. Chiew, S. Clare, and K. L. Miller. "Complex interactions of physiological noise and acceleration on tSNR in 3D EPI". *Proceedings of the 22nd Annual Meeting of the ISMRM*. Milan, Italy, 2014, p. 2995

Acknowledgements

Many people deserve acknowledgement for their contributions to the work presented in this thesis. I would like to thank Benedikt Poser for providing guidance on the implementation of the 3D EPI sequence. Thank you also to Jennifer McNab for providing me with the original TURBINE sequence code. Many thanks are to be extended to Benedikt Poser and Lauren Bains for conducting the 3D EPI comparison experiments, as well as David Norris, Peter Koopmans and Markus Barth for interesting discussions on 3D EPI. Further, I would like to thank Ruediger Stirnberg and Sam Hurley for helpful testing of 3D EPI sequences. I am also grateful to Kerrin Pine for providing motion traces and to Matthew Webster for the helpful discussions on MCFLIRT. Furthermore, I would like to thank Lucilio de Cordero-Grande and Jo Hajnal for discussing future plans for TURBINE image reconstruction with motion correction.

On a personal note, my greatest thanks go to my two supervisors Karla Miller and Mark Chiew. I am very grateful to Karla for taking me aboard the Miller group and thus allowing me to learn in such a stimulating environment. Thank you for always taking the time to talk about my work and for leading by example by being such an inspiring scientist. I would like to thank Mark for his invaluable guidance on this project, including countless hours of patient discussions and feedback. Your competence, brilliance and curiosity – even for the world's most obscure problems – is contagious and inspiring.
I would like to thank Jenn, for allowing me to continue the TURBINE legacy and for sending me Karla's way originally. Further, I want to thank Johanna for her support and for fun times measuring and discussing GIRFs, which I hope will be continued. Thank you to Sam for making working late nights enjoyable and for never leaving me hanging on a pub call. I would like to thank my fellow 2012 DPhil starters Ellie and Hannah for their friendship, support and for all the adventures – from gingerbread MRI scanners to conference road trips – that we got to experience together. Thank you to Olivia for all the great times (in Oxford and abroad) and to all of the FMRIB physics students, thank you for being such a kind and supportive bunch. FMRIB is a very inspiring and friendly community, which made it easy to enjoy going to work. Thank you for many interesting discussions, random fascinating lunch conversations, retreats as well as conference and pub trips.

Many thanks go to Annina, Lena, Carole, Ellie and Francesco for help with proofreading this thesis. A special big thank you and a lot of gratitude go to Carole, for her invaluable help with typesetting and support during long working days and nights.

There are so many people who have made this 4-year Oxford adventure very special to me. I want to thank my many Oxford housemates Sweta, Sam, Annina, Ioanna, Shelley, Jon and Emily to help make dubious Oxford student accommodations easy to bear. Special mentions to Sam for letting me "borrow" all of her food (sorry about the delicious yoghurt). To the first person I ran into at Oxford, Emma, thank you for being such a positive and supportive spirit. Special mention also goes to my "JR lunch club" pals, Hannah and Dora, for always being there to cheer me up. A big thank you to Emily for co-founding the all-chocolate-diet and being such a supportive friend. Your "are you alive?" checks and chocolate deliveries this summer have been essential.

I am very happy to have been a member of Lincoln College, specifically I greatly appreciated the time spent with the Lincoln MCR and LCBC, as well as my fellow Berrow scholars. I thank the Oxford University Women's Boat Club for countless incredible adventures and memories that will last a lifetime. I am grateful to Christine and Natasha for fostering my unexpected rowing career. I salute all of the 2013-2015 squad members, and especially my two boat race crews, for being the best team I could imagine.

I would like to dedicate this thesis to my family. Mama and Papa, you taught me to value hard work and never take anything for granted. The knowledge that you will catch me if I fall, has allowed me to take bigger leaps than I ever imagined. To my three wonderful sisters Annina, Lena and Carole, you are my best friends and a permanent source of inspiration. I hope we will keep challenging and supporting each other all our lives. Francesco, there is no way I can thank you enough. You have been my rock, my source of sanity and the light of my life. Everything I do builds on your unwavering love. Also, thank you for reminding me to "chill out" and enjoy a good glass of whisky every now and then.

Abstract

Functional magnetic resonance imaging (fMRI) has provided neuroscientists with a powerful tool to non-invasively study brain function. Typically, fMRI data acquisition is performed using the well-established multi-slice two-dimensional echo planar imaging (2D EPI) technique. While 2D EPI has the considerable advantage of robustness, it is relatively SNR inefficient, particularly at high spatial resolution. Three dimensional (3D) sampling approaches, such as multi-shot 3D EPI provide a theoretical SNR gain compared to 2D EPI and can utilize parallel imaging acceleration along multiple dimensions, leading to the potential for higher spatial and temporal resolution. However, these multi-shot acquisitions span several seconds, making them susceptible to physiological fluctuations. In particular, subject motion is a major source of image degradation.

This thesis aims to characterise and improve fMRI acquisition techniques based on 3D EPI approaches. We explored the temporal SNR characteristics of standard segmented 3D EPI for different spatial resolutions and acceleration factors. Specifically, we studied how physiological noise affects the optimal choice of imaging parameters, such as the amount of acceleration.

To address some of the shortcomings of conventional 3D EPI, we implemented a hybrid radial-Cartesian 3D EPI trajectory, called TURBINE. This scheme collects EPI "blades" which are rotated about the phase-encoding axis using a golden angle rotation increment, allowing reconstruction at flexible temporal resolution. The self-navigating properties of the sequence are used to determine motion estimates from high temporal resolution navigator images and correct for subject motion as part of the image reconstruction process. We demonstrated that this scheme reduces the impact of motion on fMRI data in the presence of subtle and large subject motions.

The techniques developed in this thesis aim to increase the flexibility and robustness of fMRI acquisitions. Ultimately, this research may help increase the utility of fMRI in difficult subjects or patient populations.

Contents

List of Figures **x**

List of Tables **xii**

List of Abbreviations **xiii**

1 Introduction **1**

1.1 Thesis Outline 3

2 MRI Background **5**

2.1 The MR Signal 6

2.1.1 NMR Basics 6

2.1.2 Excitation, Relaxation and Signal Detection 7

2.2 Image Formation 8

2.2.1 Slice Selection 10

2.2.2 Spatial Encoding and k-Space 10

2.2.3 Practical Sampling Considerations 12

2.2.4 Frequency and Phase Encoding 14

2.3 Pulse Sequences and k-Space Trajectories 15

2.3.1 Gradient Echo Pulse Sequence 16

2.3.2 Echo Planar Imaging 17

2.3.3 Non-Cartesian Imaging 17

2.3.4 Spoiling 18

2.3.5 Choice of Flip Angle 19

2.4 Image Reconstruction 19

2.4.1 Gridding 20

2.4.2 NUFFT 22

2.4.3 Parallel Imaging 22

2.4.4 Non-Cartesian Parallel Imaging 24

2.5 Typical MRI Artefacts 27

2.5.1 Chemical Shift Artefact 27

2.5.2 B_0 Inhomogeneity 28

2.5.3 Trajectory Imperfections 28

3 Functional MRI – Background and Current Methods 30
3.1 The Basics of Functional MRI . . . 31
3.1.1 Physiological Basis for fMRI . . . 32
3.1.2 The BOLD Signal . . . 33
3.2 BOLD fMRI at Ultra-High Fields . . . 36
3.3 BOLD fMRI Acquisition Strategies . . . 38
3.4 fMRI Experimental Design and Analysis . . . 40
3.5 Physiological Noise and Motion . . . 42
3.5.1 Physiological Noise . . . 42
3.5.2 Motion in fMRI . . . 45
3.6 Summary and Thesis Motivation . . . 49

4 Temporal SNR Behaviour of 3D EPI at 7 T 50
4.1 Introduction . . . 51
4.2 Theory . . . 52
4.2.1 Influence of Slice Acceleration on tSNR in 3D EPI . . . 56
4.3 Methods . . . 58
4.3.1 Sequence Development . . . 58
4.3.2 Sequence Testing and Comparison to 2D EPI . . . 59
4.3.3 Main Experiment: Impact of Acceleration on tSNR . . . 61
4.3.4 Image Reconstruction and Analysis . . . 61
4.4 Results . . . 62
4.4.1 Sequence Testing and Comparison to 2D EPI . . . 62
4.4.2 Main Experiment: Impact of Acceleration on tSNR . . . 65
4.5 Discussion . . . 67
4.6 Conclusion . . . 72
4.A Appendix: Literature Comparison and tSNR Simulation . . . 72
4.B Appendix: Investigation of the Temporal Stability of 3D EPI (Phantom Experiments) . . . 76
4.B.1 2D vs. 3D EPI Phantom Comparison . . . 76
4.B.2 λ Measurements . . . 78
4.B.3 Cross-site Comparison . . . 79
4.B.4 Discussion . . . 82

5 TURBINE Sequence for Functional MRI 83
5.1 Introduction . . . 84
5.2 TURBINE Pulse Sequence . . . 85
5.3 Image Reconstruction . . . 90
5.3.1 Gridding and NUFFT . . . 90
5.3.2 Nyquist Ghost Correction . . . 92

5.3.3 In-plane Acceleration Using GRAPPA . . . 95
5.3.4 Radial/3D Acceleration Using SPIRiT . . . 97
5.3.5 SPIRiT Algorithm Parameters . . . 97
5.3.6 Coil Compression and Reconstruction Time Management . . . 103
5.4 tSNR and fMRI Evaluation at 3 T . . . 104
5.4.1 Reconstruction and Analysis . . . 104
5.4.2 Flexible Temporal Resolution for fMRI . . . 105
5.4.3 tSNR and Preliminary Comparison to 2D/3D EPI . . . 106
5.4.4 Preliminary fMRI Comparison to 2D/3D EPI . . . 108
5.5 Preliminary 7 T fMRI . . . 109
5.6 Discussion . . . 111
5.6.1 Image Reconstruction . . . 113
5.6.2 TURBINE Trajectory Properties and Artefacts . . . 114
5.6.3 TURBINE tSNR and fMRI . . . 116
5.7 Conclusion . . . 117

6 Motion Correction for Functional MRI with TURBINE 118
6.1 Introduction . . . 118
6.2 Motion Correction Concept . . . 120
6.2.1 Correction for z-Translations Only . . . 121
6.2.2 Image-Space Based Motion Estimation . . . 124
6.2.2.1 Motion Estimation Using Simulated Motion Data . . . 126
6.2.3 Estimation of Motion Parameters in k-Space . . . 128
6.3 In Vivo Validation of TURBINE Motion Correction . . . 130
6.3.1 Methods . . . 130
6.3.1.1 Data Acquisition . . . 130
6.3.1.2 Image Reconstruction and Motion Correction . . . 131
6.3.1.3 Evaluation and FMRI Analysis . . . 133
6.3.2 Results . . . 134
6.3.2.1 Cued Motion Experiment . . . 134
6.3.2.2 fMRI with Natural Motion . . . 138
6.3.2.3 fMRI with no Explicit Motion . . . 141
6.3.2.4 Temporal SNR . . . 141
6.4 Discussion . . . 141
6.5 Conclusion . . . 146

7 Summary and Future Work **147**
7.1 Summary . . . 147
7.2 Future Work . . . 148
7.2.1 TURBINE Sequence Improvements and Variants . . . 149
7.2.2 Image Reconstruction Improvements Using Gradient Characterisation and B_0 Correction . . . 150
7.2.2.1 Preliminary Work on Trajectory Correction . . . 151
7.2.2.2 Other Reconstruction Improvements . . . 152
7.2.3 Improved TURBINE Motion Correction . . . 153
7.2.3.1 Refinement of the Navigator Approach . . . 153
7.2.3.2 Motion Correction Within the Image Reconstruction . . . 154
7.2.3.3 Motion Camera . . . 154
7.2.3.4 Static B_0 and Shimming in the Presence of Motion . . . 154
7.3 Final Remarks . . . 155

References **156**

List of Figures

2.1 Slice selection 9
2.2 2D and 3D MRI 12
2.3 Cartesian 2D k-space and image space 14
2.4 Spoiled 2D GRE sequence 16
2.5 Echo Planar Imaging 17
2.6 Spiral and radial trajectory 18
2.7 Gridding process 21
2.8 Coil sensitivities 22
2.9 Flowchart of the CG SPIRiT algorithm 26

3.1 Physiological basis of BOLD 32
3.2 Magnetic dipole field around a vessel 34
3.3 Haemodynamic response function 35
3.4 Illustration of simultaneous multi-slice acquisition 38
3.5 Illustration of the effect of intra-volume motion 48

4.1 Theoretical SNR for 2D and 3D EPI 54
4.2 Krüger model 55
4.3 Hypothesis for tSNR behaviour 57
4.4 Segmented 3D EPI sequence 58
4.5 High-resolution fMRI 63
4.6 Task fMRI comparison between 2D and 3D EPI 63
4.7 tSNR comparison between 2D or 3D EPI sequence 64
4.8 tSNR as a function of resolution R_z (example subject) 66
4.9 tSNR as a function of resolution R_z (summary) 67
4.10 Simulation of tSNR for 2D and 3D EPI 69
4.11 Simulated 3D EPI tSNR 70
4.12 Phantom tSNR 77
4.13 Phantom λ measurement 79
4.14 Cross-site tSNR comparison 81

5.1 Diagram of the 3D hybrid radial-Cartesian sampling scheme 84
5.2 Blade ordering schemes 86

5.3 TURBINE pulse sequence diagram 87
5.4 Reconstruction pipelines 90
5.5 Nyquist ghost correction 93
5.6 Example TURBINE images 96
5.7 Flexible temporal resolution of TURBINE for phantom 98
5.8 Flexible temporal resolution of TURBINE for subject 99
5.9 Iteration count in SPIRiT reconstruction 102
5.10 fMRI with flexible temporal resolutions 105
5.11 tSNR of TURBINE 107
5.12 tSNR comparison to 2D and 3D EPI 108
5.13 fMRI comparison to 2D and 3D EPI 110
5.14 7T TURBINE images 111
5.15 7 T fMRI results 112

6.1 Types of rigid body motion 122
6.2 Pilot study with z-translations only 123
6.3 Result of z-translation correction 124
6.4 Overview of motion-correction approach 125
6.5 Motion estimates for simulated step motion 127
6.6 Motion estimates for simulated natural motion 129
6.7 Overview of the motion correction and reconstruction pipeline 132
6.8 Global phase fluctuations 135
6.9 Cued motion experiment, example subject 136
6.10 Cued motion experiment, all subjects 137
6.11 fMRI with natural motion 139
6.12 fMRI with natural motion, activation maps 140
6.13 fMRI with natural motion, time courses 140
6.14 fMRI with no deliberate motion 142

7.1 Illustration of example TURBINE sequence variants 149
7.2 TURBINE trajectory measurement 152

List of Tables

4.1 Imaging parameters for in vivo experiments 60
4.2 tSNR literature values for 3D and 2D EPI 74

5.1 Imaging parameters for all experiments in chapter 5 91
5.2 SPIRiT reconstruction parameters 101
5.3 Spatial resolution estimates for reconstructions in Fig. 5.10 106

List of Abbreviations

γ	Gyromagnetic ratio of ^{1}H
$\boldsymbol{\omega_0}$	Larmor (resonant) frequency of ^{1}H
$\mathbf{B}_0$	Main magnetic field
1D	One dimensional
2D	Two dimensional
3D	Three dimensional
BOLD	Blood-Oxygenation Level Dependent
BW	Bandwidth
CG	Conjugate Gradient
CNR	Contrast to Noise Ratio
DFT	Discrete Fourier Transform
EPI	Echo Planar Imaging
FFT	Fast Fourier Transform
fMRI	Functional Magnetic Resonance Imaging
FOV	Field-Of-View
FT	Fourier Transform
GLM	General Linear Model
GRAPPA	GeneRalized Autocalibrating Partial Parallel Acquisition
M	Magnetisation
MC	Motion Correction
MRI	Magnetic Resonance Imaging
NMR	Nuclear Magnetic Resonance
PCA	Principal Component Analysis
PE	Phase Encode
R	Parallel imaging acceleration factor
RF	Radiofrequency

RO . . . . . Readout

ROI . . . . . Region Of Interest

SENSE . . . . . SENSitivity Encoding

SNR . . . . . Signal to Noise Ratio

SPIRiT . . . . . iTerative Self-consistent Parallel Imaging Reconstruction

T . . . . . Tesla (magnetic field strength)

$\mathbf{T}_1$ . . . . . Longitudinal (spin-lattice) relaxation time

$\mathbf{T}_2$ . . . . . Transverse (spin-spin) relaxation time

$\mathbf{T}_2^*$ . . . . . Transverse relaxation time including static inhomogeneity

TBWP . . . . . Time Bandwidth Product

TE . . . . . Echo Time

TR . . . . . Repetition Time

$\mathbf{TR}_{vol}$ . . . . . Volume repetition time

tSNR . . . . . Temporal Signal to Noise Ratio

TURBINE . . . Trajectory Using Radially Batched Internal Navigator Echoes

WE . . . . . Water Excitation (using a binomial pulse)

1

Introduction

The brain is the most fascinating, and still one of the most mysterious organs in the human body. In recent decades, major advances in understanding brain function and organisation have come from the development of neuroimaging tools. Functional MRI (fMRI) as a non-invasive, non-ionizing image modality has facilitated the study of brain function (especially of healthy volunteers), which has provided unique new insights. Introduced in the early 1990s, fMRI has rapidly become a standard research tool in cognitive science, and is under an active development as a clinical research and diagnostic instrument.

FMRI acquisition is challenging because it is a dynamic technique, aiming to detect very small signal changes that reflect brain activity; it therefore suffers from inherently low contrast-to-noise ratio. Measuring these activation characteristics requires relatively fast sampling, which causes the technique to be susceptible to image distortions and other imaging artefacts.

FMRI has traditionally used moderate spatial and temporal resolutions, in part because of technical limitations. There are however considerable incentives to acquire smaller voxels and sample faster. One of the advantages of higher spatial resolution is that it allows studying new aspects of the cortical and subcortical functional architecture while

high temporal resolution increases the temporal degrees of freedom of the experiment. The fMRI field is shifting towards the detection of networks, which emphasizes the need for rapid sampling to boost statistical power.

Over the last decade fMRI has made large advances driven by a range of technological development. These consist of improvements to many aspects of the scanner hardware, as well as new acquisition and reconstruction strategies, with advancements in one domain pushing developments in the other.

One important hardware advance has been the push towards higher static magnetic fields, which boosts sensitivity and specificity for fMRI. Improvements in gradient coil technology now allow faster and more accurate sampling. Finally, highly parallelised signal reception with receive arrays using up to 64 channels has allowed parallel imaging, and subsequently the speed of the acquisitions, to achieve new levels.

In parallel to the hardware development there has been a change in how sampling is approached. New and improved parallel imaging techniques and reconstruction approaches have been introduced, allowing robust image reconstruction at higher under-sampling factors. Additionally, methods based on sparsity and low-rank promoting constraints call for new and flexible encoding strategies.

A large number of new acquisition strategies are now currently being developed. An emerging theme is that the field is moving away from the traditional 2D slice-by-slice imaging to a new era of volumetric imaging. The two dominant methods in this space currently are 3D imaging and simultaneous multi slice (SMS) imaging. Both 3D and SMS imaging have improved acceleration capability and improved SNR efficiency when compared with 2D slice-by-slice imaging. Similarly, both methods allow much more flexible encoding and under-sampling yielding improvements in parallel imaging reconstruction.

Methods using 3D encoding, like 3D EPI, are fundamentally the most flexible. Any 2D imaging method (including SMS) is more likely to become limited by the fact that encoding is achieved differently in one dimension compared to the other two; for example

by reaching gradient and RF safety limits. 3D encoding allows full flexibility on the sampling pattern and aliasing can be distributed into all spatial directions in a controlled fashion. For example, it is much easier to meet the demands for incoherent sampling that many new reconstruction techniques have.

However 3D methods, such as 3D EPI, come with a number of challenges. Subject motion has always been a major source of image degradation for functional MRI, and multi-shot sequences, like 3D EPI, are especially sensitive to it. The same holds for any other temporal fluctuation, such as physiological noise or scanner instabilities.

The goal of this thesis is to address the Achilles heel of 3D EPI: its sensitivity to physiological noise and motion. We try to improve our understanding of the impact of these fluctuations on 3D EPI image reconstruction and to develop methods to mitigate their detrimental effect on fMRI data. Given the potential that 3D EPI has for enabling high-resolution fMRI, improving robustness of the method by developing a correction for nuisance fluctuations could enable 3D EPI to be adopted more broadly.

1.1 Thesis Outline

The work presented in this thesis is concerned with the development and evaluation of fMRI acquisition methodology for 3D EPI. This body of the thesis consists of two background chapters and three research chapters as described below:

Chapter 2: MRI Background provides a brief overview of the relevant MRI background to this thesis. Topics covered include basic encoding principles, image reconstruction and parallel imaging.

Chapter 3: Functional MRI – Background and Current Methods gives an introduction to functional MRI. It explains the basic principles of fMRI contrast origins, and presents current acquisition strategies and challenges.

The three research chapters that follow are organised chronologically in the order the work was done.

Chapter 4: Temporal SNR Behaviour of 3D EPI at 7 T documents basic testing of the standard 3D EPI sequence and explores its behaviour at different spatial resolutions and acceleration factors. Specifically, we investigate potential interactions between acceleration and physiological noise and how the temporal SNR reflects a balance of these processes.

In **Chapter 5: TURBINE Sequence for Functional MRI** we introduce a novel 3D hybrid radial-Cartesian acquisition strategy for functional MRI. The pulse sequence is called TURBINE and uses EPI planes rotated about the slab-encoding axis. In this chapter we document the implementation and testing of the sequence and demonstrate its suitability for fMRI data acquisition.

Chapter 6: Motion Correction for TURBINE fMRI presents a motion correction scheme using the TURBINE sequence. The approach makes use of the flexible temporal resolution property of the golden angle scheme. We demonstrate that the motion correction works robustly and that we can achieve substantial artefact reduction as well as improvements in the detection of fMRI activation in the presence of subject motion.

Finally, **Chapter 7: Summary and Future Work** provides a summary of the thesis and discusses possible extensions and future directions for this work.

2

MRI Background

Contents

2.1 The MR Signal 6
2.1.1 NMR Basics 6
2.1.2 Excitation, Relaxation and Signal Detection 7
2.2 Image Formation 8
2.2.1 Slice Selection 10
2.2.2 Spatial Encoding and k-Space 10
2.2.3 Practical Sampling Considerations 12
2.2.4 Frequency and Phase Encoding 14
2.3 Pulse Sequences and k-Space Trajectories 15
2.3.1 Gradient Echo Pulse Sequence 16
2.3.2 Echo Planar Imaging 17
2.3.3 Non-Cartesian Imaging 17
2.3.4 Spoiling 18
2.3.5 Choice of Flip Angle 19
2.4 Image Reconstruction 19
2.4.1 Gridding 20
2.4.2 NUFFT 22
2.4.3 Parallel Imaging 22
2.4.4 Non-Cartesian Parallel Imaging 24
2.5 Typical MRI Artefacts 27
2.5.1 Chemical Shift Artefact 27
2.5.2 B_0 Inhomogeneity 28
2.5.3 Trajectory Imperfections 28

This chapter provides a short introduction to NMR and MRI image formation, with the main focus on introducing the key vocabulary and concepts relevant to this thesis. More

complete and refined descriptions of MRI can be found in many standard text books (for example by Bernstein et al. [1], Haacke et al. [2] or Nishimura [3]).

2.1 The MR Signal

Magnetic Resonance Imaging (MRI) is a tomographic imaging method based on the principle of Nuclear Magnetic Resonance (NMR), first discovered by Bloch and Purcell in 1946 [4, 5]. All nuclei have a quantum mechanical property called spin, which assumes different values depending on the chemical element and isotope. Particles with a non-zero spin interact with an external magnetic field $\boldsymbol{B_0}$, which forms the basis of NMR. The proton of the hydrogen atom is the most frequently used nucleus in MRI for medical applications, as water is abundant in almost all biological tissues.

2.1.1 NMR Basics

A proton has two possible quantum spin states, parallel alignment (lower energy state) with $\boldsymbol{B_0}$ or anti-parallel alignment (higher-energy state). The frequency associated with an energy transition between the two is given by the Larmor equation $\omega_0 = \gamma \cdot B_0$, where $B_0 = |\boldsymbol{B_0}|$ is the magnetic field strength in Tesla (T) and γ the gyromagnetic ratio [rad/s/T] of the nucleus used. For an ensemble of weakly coupled spins, such as the water proton, the occupancy of the two states can be described by a Boltzmann distribution. In living biological tissue, where the thermal energy is much higher than the energy difference between the two states, there is only a small surplus of protons in the parallel state. Under these assumptions, the maximum magnetisation $\boldsymbol{M}$ available for signal formation can be derived to be proportional to the magnetic field strength B_0 and inversely proportional to the absolute temperature T [K]:

$$M = |\boldsymbol{M}| \approx N \left(\frac{\hbar\gamma}{2}\right)^2 \frac{B_0}{k_B T} \tag{2.1}$$

N corresponds to the number of spins in the ensemble, $k_B = 1.3807 \cdot 10^{-23} J/K$ is the Boltzmann constant, and $\hbar$ corresponds to Planck's constant $h = 6.626 \cdot 10^{-34} J \cdot s$

divided by 2π. The number of spins per unit volume is called the proton density ρ. An important practical conclusion from equation 2.1 is the potential for increased signal at higher magnetic field strengths.

Being able to describe an ensemble of spins in terms of a bulk magnetisation oriented along B_0 allows us to transition from a quantum mechanical to a classical description of MR. Most elements of an MRI experiment can be understood using classical physics only, where the protons are viewed as tiny magnetic dipoles with "intrinsic angular momentum". If a system is in thermal equilibrium, $\boldsymbol{M}$ is perfectly aligned with the $\boldsymbol{B_0}$ axis, which is also referred to as the longitudinal axis or z-axis. The plane perpendicular to the longitudinal axis is referred to as the transverse plane (x-y-plane). If $\boldsymbol{M}$ is made to point in any different direction to $\boldsymbol{B_0}$, the magnetisation experiences torque pushing it towards alignment with $\boldsymbol{B_0}$. The equation of motion for the magnetisation M can be written as:

$$\frac{\mathrm{d}\boldsymbol{M}}{\mathrm{d}t} = \boldsymbol{M} \times \gamma \boldsymbol{B_0}. \tag{2.2}$$

This first order differential equation describes the precession of the magnetisation vector about the $\boldsymbol{B_0}$ axis, with the precession frequency corresponding to the Larmor frequency mentioned above. This behaviour is analogous to the motion of a spinning top in a gravitational field, which is frequently used as an intuitive illustration of Larmor precession.

2.1.2 Excitation, Relaxation and Signal Detection

The precessional motion described above forms the basis of NMR and allows the excitation of the magnetisation by using a transmit coil to create an additional external magnetic field B_1 varying at the resonance Larmor frequency. Because Larmor frequencies are in the radio frequency (RF) range for the B_0 values typically used in MRI, the excitation magnetic field is also called an RF pulse. Applying an excitation pulse causes M to spiral out of alignment with B_0. The angle between the final magnetisation and the longitudinal axis is called the flip angle and is determined by the amplitude and duration of the RF pulse.

In a process called relaxation the magnetisation spirals back into equilibrium. This consists of the transverse component of the magnetisation decaying away (governed by

the time constant T_2) and the longitudinal component returning to thermal equilibrium (governed by the time constant T_1, which is always longer than T_2). The physical processes underlying relaxation are complex but generally speaking they describe the energy loss due to interaction of spins with their environment. Equation 2.2 can be extended to take into account these two relaxation mechanisms to yield the Bloch equation:

$$\frac{\mathrm{d}\boldsymbol{M}}{\mathrm{d}t} = \boldsymbol{M} \times \gamma \boldsymbol{B_0} + \begin{pmatrix} -M_x/T_2 \\ -M_y/T_2 \\ (M_0 - M_z)/T_1 \end{pmatrix}. \tag{2.3}$$

Tissue types have different relaxation times and different proton densities. Knowledge of these properties enables modulation of the signal with the excitation and imaging parameters to achieve different tissue contrasts. If inhomogeneities are present in the magnetic field, the transverse magnetisation decays faster. This decay is governed by the time constant T_2^* and will be discussed in more detail in chapter 3. Besides the relaxation constants and proton density many more factors can generate contrast in MRI, among others the magnetic susceptibility as well as the diffusion of water and blood flow. The contrast mechanism used for fMRI is covered in chapter 3.

The rotating transverse magnetisation generates a change in magnetic flux, which by Faraday's law induces an electromotive force in a coil tuned to the Larmor frequency. The same coil can perform RF transmission and reception, but frequently designated coils are used for the two tasks.

2.2 Image Formation

In order to generate images from the NMR signal we have to spatially localize the signal. There are two main mechanisms (normally used in combination) that allow image formation: (1) spatially selective RF excitation (2) gradient (Fourier) encoding. Both (1) and (2) rely on the use of additional linear gradient magnetic fields (called gradients) to spatially manipulate the resonance frequency. Under the influence of a gradient the magnetic field becomes a function of the position:

$$\boldsymbol{B(r)} = (B_0 + \boldsymbol{G} \cdot \boldsymbol{r})\hat{\boldsymbol{z}}, \text{with } \boldsymbol{G} = \begin{pmatrix} G_x \\ G_y \\ G_z \end{pmatrix} = \begin{pmatrix} \mathrm{d}B_z/\mathrm{d}x \\ \mathrm{d}B_z/\mathrm{d}y \\ \mathrm{d}B_z/\mathrm{d}z \end{pmatrix}, \tag{2.4}$$

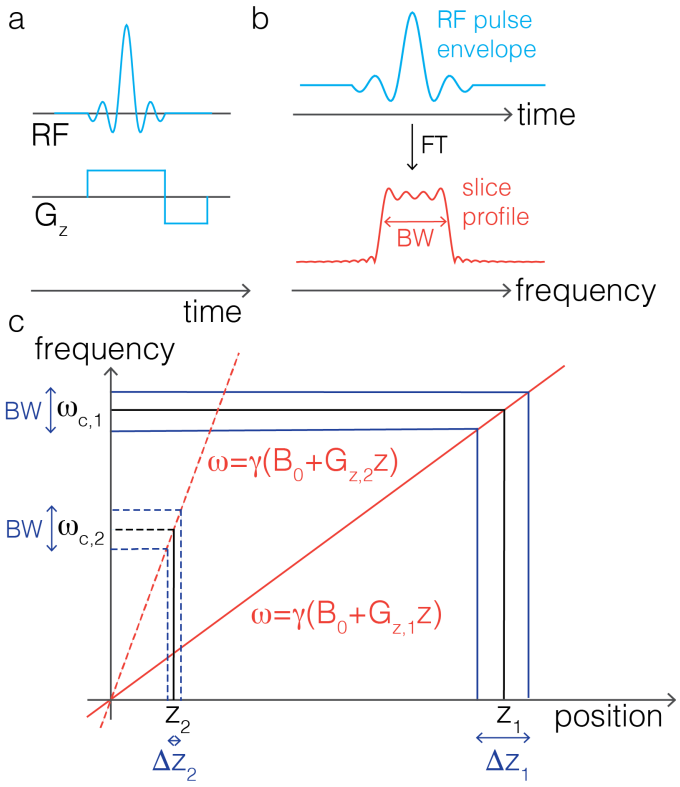


Figure 2.1: Illustration of the concept of slice selection. (a) A RF pulse is applied in combination with a constant gradient (in this example along the z-axis). (b) The slice profile is approximately defined by the Fourier transform of the RF pulse. (c) The slice select gradient amplitude determines the position and thickness of a slice excited by a given RF pulse with bandwidth BW and carrier frequency ω_c. A stronger gradient $G_{z,2}$ excites a thinner slice than the weaker $G_{z,1}$ gradient.

where $\boldsymbol{r}$ corresponds to the position vector and $\hat{\boldsymbol{z}}$ to the unit vector in the z-direction. It is important to note that the magnetic field generated by the gradient coils points in the z-direction, but its strength now varies with position along all three directions (under independent control in time). The time domain signal under the influence of a gradient can be Fourier transformed to obtain the signal contributions as a function of frequency. The mechanism of localising the signal to position was developed by Lauterbur [6], Ernst [7] and Mansfield [8] in the mid 1970s, and was originally called NMR Fourier zeugmatography. The use of multiple receive coils can also aid encoding, which is discussed in the parallel imaging section in 2.4.3.

2.2.1 Slice Selection

The first step in an MR experiment consists of excitation. The entire sample can be excited using non-selective excitation but more typically a part of the signal localisation is done by spatially selective excitation. The goal of this is to choose which part of the sample, typically a thin slice (or thicker slab), should contribute to the MR signal. This is achieved by applying a gradient in the direction perpendicular to the slice during the RF pulse (Fig. 2.1). The carrier frequency of the RF pulse defines the position of the slice, while the combination of the RF bandwidth used and the gradient strength determine the thickness of the excited volume (Fig. 2.1c). Due to the finite length of RF pulses, the achieved slice profile is broadened and contains imperfections such as additional side lobes and ripples in the flat top. A measure of the selectivity of the pulse or steepness of the slice profile is its time-bandwidth product $TBWP = T \cdot \Delta f$, where T corresponds to the duration of the RF pulse and Δf can be estimated as the inverse of the full-width at half maximum (FWHM) of the central lobe. There are many practical considerations, which need to be taken into account for RF pulse design, which cannot be covered here. However, an important issue related to excitation is how to avoid the chemical shift artefact related to fat signal, which is briefly discussed in chapter 2.5.1.

2.2.2 Spatial Encoding and k-Space

In this section we are neglecting relaxation and off-resonance effects, which is not relevant to the understanding of image encoding. Assuming that the initial phase after excitation is zero and ignoring the static component of the field (demodulating the signal by the Larmor frequency) we can write the signal picked up by the RF coil as:

$$s(t) \propto \int_V \rho(\boldsymbol{r}) e^{-i2\pi\phi(\boldsymbol{r},t)} \mathrm{d}\boldsymbol{r} \tag{2.5}$$

where V is the volume of the excited spins and $\rho(\boldsymbol{r})$ is the proton density of the object at the location $\boldsymbol{r}$. ϕ represents the phase of the spin precession, which amounts to the time-integral of the frequency since excitation $\phi(\boldsymbol{r},t) = \gamma \int_0^t \omega(\boldsymbol{r},t')\mathrm{d}t'$. Fourier encoding

is achieved by applying additional magnetic fields, as described previously. This allows spatial modulation of the precession frequency within the volume of interest, creating a well-defined relationship between any position r and the frequency.

$$\omega(\boldsymbol{r}, t') = \omega_0 + \gamma \boldsymbol{G}(t) \cdot \boldsymbol{r} \tag{2.6}$$

The associated phase accumulation is:

$$\phi(\boldsymbol{r}, t) = \int_0^t \omega(\boldsymbol{r}, t') \mathrm{d}t' = \boldsymbol{r} \cdot \gamma \int_0^t \boldsymbol{G}(t') \mathrm{d}t' \stackrel{\text{def}}{=} \boldsymbol{r} \cdot \boldsymbol{k}(t), \tag{2.7}$$

where $\boldsymbol{k}(t)$ corresponds to the linear phase accumulated across the sample. With this definition equation 2.5 becomes:

$$s(t) \propto \int_V \rho(\boldsymbol{r}) e^{-i2\pi \boldsymbol{r} \cdot \boldsymbol{k}} \mathrm{d}\boldsymbol{r} \tag{2.8}$$

This corresponds to a Fourier transform and the domain spanned by $\boldsymbol{k}$ is called k-space, which is a reciprocal space to image space. Acquiring a specific location in k-space corresponds to sampling a specific spatial frequency of the spin distribution of the object. With perfect knowledge of k-space (entire spatial frequency domain), an image of the transverse magnetisation could be perfectly recovered using an inverse Fourier transform. The k-space formalism was introduced for MRI in the 1980s by Brown et al. [9], Ljunggren et al. [10] and Twieg et al. [11].

The centre of k-space contains the "DC" component of the image, i.e. information with zero spatial frequency. Low spatial frequencies are located close to the k-space centre and the higher spatial frequency at the periphery. An image reconstructed from a central region lacks edge detail, while images reconstructed using only the outer regions of k-space are missing the contrast and slowly varying image information.

When position encoding using gradients is performed in three directions we call this 3D MRI and the excited volume, which is typically thicker, is called a slab. In 3D imaging the different k-space planes are sometimes called partitions instead of slices. In 2D MRI slice selection takes care of the spatial encoding in one direction (by exciting

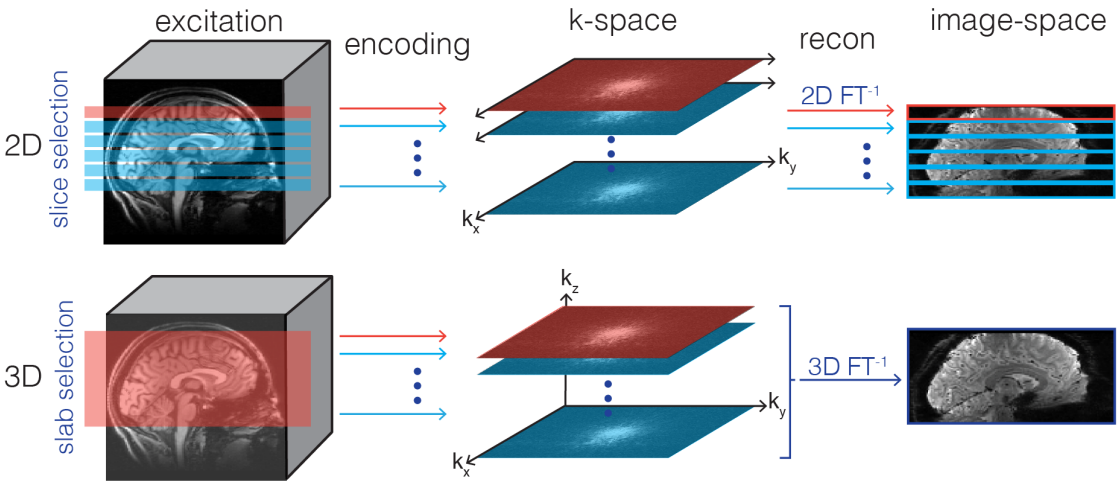


Figure 2.2: Illustration of the difference between 2D (top) and 3D (bottom) MRI using the example of 2D and 3D EPI. In 2D EPI, each TR one plane of k-space is excited (shown in red), encoded and reconstructed, using a 2D inverse FT. This process is repeated for all slices until the stack of 2D images forms the entire 3D volume. In 3D EPI each TR the entire slab is excited and one plane of k-space is encoded (red). Once the entire 3D k-space has been built up the 3D image is reconstructed using a 3D inverse FT.

a series of slices with the thickness of the desired resolution) and Fourier encoding is used in two dimensions for every slice (Fig. 2.2).

Due to symmetry properties of the Fourier transform, half the information in the k-space is related to the other half by the complex conjugate (if the object is real). Therefore, if the object is real or if the object phase is known, it is possible to only sample a part of k-space; this method is called Partial Fourier imaging [12].

2.2.3 Practical Sampling Considerations

Practically, k-space cannot be sampled continuously and to infinite extent; this has a number of important consequences for MRI. For simplicity all the equations are noted for a 1D case (using the example $k = k_x, \Delta k = \Delta k_x$ and $\mathrm{FOV} = \mathrm{FOV}_x$) but the extension to 2D and 3D sampling is straightforward. For uniform sampling the sampling process can be expressed as a multiplication of the continuous k-space $s(k)$ with a Dirac comb $\frac{1}{\Delta k}\mathrm{III}\left(\frac{k}{\Delta k}\right) = \sum_i \delta(k - i \cdot \Delta k)$ yielding the sampled k-space $s_{\mathrm{sampled}}(k)$ as

$$s_{\mathrm{sampled}}(k) = s(k) \cdot \frac{1}{\Delta k}\mathrm{III}\left(\frac{k}{\Delta k}\right). \tag{2.9}$$

Taking into account that sampling occurs over a finite extent we can expand the equation above by multiplying it with a rectangular function $\sqcap$ of $2 \cdot k_{max}$ width:

$$s_{\text{sampled}}(k) = s(k) \cdot \frac{1}{\Delta k} \text{Ш}\left(\frac{k}{\Delta k}\right) \cdot \sqcap\left(\frac{k}{2k_{max}}\right) \tag{2.10}$$

We discussed above that image space and k-space are related via the Fourier transform $\mathcal{F}$. Inverse Fourier transforming equation 2.10, the resultant function $\mathcal{F}^{-1}\{s_{\text{sampled}}(k)\} = \rho_{\text{sampled}}(x)$ is

$$\rho_{\text{sampled}}(x) = \rho(x) * \text{Ш}(\Delta k x) * 2k_{max}\text{sinc}(2k_{max}x), \tag{2.11}$$

where $\mathcal{F}^{-1}\{s(k)\} = \rho(x)$. The train of delta functions from the Dirac comb (with a spacing of $1/\Delta$k) is convolved with a sinc function resulting in a train of equally spaced sinc functions. This has two important implications for MRI: Firstly, the sampling process generates infinite replica of the sampled object in image space. For a desired field of view (FOV) in image space, the sampling increments have to meet the condition

$$\Delta k = \frac{1}{FOV} \tag{2.12}$$

to prevent overlap between adjacent replica (aliasing). This requirement is equivalent to the Nyquist sampling theorem, which states that the sampling rate has to be at least twice the (highest) frequency to be recovered. In MRI aliasing results in a "wrap-around" artefact because of mis-mapping of high frequencies onto lower frequencies, leading to spatial ambiguity.

Secondly, the convolution with the sinc function in 2.11 limits the ability to resolve structures and therefore defines the sampling point-spread-function of the acquisition. A common measure of resolution is the FWHM of the main lobe of the sinc function, but for simplicity the first zero crossing $1/2k_{max}$ is used. This defines the spatial resolution of the acquisition

$$\Delta x = \frac{1}{2k_{max}}, \tag{2.13}$$

which is equivalent to the “Fourier resolution” defined as the FOV divided by the number of sampling points (Fig. 2.3). The truncation of k-space leads to the Gibbs ringing artefact,

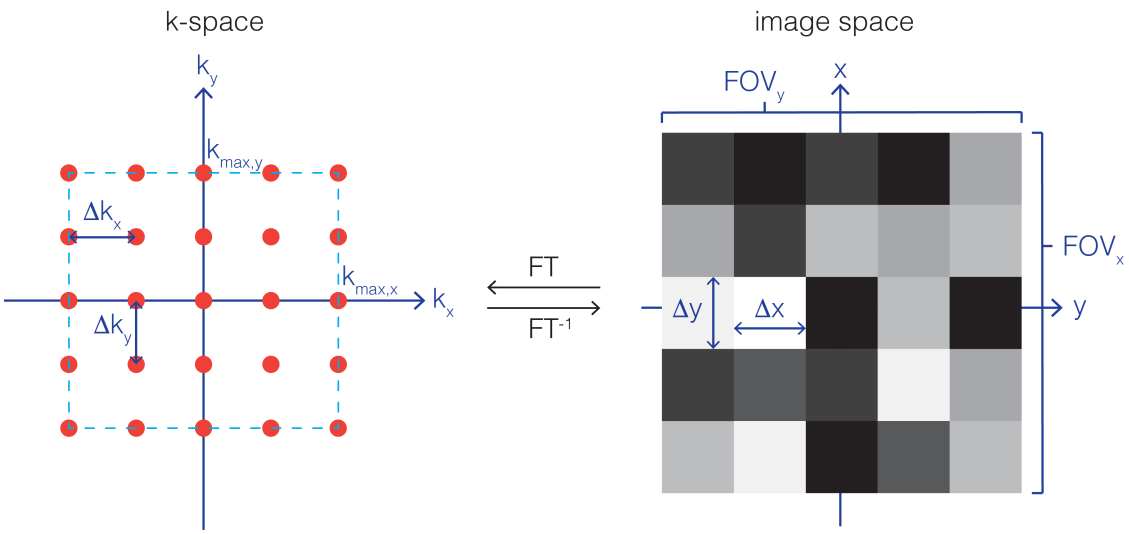


Figure 2.3: Illustration of a Cartesian 2D k-space and image space. The spacing between the k-space locations Δk and the sampling extent k_{max} in k-space define the FOV and the resolution in image space.

which manifests itself as characteristic ripples around image boundaries. The artefact can be reduced by applying an apodisation filter to k-space prior to reconstruction or by increasing the spatial resolution. It should be noted, that the effective voxel size is also influenced by signal decay over the acquisition window, but conventionally the nominal voxel size (as defined above) is reported.

2.2.4 Frequency and Phase Encoding

In most commonly used encoding schemes one dimension of Fourier encoding is called the frequency encoding direction or the readout direction, while all remaining (1-2) directions are referred to as phase encoding axes. Mathematically, frequency and phase encoding are the same processes as in both cases the position is encoded by the amount of phase accumulation. The differentiation comes from the way the phase modulation is achieved in practice. In frequency encoding a constant gradient is played which creates a position dependent resonance frequency, as described for slice selection above. This gradient is called the readout gradient, because the signal is sampled (read out) while the readout gradient is on. In phase encoding a gradient is only turned on for a brief period of time prior to the readout. During this duration spins at different spatial locations accrue

different amount of phase, which they keep as a "tag" when they return to their original precession frequency. This process has to be repeated multiple times, every time using a different gradient moment and introducing a different amount of phase.

Due to the differences in sampling, aliasing can be mitigated more easily in the frequency encoding direction. Along the readout direction the signal acquisition samples an analog signal and the sampling speed is only limited by the dwell time (i.e. the interval between digitised samples) of the analogue to digital converter (ADC), which is typically in the range of 10s of microseconds. Therefore the readout direction of k-space data can be filtered prior to reconstruction to prevent aliasing. This is not possible in the phase encoding directions, where the sample spacing is controlled by the gradient moment applied.

2.3 Pulse Sequences and k-Space Trajectories

Above we mentioned the standard acquisition strategy using Fourier encoding with readout and phase encoding gradients. However, there are many different ways to spatially encode the MR signal, each coming with a unique set of benefits and challenges. The pattern and order in which k-space is covered is called the k-space trajectory. A pulse sequence contains the information about all the hardware components of the MR system (excitation, gradients and reception) and their respective timings. If (2D or 3D) k-space is covered over multiple readouts this is called multi-shot imaging whereas collecting the entire k-space after one excitation is called single-shot imaging. When samples lie on a rectangular grid in k-space this is denoted Cartesian imaging, all other trajectories are non-Cartesian.

The time between subsequent RF excitation of the same spins is called the repetition time TR, while the time required to acquire an entire 3D image is called the volume TR (TR_{vol}). The echo time (TE) refers to the time between the excitation and the peak of the signal induced in the coil.

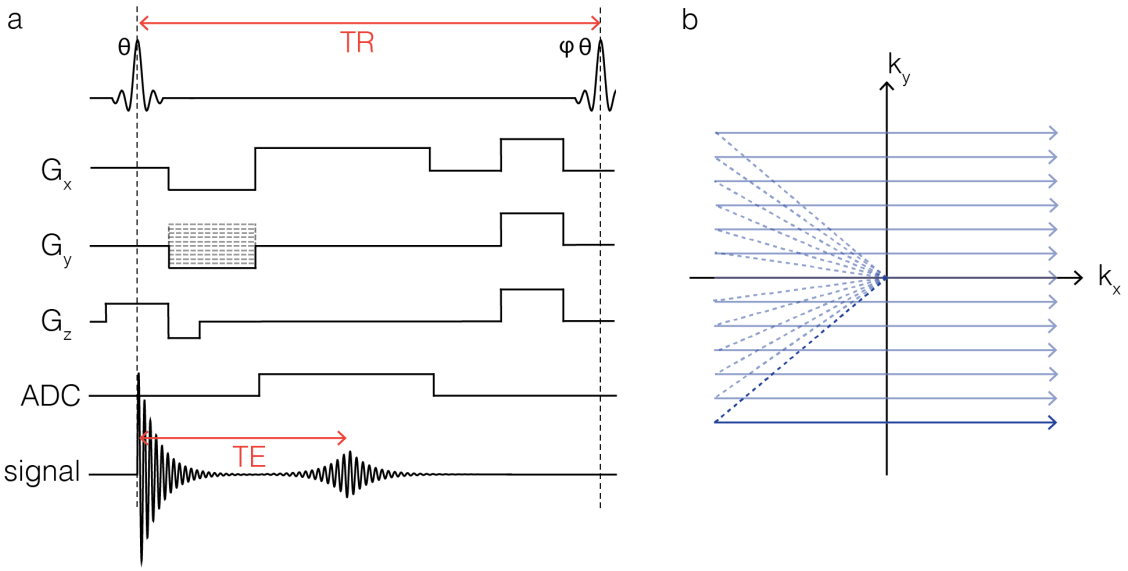


Figure 2.4: Illustration of the basic spoiled 2D GRE pulse sequence and the corresponding k-space trajectory. In this example slice encoding is along z, frequency encoding along x and the phase encoding along the y-direction. (a) The diagram shows one TR block, which samples one line of k-space. This block is repeated multiple times with G_y changing each TR as indicated by the dashed lines. In this example the residual transverse magnetisation is spoiled for every TR. (b) The resulting k-space trajectory traces out a different line along k_x for each TR using gradient and RF spoiling. A solid blue line shows the line collected in the first TR, and all subsequent lines are transparent. The dashed lines correspond to the prephasers, during which no data is acquired.

2.3.1 Gradient Echo Pulse Sequence

One of the most basic MRI sequences is the gradient recalled echo (GRE) sequence (originally spin warp sequence [13], shown in Fig. 2.4). Following excitation, a gradient in the readout direction (pre-winder) and the phase encode gradient are played to bring the k-space vector to the edge of the k-space line to be sampled. Then, a continuous readout gradient is applied while the ADC is turned on, sampling one line of k-space. To cover the whole of k-space this is repeated with slightly modified phase encoding gradient strength. This most basic acquisition scheme has the benefit that the phase of the signal is reset prior to collecting each line of k-space. This is helpful, since there are several practical problems that can cause loss of control over the phase of the signal over time. The principle drawback is that it is very slow and therefore not suitable to image dynamic processes. When an additional 180° RF pulse, which inverts the phase of the magnetisation, is inserted into this sequence it is called a spin echo acquisition. Spin

echo sequences have different properties to gradient echo sequences and are not covered in this chapter as we are only using GRE sequences in this thesis

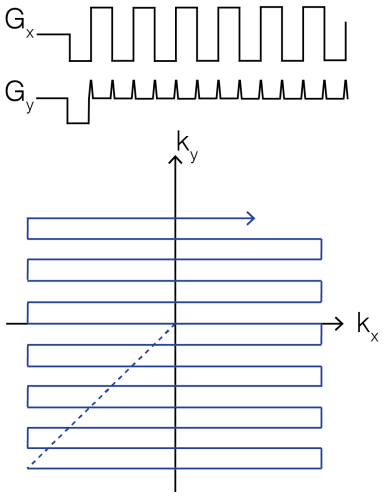


Figure 2.5: Gradient waveforms (top) and k-space trajectory (bottom) for EPI. After the prephaser (dashed line), a readout gradient G_x is used to sample the first line of k-space. Subsequently, a G_y gradient "blip" moves the k-space vector to the next line. This is repeated until the desired amount of k-space has been traversed.

2.3.2 Echo Planar Imaging

The echo-planar imaging (EPI) readout (first suggested by Mansfield in 1977 [8]) speeds up imaging by covering multiple lines of k-space per TR (Fig. 2.5). As above the readout gradient is played while the ADC is on, collecting signal from a single line of k-space. Then the phase encoding gradient is played very briefly ("blipped") with an area sufficient to move to the adjacent desired k-space line, and this process is repeated until the desired number of k-space lines are acquired (Fig. 2.5). This type of readout has the benefit of reduced acquisition time at the expense of more image artefacts such as distortion from magnetic field inhomogeneity. Typically, all the lines of k-space are acquired after one excitation pulse (single shot EPI), resulting in TRs on the order of 10s of milliseconds.

2.3.3 Non-Cartesian Imaging

The trajectories described above collect equidistant samples along every dimension of k-space. However, there are numerous other ways to sample k-space, for example using a

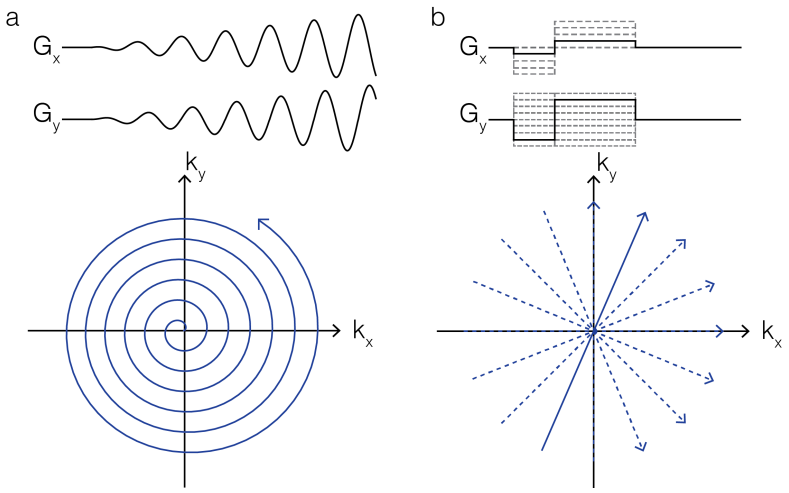


Figure 2.6: Gradient waveforms and trajectories for two example 2D non-Cartesian acquisitions: (a) Single shot spiral (b) Multi-shot radial. The solid lines show the waveforms and acquired k-space for one shot; all other shots are depicted with dashed lines.

radial or spiral pattern (Fig. 2.6). Radial trajectories for example produce less coherent artefacts and have reduced motion sensitivity. Spiral trajectories are very efficient and provide an alternative to EPI for fast imaging. Generally the demands on the accuracy of the gradient hardware are higher and the reconstruction is more complicated (see 5.3). Sampling SNR efficiency of non-Cartesian trajectories is intrinsically lower than for uniform sampling [14]. Radial imaging for example requires $\frac{\pi}{2}$ more samples than an equivalent Cartesian acquisition to meet the Nyquist criterion at the edges of k-space. However, due to the more incoherent aliasing artefacts many non-Cartesian trajectories tolerate higher under-sampling factors. Because of the challenges associated with non-Cartesian imaging most MR acquisitions are performed using Cartesian sampling.

2.3.4 Spoiling

Frequently in MRI the TR is not long enough for the transverse magnetisation to naturally decay to zero by the end of the TR. The goal of spoiling is to make sure that immediately before each RF-pulse, the transverse component of the magnetisation is zero, to avoid any signal contribution from residual transverse magnetisation. Two main spoiling strategies exist, gradient spoiling and RF spoiling. In gradient spoiling an additional gradient is added that dephases the spins, causing the residual transverse signal to decay. However,

the transverse magnetisation is not destroyed and can be refocused in subsequent TRs (to produce a spin echo or stimulated echo). In RF spoiling the RF phase is varied according to a predefined scheme, introducing a phase offset between the freshly excited and persistent magnetisation. If a good phase cycling scheme is chosen the magnetisation for preceeding TRs will cancel while magnetisation excited by the most recent RF pulse is preserved. A popular phase cycling scheme is a quadratic scheme proposed by Zur et al. [15] $\phi_n = 0.5 \cdot n(n+1)\Delta\phi$, where ϕ_n is the phase of the nth RF pulse and $\Delta\phi$ is the phase increment, with a popular choice being 117°. Spoiling strategies can be optimised for specific acquisition strategies. For example for radial imaging it has been shown that a scheme using random RF phase and gradient moments provides good spoiling [16].

2.3.5 Choice of Flip Angle

In the short TR regime $TR \ll T_1$ the longitudinal magnetisation cannot fully return to equilibrium, before the next excitation. Instead, after a number of TRs a steady state magnetisation is reached, which depends on the TR, T_1 and flip angle θ. Assuming perfect spoiling, the optimal flip angle for a gradient echo sequence is defined by the Ernst angle $\theta_E = cos^{-1}(e^{-TR/T_1})$ [17]. Using this flip angle provides the maximum possible signal for a given TR and T_1.

2.4 Image Reconstruction

As described above, the MR image and the raw k-space data are related via the Fourier transform. Calculating the inverse discrete Fourier transform (DFT) is very time consuming, therefore image reconstruction is usually performed using the inverse Fast Fourier Transform (FFT) [18], an efficient implementation of the inverse DFT. A requirement for the FFT is that the input samples are equidistant in each dimension. Thus, in the case of non-Cartesian data, other methods must be implemented. Some of the first MR acquisitions used radial sampling and employed back projection [6] to reconstruct the image (inspired by the then already established computed tomography

reconstruction techniques). However, this approach works only for the radial case. More generally, non-Cartesian reconstruction is tackled using an initial interpolation step in k-space, followed by the regular FFT reconstruction. The most commonly used algorithm for this is called gridding (sometimes called regridding) and is described in 2.4.1 below.

On a side note, linear gradient fields form the kernel of the Fourier transform integral. Hence, in the case of deviation from the linearity of gradients, either deliberate or due to gradient hardware imperfection, the Fourier transform no longer accurately describes the transformation between the raw-data space and the image space. For nonlinear gradients a generalised reconstruction framework using more general basis functions needs to be used.

2.4.1 Gridding

In short, gridding [19, 20] interpolates data from an arbitrary sampling pattern onto a uniform grid (Fig. 2.7). The first step in gridding usually consists of compensating for the non-uniform k-space weighting due to a non-Cartesian sample distribution (density pre-compensation). The non-uniform k-space is then convolved with a gridding kernel of choice and evaluated at the Cartesian target grid locations (typically choosing an oversampled target grid). Subsequently, the image is reconstructed by applying the inverse FFT. In a final step any shading of the image, introduced by the kernel convolution, is compensated via multiplication with a de-apodisation function (inverse FT of the gridding kernel) and the oversampling is removed by cropping the image.

There are two main practical choices which influence the quality of the gridding reconstruction: the choice of the gridding kernel and the amount of grid-oversampling. In theory an infinite sinc function, corresponding to a rectangular function in image space, would be the ideal kernel. In practice the choice of kernel is a trade-off between artefact level, memory and speed. Currently, the Kaiser-Bessel kernel is the preferred kernel in the MR field as it provides best results at minimal oversampling factors [21]. Typically, the 2D or 3D gridding kernels are designed to be separable (the product of 1D kernels), but it is also possible to design kernels explicitly for higher dimensions. Oversampling the

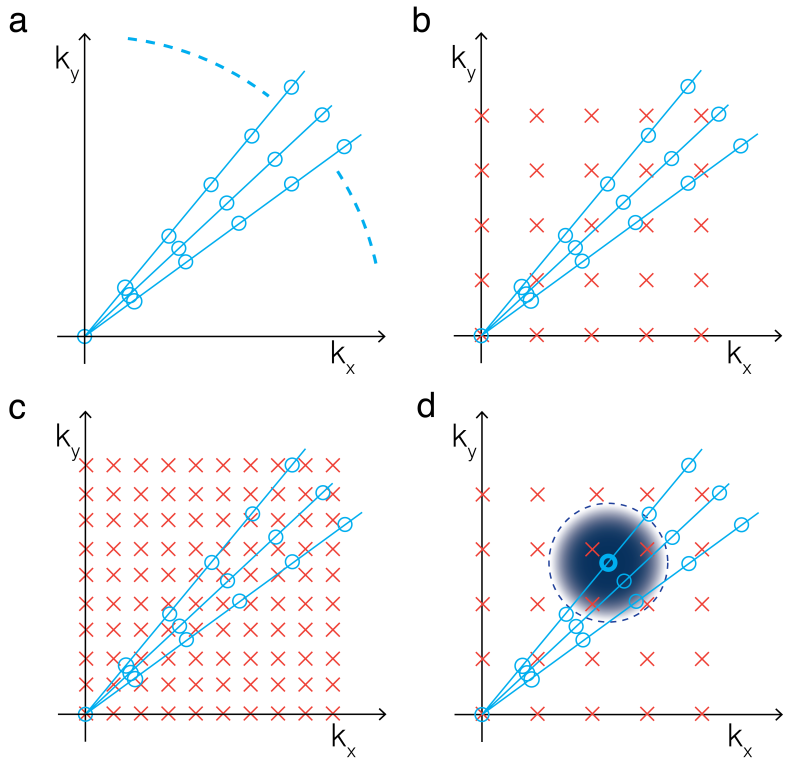


Figure 2.7: Illustration of the gridding process, exemplified using a radial trajectory. (a) The data points (depicted with blue circles) are acquired along radial spokes. (b) The target Cartesian grid is shown with red crosses. (c) The choice of the target grid is arbitrary and a denser target grid (compared to (b)) can be chosen. (d) In the gridding algorithm each data point is convolved with a gridding kernel (shown in dark blue), and that convolution evaluated is at the target grid points.

Cartesian grid reduces gridding artefacts, because the image replicas generated by the re-sampling are pushed further out in image space, hence reducing overlap of potential side lobes with the main image. Grid oversampling also reduces shading of the image.

Inverse gridding calculates the non-uniform k-space data from the (Cartesian sampled) image. This is important for iterative reconstruction algorithms, which alternate between image space and k-space. Iterative reconstruction methods for MRI have become more important in recent years, mainly due to the development of more sophisticated parallel imaging methods as well as to cope with problems such as field inhomogeneity and incomplete sampling (for example using compressed sensing).

2.4.2 NUFFT

The non-uniform FFT (NUFFT) is essentially a generalized version of the gridding algorithm. It uses similar ideas to gridding, but provides a fast implementation and is therefore better suited for incorporation in iterative algorithms. The most popular implementation used by the MR community is incorporated in the image reconstruction toolbox by Fessler et al. [22].

2.4.3 Parallel Imaging

In 1990 Roemer et al. [23] introduced the phased array coil for signal reception in MRI. These coils consist of a set of smaller coils, whose signals are combined to cover a larger volume, which can provide a SNR advantage due to the use of smaller coils. Each channel records the underlying magnetisation weighted by the coil's sensitivity profile (each coil is only sensitive to magnetisation close to the coil, see Fig. 2.8). The goal of parallel imaging (originally called partially parallel imaging) is to use this additional spatial information to reduce the amount of samples required for spatial encoding. Needing fewer samples to form the image can be used to acquire data with shorter TE/TR and reduced distortion, depending on what is important for the application. It is especially important in fMRI to allow optimal echo time and faster volume acquisition times (see chapter 3).

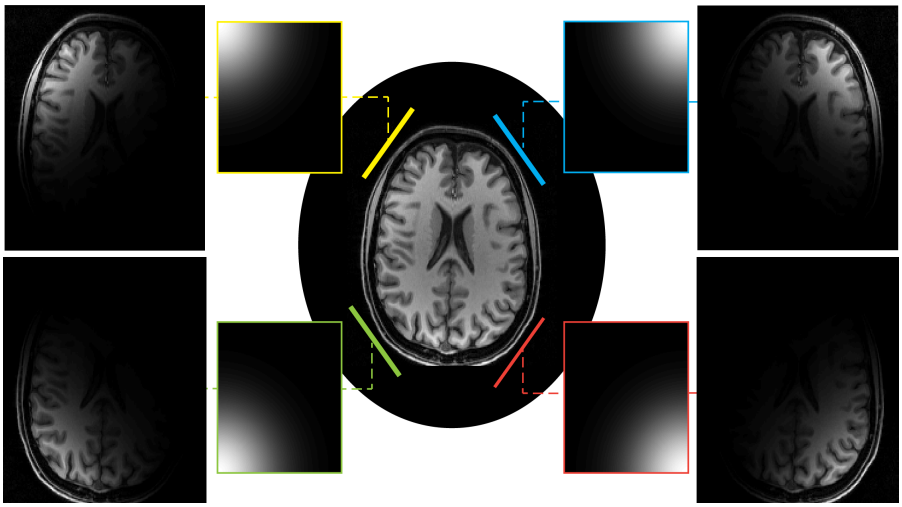


Figure 2.8: Illustration of acquisition using multiple receive coils. The signal from each channel is weighed by the coil's sensitivity profile.

The amount by which the number of acquired samples is reduced with respect to a fully sampled acquisition is called the acceleration factor or reduction factor R. The SNR of the reconstructed images is reduced by a factor $\sqrt{R}$, due to the reduction in Fourier averaging, and by the geometry factor (short g-factor), which takes into account the parallel imaging performance of the coil geometry:

$$SNR_{\text{accel}} = \frac{SNR_{\text{full}}}{\sqrt{R} \cdot g} \tag{2.14}$$

The g-factor varies spatially and depends on the sampling pattern, the coil sensitivities as well as the reconstruction used. A large range of parallel imaging techniques has been developed over the last 25 years (see [24] for a review), which broadly speaking fall into two families of methods: Firstly image space based parallel imaging methods like SENSitivity Encoding (SENSE) [25] and secondly k-space based methods like Simultaneous Acquisition of Spatial Harmonics (SMASH) [26] or GeneRalized Autocalibrating Partial Parallel Acquisition (GRAPPA) [27].

In SENSE, images are generated from the under-sampled k-space data containing aliasing. Knowledge of the coil sensitivity profiles allows removing the aliasing from the images.

In GRAPPA the under-sampling problem is solved in k-space. The missing data points are synthesised using a kernel, which has been calibrated on fully sampled data (the so called auto calibration signal (ACS) or reference data). The kernel consists of the GRAPPA weights g, which are introduced in the following sum over the individual coil signals:

$$S_n(k_x, k_y + r\Delta k_y) = \sum_{n'=1}^{N_c} \sum_{a=-N_x^-}^{N_x^+} \sum_{b=-N_y^-}^{N_y^+} g_{n,n'}^{a,b,r} \cdot S_{n'}(k_x + a\Delta k_x, k_y + (bR+1)\Delta k_y) \tag{2.15}$$

In this linear combination, the index n' counts through the individual coils (with the total number of coils equal to N_c). The indices a and b count through each reconstruction block of size $[N_x^+ + N_x^- + 1, N_y^+ + N_y^- + 1] = [N_x, N_y]$. R is the reduction factor and $g_{n,n'}^{a,b,r}$ corresponds to the GRAPPA kernel weights for coil n and the r$^{\text{th}}$ missing line of the block.

This equation is first used to calibrate the GRAPPA kernel, i.e. find the GRAPPA weights and subsequently to find missing k-space points in the undersampled data. Equation

2.15 can be expressed more compactly in matrix notation:

$$\boldsymbol{S_i} = \boldsymbol{GS} \tag{2.16}$$

In the ACS data the target points $\boldsymbol{S_i}$ and the source points $\boldsymbol{S}$ are known, therefore the GRAPPA weights $\boldsymbol{G}$ can be estimated using

$$\hat{\boldsymbol{G}} = (\boldsymbol{S^H S})^{-1} \boldsymbol{S^H S_i}, \tag{2.17}$$

where $\hat{\boldsymbol{G}}$ is the linear least-squares estimate of $\boldsymbol{G}$. After calibration the missing data $\boldsymbol{S_i'}$ can be estimated using the acquired data $\boldsymbol{S'}$:

$$\boldsymbol{S_i'} = \hat{\boldsymbol{G}} \boldsymbol{S'} \tag{2.18}$$

This yields the complete k-space for each coil, which are individually reconstructed and then combined; for example using sum-of-squares coil combination.

Theoretically, the SENSE and GRAPPA concepts are equivalent [28]. The multiplication of the image with a coil sensitivity profile corresponds to the convolution of the k-space with the Fourier transform of the coil sensitivity profile (which is the theoretical ideal GRAPPA kernel). If the coil sensitivities are slowly varying the use of small GRAPPA kernels is justified. In practice, SENSE is considered the method of choice when used with exact sensitivities, otherwise GRAPPA frequently performs better[27, 29].

2.4.4 Non-Cartesian Parallel Imaging

Many non-Cartesian acquisition strategies tolerate low levels of undersampling without detrimental artefacts. This is because of the more distributed sampling point spread function and because many non-Cartesian trajectories over-sample the centre of k-space. However, for larger reduction factors the coil information has to be used to generate images of acceptable quality. The more complex aliasing patterns make parallel imaging more complicated; for example the unfolding matrices in SENSE become extremely large, because every pixel can be aliased onto almost every other pixel. Due to this increased complexity, most non-Cartesian parallel imaging methods use iterative approaches. Many

methods are also highly specific to the sampling method used (for example radial GRAPPA [30]). A comprehensive overview of the current methods available is beyond the scope of this chapter, but can be found for example in [31]. The iTerative Self-consistent Parallel Imaging Reconstruction (SPIRiT) [32] algorithm was used in this thesis and is therefore briefly described here below.

SPIRiT can be considered an extension of GRAPPA and is also based on calibrating a shift invariant kernel on calibration data. A key difference of SPIRiT, compared to GRAPPA, is that it uses the sampled k-space points as well as the estimated values (from the previous iteration) to synthesise the non-acquired samples. The algorithm is based on two concepts: (a) self-consistency of the reconstructed data with the acquired k-space data and (b) calibration consistency. The SPIRiT reconstruction aims to minimize the calibration consistency term $||(\boldsymbol{G}-\boldsymbol{1})\boldsymbol{x}||^2$, while fulfilling the data consistency condition $||\boldsymbol{y}-\boldsymbol{D}\boldsymbol{x}||^2 < \epsilon$. In Lagrangian form, the SPIRiT reconstruction can be written as:

$$\text{argmin}_{\boldsymbol{x}}(||\boldsymbol{y}-\boldsymbol{D}\boldsymbol{x}||^2 + \lambda(\epsilon)\cdot||(\boldsymbol{G}-\boldsymbol{1})\boldsymbol{x}||^2). \tag{2.19}$$

The vector $\boldsymbol{x}$ contains the data sampled onto a Cartesian grid, whereas $\boldsymbol{y}$ contains the k-space samples along the original trajectory. $\boldsymbol{D}$ is the sampling operator, which selects only acquired k-space locations. In the case of a non-Cartesian trajectory the operator $\boldsymbol{D}$ also contains a gridding operator. $\boldsymbol{G}$ is an operator, which applies the SPIRiT kernel weights and performs the required summations (see Fig. 2.9 for a an overview of the SPIRiT algorithm). Note that the SPIRiT kernel is always applied on the Cartesian data. The parameter λ controls the trade-off between calibration consistency and data consistency: smaller λ values favour data consistency, while larger values more heavily enforce calibration consistency.

The kernel weights are calibrated on a Nyquist sampled region of k-space. This can either be performed using separately acquired (Cartesian) data, or a densely sampled region can be interpolated onto a Cartesian grid. The SPIRiT reconstruction can be performed either in image space or in k-space and a number of different approaches can be used to minimise the objective function. In this thesis we used the image-space variant

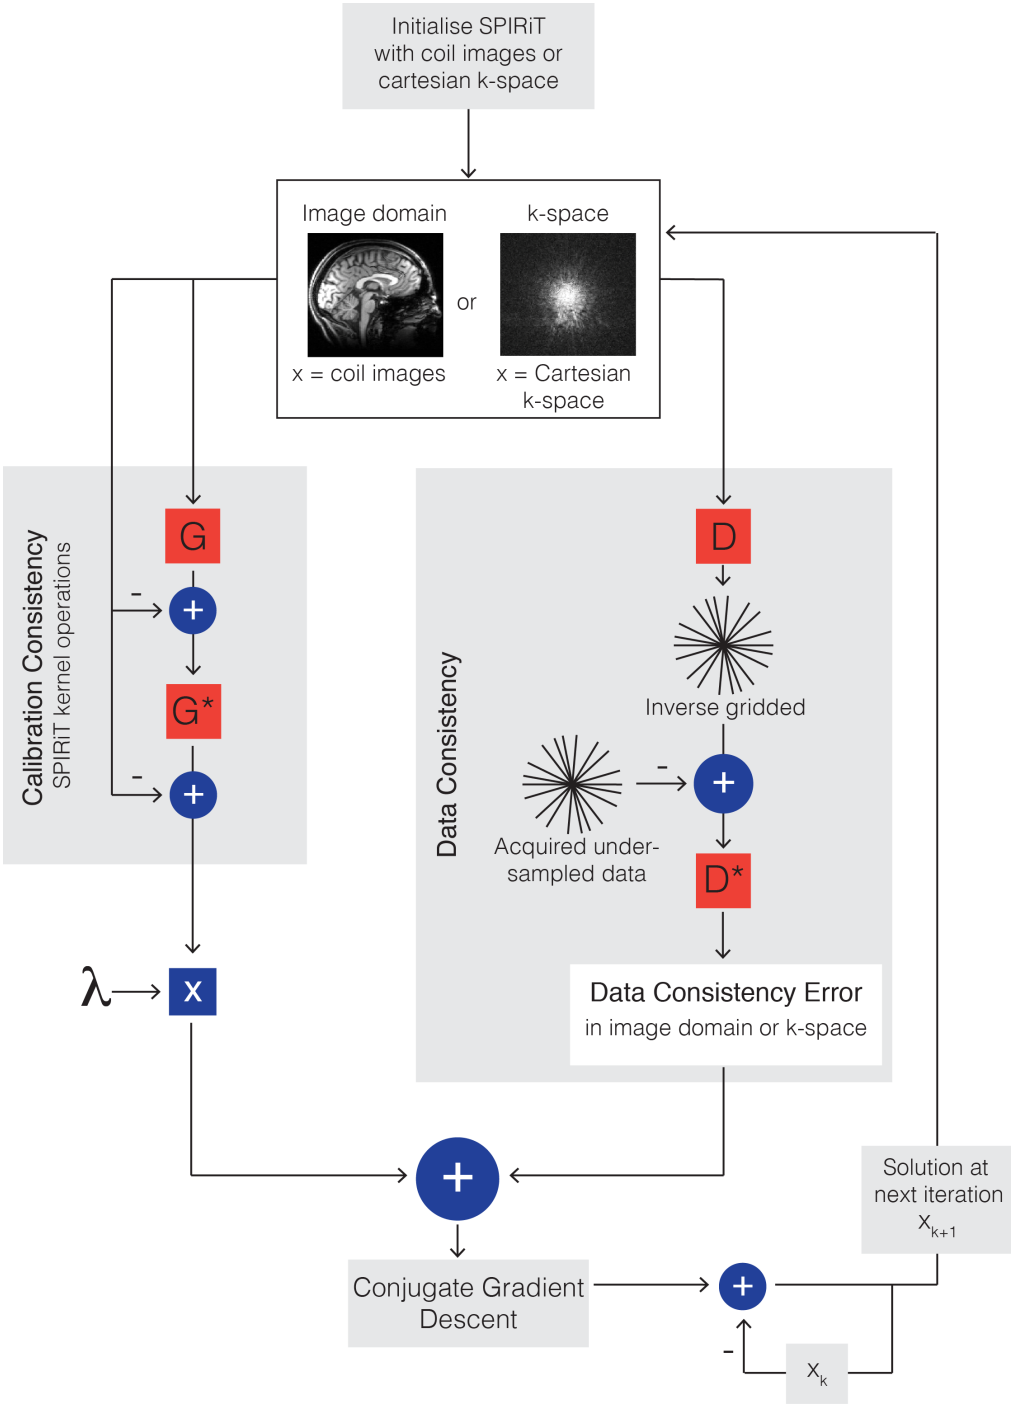


Figure 2.9: Flowchart of the CG SPIRiT algorithm for non-Cartesian samples, showing both the k-space domain and image space domain approach. The reconstruction is initialised with either Cartesian k-space data or images for each coil. The operator $\boldsymbol{D}$ transforms the input data to k-space samples along the non-Cartesian trajectory used. $\boldsymbol{D}^*$ is its adjoint operator. The calibration consistency step is performed using the operator, $\boldsymbol{G}$, and its adjoint, $\boldsymbol{G}^*$. Figure is modeled after similar illustrations in [31, 32].

of the algorithm, using a conjugate gradient (CG) descent approach: The algorithm is typically initialised with the zero-filled reconstruction of the under-sampled data. Every iteration the CG descent returns an improved estimate of the coil images, which are fed into the next iteration. The algorithm terminates when the maximum number of iterations or another stopping criterion is met.

2.5 Typical MRI Artefacts

This section introduces some commonly seen artefacts in MRI. Artefacts caused by subject motion and physiological noise are covered in chapter 3.

2.5.1 Chemical Shift Artefact

The effective magnetic field experienced by a nucleus depends on its chemical environment as well as a process called electron shielding, which can alter the Larmor frequency of the nucleus. While the chemical shift provides information in NMR spectroscopy, in MRI we use frequency to encode position and therefore chemical shift is considered a nuisance. Fat tissue has a relatively large chemical shift, therefore "fat suppression" techniques are normally used to avoid misplaced fat signal.

The commonly used fat saturation (FatSat) scheme is a variant of the chemical shift selective (CHESS) [33] method. FatSat works by applying a RF-pulse tuned to resonance frequency of fat, rotating the fat magnetisation into the transverse plane. The fat magnetisation is then spoiled using crusher gradients along one or more axes, leaving no magnetisation from the fat available for the subsequent excitation.

Binomial pulses are a class of composite pulses, which can be used to selectively excite the water protons [34]; typically called water excitation (WE) pulses. The sum of their flip angles corresponds to the desired flip angle, while their flip angle ratios follow a binomial expansion $(a+b)^n$, $n \in \mathbb{N}$. For example, a 1-1 pair, a 1-2-1 triplet or a 1-3-3-1 quadruplet can be used. Longer pulse chains are more selective in the frequency domain

but take longer to play. Binomial pulses are especially useful at high fields because they are relatively insensitive to B_1 inhomogeneities.

2.5.2 B_0 Inhomogeneity

Placing an object in a static magnetic field distorts the field, depending on the magnetic susceptibility present in the object. This effect is field strength dependent (getting worse with higher field strength) and depends on the local magnetic susceptibility gradient. On a microscopic level, B_0 inhomogeneities result in the BOLD contrast (discussed in chapter 3), but on a macroscopic level they are mainly a source of image artefacts. The field differences cause signal dropout, distortions [35] and for some encoding schemes also blurring. The best way to reduce artefacts due to B_0 inhomogeneity is to reduce the inhomogeneity itself, by performing shimming. Shimming is a process, which uses the gradient coils as well as designated higher order shim coils to apply spatial magnetic fields to counter-act the inhomogeneity. However, very large and focal susceptibility mismatches, as seen for example at air tissue interfaces in the human head (in the sinus and ear cavities), are difficult to remove completely. Avoiding long echo trains, for example using parallel imaging, can also reduce distortion. Another approach is to measure the B_0 field and use this information to correct for the distortion in post-processing [36, 37].

2.5.3 Trajectory Imperfections

For a number of reasons the dynamic gradient fields frequently deviate from the desired ones. One of the main mechanisms generating these imperfections are eddy currents, generated in the conductive structures of the scanner (for example in the low-temperature heat shields around the main magnet) [38]. Mechanical oscillations of the gradient coils can lead to a current induction in the coils and subsequently oscillation in the generated fields [39]. Other hardware related mechanisms are timing mismatches, bandwidth limitations of the gradient amplifiers, as well as resistance and self-inductance in the coils. Thermal drifts can also change the state of the main magnet and thereby introduce a field drift in the main magnetic field. Despite great efforts in gradient design and correction schemes

(such as using gradient pre-emphasis) a certain amount of discrepancy to the nominal trajectory remains in practice [40, 41].

The appearance and the severity of artefacts, caused by trajectory imperfections, depend on the sampling scheme used [42]. Slow sampling schemes, such as the gradient echo sequence, operate less at the limits of the gradient system and are therefore generally more robust. In EPI the high gradient amplitudes and fast switching (and therefore fast slew rates) cause more eddy currents. Due to the alternating readout polarity used in EPI, delays cause shifts in a zigzag pattern [43], which results in ghosts shifted by half the FOV. This artefact is called the Nyquist ghost or N/2 ghost. Non-Cartesian trajectories frequently also have high hardware demands. In spiral trajectories, for example, trajectory imperfections typically lead to blurring.

3

Functional MRI – Background and Current Methods

Contents

3.1 The Basics of Functional MRI 31
3.1.1 Physiological Basis for fMRI 32
3.1.2 The BOLD Signal 33
3.2 BOLD fMRI at Ultra-High Fields 36
3.3 BOLD fMRI Acquisition Strategies 38
3.4 fMRI Experimental Design and Analysis 40
3.5 Physiological Noise and Motion 42
3.5.1 Physiological Noise 42
3.5.2 Motion in fMRI 45
3.6 Summary and Thesis Motivation 49

This chapter provides an introduction to functional MRI based on the Blood Oxygen Level Dependent (BOLD) contrast. It summarises the basic mechanisms behind the BOLD response, methods for BOLD fMRI acquisition as well as physiological noise and motion correction techniques for fMRI.

3.1 The Basics of Functional MRI

Brain activation typically refers to the increase in the electrophysiological activity of a population of neurons. In vitro studies and studies on animal models can use invasive patch clamp or single unit recordings to study the neurophysiological basics of information processing in the brain. However, for studying the human brain as a whole it is important to have non-invasive brain imaging techniques available. A large range of imaging methods have been developed [44], which aim to detect changes in neural activity. Some of the methods measure electrical activity and associated magnetic fields directly, like electro-encephalography (EEG) and magneto-encephalography (MEG). Most techniques, however, use indirect measures by measuring secondary physiological changes caused by the brain activity. These include metabolic changes, changes in cerebral blood volume (CBV), changes in the cerebral blood flow (CBF) or changes in blood oxygenation. Positron emission tomography (PET) measures brain metabolism, by looking at the decay activity of a previously injected metabolic tracer. Optical methods, such as near-infrared spectroscopy, can measure information on blood flow and oxygenation state. PET has the drawback that it is invasive, due to the injection of a radioactive tracer. Furthermore PET, EEG/MEG and optical methods are limited by poor spatial resolution. The need for a non-invasive flexible brain imaging method has led to the rise fMRI, which is currently the most commonly used technique to study human brain activity.

The family of MRI based methods aiming at measuring brain function are called functional MRI. Functional MRI makes use of coupling between neuronal activity and haemodynamics, to map brain activity [45]. It has become the favoured method for neuroscience investigations on healthy volunteers: it is non-invasive, MRI scanners are widely available and it provides a good balance of spatial and temporal resolution. Methods based on blood volume contrast, such as vascular space occupancy (VASO) [46], and methods measuring blood perfusion, such as arterial spin labelling (ASL) [47, 48], also fall under the umbrella of functional MRI. However, here we focus on fMRI using the blood oxygen level dependent (BOLD) contrast, which was used for fMRI in this thesis.

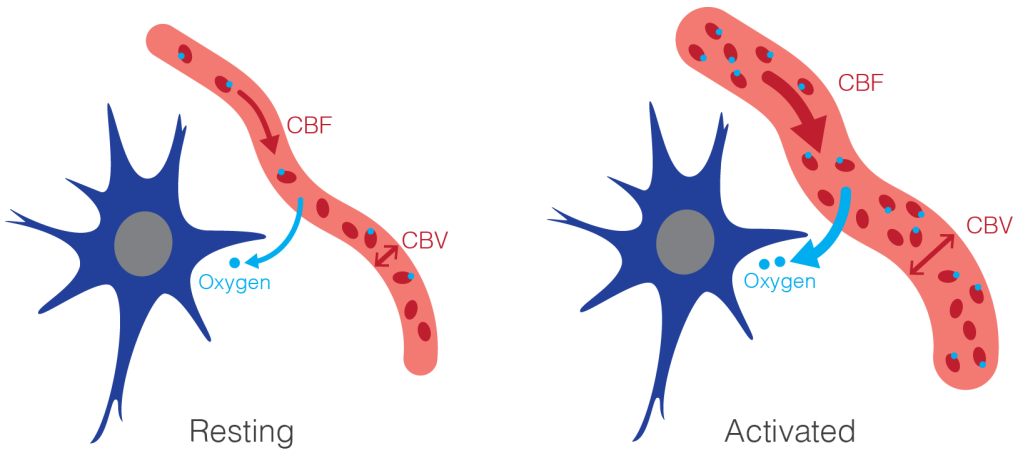


Figure 3.1: A simple illustration of the physiological changes at the basis of fMRI. The blood supplies oxygen (illustrated with blue circles) to the brain cells. In a region of active neurons, the $CMRO_2$ is higher. The increased oxygen demand is met by increasing CBV and CBF, resulting in larger amounts of oxygenated blood in the blood vessels near the active area.

3.1.1 Physiological Basis for fMRI

The brain produces its energy using a mechanism called the mitochondrial oxidative metabolism, which requires a constant and high influx of glucose and oxygen. Oxygen is extracted from oxyhaemoglobin (oxygenated form of haemoglobin), which is attached to the red blood cells. This exchange takes place in the capillaries close to the site of activation. Because we know that neurons require oxygen to function, we use the amount of oxygenated blood transported to a certain brain region as a proxy for brain activity. The relationship between the influx of blood and the neural activation is called the neurovascular coupling.

An increase in neuronal activation results in increased energy utilisation, therefore increasing the cerebral metabolic rate of oxygen consumption ($CMRO_2$). This triggers the release of signalling molecules, which cause vasodilation. The subsequently larger diameter of the blood vessel (increased CBV) reduces the vascular resistance and causes an increase in CBF. Higher CBV and CBF allow more arterial blood, which is highly oxygenated, to flow into the area (see Fig. 3.1).

The fact that neural activity gives rise to haemodynamic changes has been well established; for example through experiments showing correlation between the BOLD signal and local

field potentials [49]. The precise mechanisms that cause vasodilation and vasoconstriction are complex, potentially involve many different cell types and are an active area of research.

3.1.2 The BOLD Signal

There are two mechanisms causing the relaxation of the transverse magnetisation: spin-spin interactions (governed by the decay constant T_2) and changes in spin precession due to local magnetic field inhomogeneities (given by T_2'); the combined effect of the two is characterised by T_2^*:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{3.1}$$

While T_2' and T_2 decays are both driven by the loss of phase coherence, the T_2 decay is non-refocusable, as the phase changes are caused by stochastic effects, such as water diffusion. The T_2' decay, however, can be refocused using a spin-echo. The BOLD contrast is caused by the influence of blood oxygenation on both the T_2 and the T_2' decay.

Haemoglobin (Hb) is the protein that carries oxygen in the blood. It contains four haem groups each of which includes an iron atom. The iron atom contains four unpaired electrons, creating an orbital magnetic moment and having a paramagnetic effect on the surrounding water spins. If oxygen binds to haemoglobin the resulting Fe-O_2 complex is diamagnetic and therefore oxyhaemoglobin has very similar magnetic susceptibility as the surrounding tissue water, which is also diamagnetic. On the other hand deoxyhaemoglobin, which is unsaturated with oxygen molecules, is paramagnetic. It has been shown that the magnetic susceptibility of blood depends linearly on its oxygenation level [50].

The observation that changes in deoxyhaemoglobin concentration modulate the MRI contrast was first described in the in the 1980s and 90s [51–54]. A number of different mechanisms contribute to the BOLD contrast, they can be separated into extravascular and intravascular effects and their relative importance depends on the field strength and the acquisition strategy used.

Extravascular dephasing is caused by distortions in the local magnetic field around the blood vessels. The amount of field perturbation depends on the angle of the vessel to B_0, the magnetic susceptibility of the blood inside the vessel (i.e. the oxygenation level of the blood), as well as the size of the vessel (see Fig. 3.2). An important feature is that there is no extravascular field change or field distortion for a vessel oriented perpendicular to B_0. These dephasing effects shorten T_2^*, contributing to both T_2 and T_2' as described below.

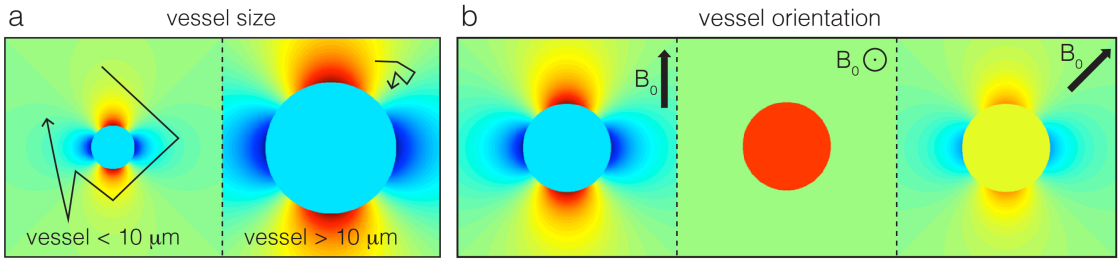


Figure 3.2: Magnetic dipole field around a vessel: (a) illustrates the mechanisms for extravascular dynamic (small vessel) and static (large vessel) dephasing, (b) shows the impact of the orientation of the vessel with respect to B_0.

For large vessels (diameter > 10 μm) the average path length of the water spins diffusing around in the inhomogeneous field surrounding the vessel is small with respect to the extent of the field distortions. This effect is called "extravascular static dephasing" and can be refocused using a spin-echo acquisition. For small vessels (diameter < 10 μm), the spins experience large field variations over a typical diffusion distance. Therefore the loss of phase coherence becomes stochastic and cannot be refocused; this is called "extravascular dynamic dephasing".

Intravascular effects also influence the BOLD signal. Firstly, the oxygenation level can directly influence the T_2 of the blood, because of the stronger interaction of the water spins with deoxygenated blood, resulting in shorter T_2. Secondly, there is vascular dephasing due to the frequency offset within the vessel compared to it's surrounding. How this effect influences the BOLD signal from a voxel depends critically on the orientation and distributions of vessels within. Generally speaking, the effect is a larger scale effect (similar to the static extravascular dephasing described above) and can be refocused [55, 56].

When the brain region is activated, somewhat counter intuitively, oxygenation increases despite increased demand. Therefore, the dHb concentration is reduced and the T_2^*

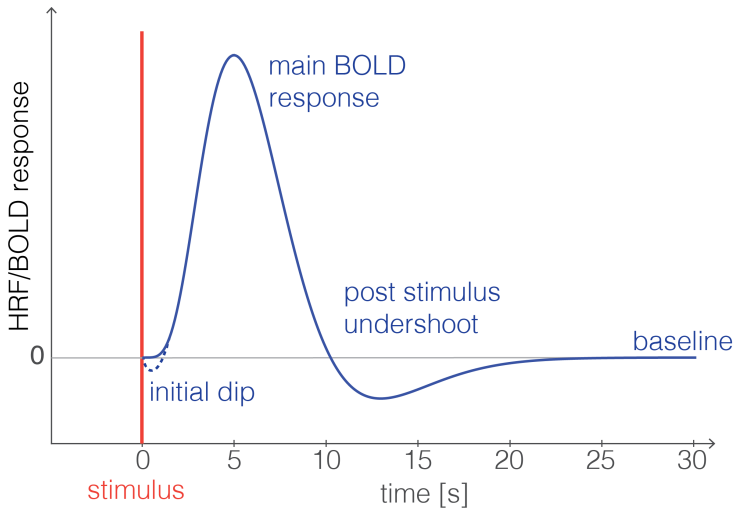


Figure 3.3: Illustration of the haemodynamic response function (or BOLD signal response to a brief stimulus), modelled by two gamma distributions. The controversial initial dip is indicated with a dashed line.

lengthened, leading to an increase of the MR signal in the corresponding regions. The temporal dynamics of the BOLD response are characterised by the so-called haemodynamic response function (HRF). An "initial dip" has been observed in some cases, due to increased oxygen uptake, before hemodynamic changes occur [57]. However, the blood flow effect quickly dominates and causes the positive main BOLD response to peak after approximately 4–5 s. The BOLD signal then drops lower than the baseline (post-stimulus undershoot), and slowly returns to baseline. A simple schematic of the HRF is shown in Fig. 3.3. The temporal characteristics of the BOLD response are influenced by the temporal dynamics of CBF, CBV and the $CMRO_2$. There are several biophysical models for explaining the BOLD response; the most well known is the so-called Balloon model [58].

The MR signal measured in a BOLD fMRI experiment depends on the signal S_0, available immediately after excitation, and the subsequent signal decay, governed by the decay rate $R_2^* = \frac{1}{T_2^*}$:

$$S = S_0 e^{-TE \cdot R_2^*} \tag{3.2}$$

The BOLD contrast is proportional to the change in R_2^*, it can therefore be expressed as the differential:

$$\frac{\mathrm{d}S}{\mathrm{d}R_2^*} = -TE \cdot S_0 e^{-TE \cdot R_2^*} = -TE \cdot S \tag{3.3}$$

In order to be able to detect the signal difference between activated and resting areas of the brain, the BOLD contrast resulting from very small changes in T_2^* has to be maximised. For a BOLD induced change in R_2^* of ΔR_2^* the signal change is

$$\Delta S = -\Delta R_2^* \cdot TE \cdot S. \tag{3.4}$$

Taking the differential with respect to TE and setting this to zero can find the optimal TE for maximal BOLD weighting.

$$\frac{\mathrm{d}(\Delta S)}{\mathrm{d}(TE)} = \frac{\mathrm{d}}{\mathrm{d}TE}(-\Delta R_2^* \cdot TE \cdot S_0 e^{-TE \cdot R_2^*}) = (TE \cdot R_2^* - 1) \cdot \Delta R_2^* \cdot S \tag{3.5}$$

This results in maximal BOLD sensitivity for $TE = T_2^*$. Typical values for T_2^* of grey matter are ∼60 ms at 1.5 T, ∼45 ms at 3 T and ∼25 ms at 7 T.

3.2 BOLD fMRI at Ultra-High Fields

The majority of existing clinical MR systems operate in a B_0 range of 0.5 T to 3 T. For human brain research applications 3 T is the standard, but 7 T MRI is gaining popularity [59] and there is a push to even higher field strength. A few 9.4 T scanners are already currently in use [60], and systems up to 11.7 T are currently being developed [61, 62]. For historic reasons 3 T MRI is called high-field MRI, while everything beyond is currently referred to as ultra-high field MRI.

As described in chapter 2, the magnetisation available for NMR signal formation is proportional to the static field strength. Therefore a 7 T scanner is expected to provide just over twice the thermal SNR compared to a 3 T scanner. Functional MRI specifically profits from higher field strength in more ways than just by an increased SNR. The T_2^* weighted signal scales with B_0, leading to larger BOLD signal changes. Theoretically a linear dependence is expected for larger veins and supra-linear dependence for capillaries [63, 64]. Larger BOLD responses at higher fields have been demonstrated in vivo [60, 61].

One of the challenges in BOLD fMRI is that a large proportion of the signal can come from large veins [65]. These draining vessels are potentially far away from the site of

activation, causing a loss in spatial specificity. The T_2^* of venous bloods drops rapidly with field strength and is very short (< 5 ms) at 7 T, This means that, at typical echo times, the intravascular contribution from venous blood becomes negligible and the signal comes almost exclusively from the capillaries, which are very close to the site of activation. Frequently, the activation maps therefore look sparser at 7 T, showing a reduced number of active voxels, but gaining spatial specificity. Because the extravascular dephasing effect of large vessels can be refocused, the specificity can be further improved by using a spin-echo acquisition [66].

The higher field strength comes with a number of challenges. For example geometrical distortions in EPI images get considerably worse because of the larger amount of field inhomogeneity; especially for high spatial resolution scans with long echo train lengths. Furthermore, physiological noise is generally assumed to be proportional to to the detected MR signal[67], which scales with B_0^2 [68]. This can eliminate some of the benefits of ultra-high field fMRI. Physiological noise is discussed in section 3.5.1 and in more detail in chapter 4. In short, appropriate imaging parameters and correction schemes must be chosen to achieve the expected BOLD contrast gains. This normally requires imaging at high resolution ($\leq$ 1.5 mm isotropic).

There are several reasons why high field strength and high resolution are a good match: Higher fields enable higher resolution, because the SNR increase partially offsets the large SNR penalty of smaller voxels. As pointed out above, the BOLD response at 7 T has a narrower point response [69] encouraging the use of smaller voxels. Large voxels are likely to contain larger veins as well as contribution from the microvasculature. High resolution reduces this "partial volume effect" and thus allows to better identify the contribution that stems solely from the capillary bed. Finally, the use of small voxels ensures that imaging is performed in a thermal noise dominated regime.

Overall, if the challenges of high fields are adequately addressed, high field MRI offers increased sensitivity and specificity for BOLD fMRI.

3.3 BOLD fMRI Acquisition Strategies

The first BOLD experiments were conducted using a spin warp sequence [52]. While this yields T_2^* weighted BOLD contrast, it quickly becomes limiting because of the slow acquisition speed. Therefore, shortly after the discovery of the BOLD signal, EPI was introduced as the method of choice for fMRI [70, 71]. To date conventional 2D multi-slice EPI is the most frequently used acquisition strategy for BOLD fMRI.

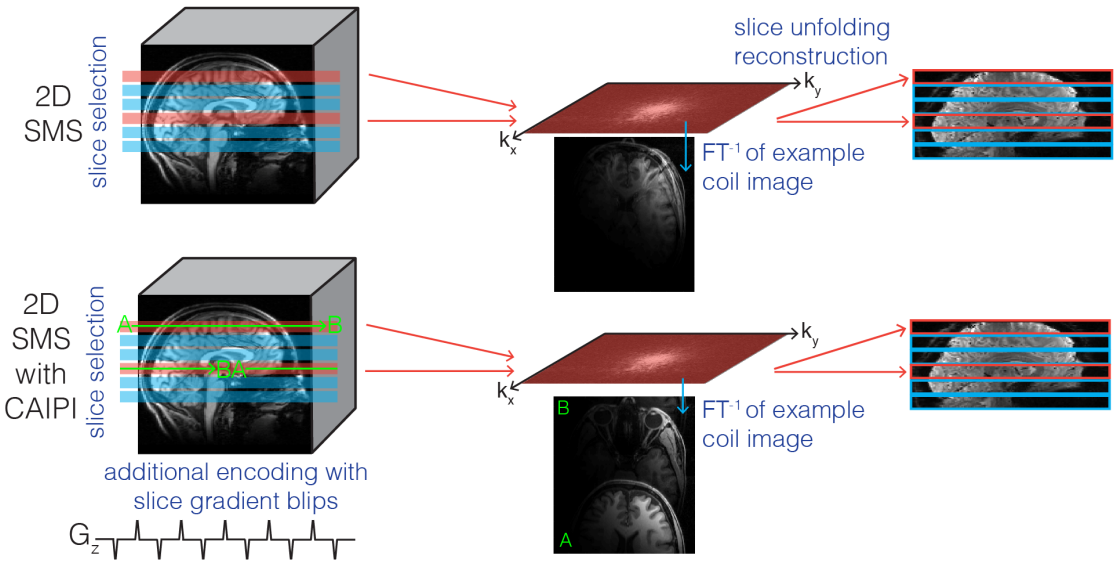


Figure 3.4: Illustration of simultaneous multi-slice acquisition, where multiple slices are excited at once and separated with the aide of multiple receive coils. If CAIPI blips are used (bottom panel) the images from the different slices can be unfolded more easily. For illustration of 2D and 3D EPI see Fig. 2.2.

2D multi-slice spiral imaging has long been considered the main competitor for EPI [72, 73]. It can be nominally faster than EPI, when the acquisition is limited by hardware constraints, because the gradient load is distributed more evenly over the acquisition window. However, on modern systems the gradients are often sufficiently powerful and limitations are more likely set by peripheral nerve stimulation (PNS) limits. This, and the fact that spiral imaging requires more sophisticated reconstruction and artefact correction schemes, have caused spirals to not yet become a mainstream technique.

The imaging requirements for BOLD fMRI have evolved over recent years. Originally, due to the inherently broad spatial point spread function of the BOLD response, in-

plane spatial resolutions of 3–5 mm were considered sufficient. Similarly, it was thought that there was no benefit in imaging with volume TRs faster than a few seconds, due to the slow haemodynamic response. This traditional combination of temporal and spatial resolution can easily be achieved with 2D multi-slice EPI. However, the advent of resting state fMRI and high field fMRI has created new incentives for both high temporal and high spatial resolution.

Feinberg et al. [74] showed that faster sampling increases the sensitivity in detecting brain networks in resting-state fMRI data. High temporal resolution allows detecting more networks [75] and for very high temporal resolution, it can reveal dynamic changes across networks [76]. Imaging at high spatial resolution reveals fine anatomical structures, enabling to examine the columnar [77] and laminar architecture [78] of the cerebral cortex. Recent studies (using optical imaging/microscopy [79]) have also indicated that the spatial scales of the haemodynamic response might be smaller than originally assumed.

High field fMRI also creates a technical incentive for efficient k-space readouts, to mitigate the increased distortion and signal loss produced by increased magnetic field non-uniformity as well as reduced T_2^* values. The main limitation of 2D multi-slice EPI is that the acquisition can only be accelerated in one dimension. The efforts to address this limitation have resulted in the development of two methods, which have the potential to dramatically speed up fMRI acquisitions: 2D simultaneous multi slice imaging (SMS) [80] and 3D segmented EPI [81].

In SMS, also called multiband imaging, multiple slices are excited at the same time. The recorded k-space is associated with multiple slices and the spatially varying coil sensitivities are used to separate the images post acquisition [80], see Fig. 3.4. If the regions of overlap, between images from different slices, is minimised, there is very little SNR penalty from SMS acceleration (as there is no omission of k-space lines and therefore no reduction factor penalty). The overlap can be reduced using the Controlled Aliasing In Parallel Imaging Results IN Higher Acceleration (CAIPIRINHA) method [82, 83], in

short called "CAIPI". The goal of the scheme is to increase the distance between aliased voxels in image-space, by manipulating the sampling in k-space.

3D segmented EPI [81] (frequently just called 3D EPI) provides high SNR and allows flexible under-sampling. The properties and potential benefits of 3D EPI are discussed in much more detail in chapter 4. CAIPI style sampling has also been recently combined with 3D EPI (sometimes called 3D CAIPI EPI) [84, 85]. Again, by manipulating the sampling in k-space, CAIPI reduces the burden on the coils for reconstruction, allowing higher acceleration factors.

The methods mentioned above are not an exclusive list and a multitude of other approaches have been suggested, for example the 2D methods 'radial EPI' [86] and PROPELLER [87]. Apart from 3D segmented EPI, many other 3D trajectories have also been tested for fMRI: In echo voluminar imaging (EVI) [88, 89], the entire 3D k-space is sampled in one shot. Other trajectories include 3D radial EPI [90], 3D stack of spirals [91, 92], 3D rosettes [93] and 3D concentric shells [94]. TURBINE [95], the 3D trajectory we have been developing for fMRI, is introduced in chapter 5.

In most neuroscience laboratories 2D multi-slice EPI is still considered the workhorse for fMRI studies, but both SMS and 3D-EPI have been successfully applied to provide high-quality fMRI data, and are rapidly gaining popularity. SMS is being used in a number of large studies, such as the Human Connectome Project [96]. Another large population study, the Rhineland study [97], has chosen to use 3D EPI [98]. The field is rapidly evolving and it remains to be seen if a method will emerge as the new gold standard.

3.4 fMRI Experimental Design and Analysis

Task fMRI studies the brain's response to specific tasks and/or stimuli. In contrast, resting state fMRI looks at how different brain regions communicate while the brain is "at rest" (i.e. not performing any active task). While the same biophysical contrast

mechanisms govern both cases, their experimental conditions and analysis methods are quite different, and we focus this discussion on task-based fMRI.

The contrast to noise ratio in fMRI is driven by the difference in BOLD signal in different brain states. In the simplest so-called block design, the stimulus or task is presented continuously over the active block, followed by a block of rest. Optimal block durations are between 20-30 seconds. The downside of block designs is that no information about individual trials or events can be recovered. If information about individual events is required, the stimuli need to be spread out in an event related design.

The 4D fMRI data set is pre-processed in preparation for the statistical analysis. Typical steps include the following (in no particular order):

- Head motion correction and physiological noise removal, as described in more detail in section 3.5
- Brain extraction, consisting of the masking of non-brain voxels
- Slice-timing correction in 2D imaging
- Registration to a structural template brain for multi-subject analysis
- Temporal (high and/or low-pass) filtering
- Correction for temporal autocorrelation (sometimes called "pre-whitening")
- Spatial smoothing, performed to boost the SNR (this is avoided in high-resolution fMRI)

Most commonly a General Linear Model (GLM) approach is used to detect BOLD activation. GLMs are univariate, i.e. they do not exploit covariance across neighbouring voxels to infer brain function. Multivariate methods, such as multi-voxel pattern analysis [99] and independent component analysis (ICA) [100], are becoming increasingly important but are not used or discussed in this thesis.

The GLM uses a model (i.e. a set of patterns you expect to see in the data) and fits it to the data. Typically, the model consists of the stimulus paradigms convolved with the haemodynamic response function, as well as potential nuisance regressors (for example

taking into account motion or physiological noise). Expressed as an equation, model $\boldsymbol{X}$ is fitted to the data $\boldsymbol{Y}$ to determine the parameter estimates (PEs) $\boldsymbol{\beta}$ and the error term $\boldsymbol{\epsilon}$:

$$\boldsymbol{Y}[t] = \boldsymbol{\beta} \cdot \boldsymbol{X}[t] + \boldsymbol{\epsilon} \tag{3.6}$$

The PEs are determined by minimizing the squared differences between $\boldsymbol{Y}$ and the estimated data $\boldsymbol{\beta} \cdot \boldsymbol{X}[t]$, also yielding $\boldsymbol{\epsilon}$. To test whether the fit is significant, the PE is divided by the standard error of the estimate. This results in a t-value, which can be converted to z-values (using statistical transformations) and displayed as z-maps. These statistical maps are normally thresholded according to a chosen level of significance, to show which voxels are active under these assumptions. As a typical fMRI data set contains many voxels, it is important to be account for the multiple comparison problem, stating that a certain amount of voxels will be "activated" by chance, even in the absence of a stimulus. This can be corrected for by adjusting the statistics for each voxel to account for the number of voxels being considered (potentially in a data-driven way, such as with permutation testing).

3.5 Physiological Noise and Motion

Additional to the typical image artefacts (caused by the hardware or the acquisition strategy), which are described in chapter 2, MRI data also suffers from noise and artefacts introduced by the subject. These subject-related challenges fall into two categories, physiological noise (3.5.1) and motion (3.5.2). Note that the topic of physiological noise is covered more in-depth in chapter 4 and motion correction is further explored in chapter 6 of this thesis.

3.5.1 Physiological Noise

Physiological noise is structured temporal variation introduced by the subject, which is typically temporally correlated to the subject's cardiac and respiratory cycle. Physiological noise encompasses a large range of effects: It includes changes in the flow of blood and

cerebrospinal fluid (CSF), driven by cardiac pulsation as well as subtle brain pulsation [101]. Respiration causes changes in lung volume, which alter the susceptibility and therefore the B_0 field; this can cause phase distortion [102, 103]. Furthermore, respiration can introduce subtle motion artefacts. Lastly, changes in respiratory and cardiac rates can give rise to fluctuations in the blood oxygenation levels [104].

These changes in the fMRI signal have a complex spatial and temporal structure and they can be very difficult to separate from the BOLD response [105]. The problem is even more complicated, as physiological variations can be coupled with increased task demands. In that case, removing the physiological fluctuations often removes a part of the signal of interest as well. Physiological noise limits the temporal signal-to-noise ration (tSNR), can substantially hamper the ability to detect activation and can introduce spurious resting-state functional connectivity (see [106] for a review). Here, we briefly introduce a few methods available for physiological noise reduction.

Respiratory and cardiac noise signals are generally very periodic, with a frequency of ∼1 Hz (∼60 beats per minute) for the cardiac fluctuations and ∼0.2-0.3 Hz (∼3-5 s per breath) for respiratory fluctuations in typical subjects. Therefore, a simple and fast correction is band-stop filtering, by removing the frequency bands of the cardiac and respiration signals [107]. This approach only works when the acquisition is fast enough to Nyquist sample these signals. Typical fMRI acquisitions sample much slower, resulting in the fluctuations being aliased into the data and the signal being spread throughout the frequency range. It also assumes that the signal is stationary (i.e. ignoring the signal phase), which is not valid for cardiac activity or respiration changes.

The most commonly used methods require external measurement of respiration and heartbeat and many of them are based on the assumption that physiological fluctuations are quasi-periodic. In the method RETROICOR [108], low-order Fourier terms are fitted to the data based on the time of each image acquisition, relative to the phase of the cardiac and respiratory cycles. These regressors are then included in the GLM, to remove physiological noise, and thereby improving statistical significance. This approach

was initially proposed as a correction in k-space, RETROKCOR [109], but the image space version RETROICOR is much more frequently used today. A number of variants of RETROICOR have been introduced, including for example optimised regressors for different brain region [105, 110]. RETROICOR and related methods have been shown to substantially improve the BOLD sensitivity, especially at 7 T [111].

RETROICOR has been used to correct 3D data [112, 113], but it is limiting because the regressors are evaluated at one time point per volume; assuming that each image is acquired instantaneously. For 3D data with typical volume TRs of several seconds, this assumption is likely not valid. Therefore intuitively we would expect RETROKCOR to improve correction in 3D data. However, a recent comparison on 3D EPI data [114] concluded that RETROKCOR does not provide any improvement over RETROICOR, so the impact of physiological noise on 3D data remains quite significant.

A limitation of RETROICOR is that it cannot cope well with sudden changes in physiological signal, and that it is based on relatively strict model assumptions (such as a linear relation between cardiac/respiratory phase and the induced signal change). A range of other methods exist, which try to address some of these issues; for example the Bayesian DRIFTER [115] technique or data driven approaches (requiring no external measurements of heartbeat and respiration) such as spatial ICA [116].

Another strategy is to incorporate navigator echoes into the acquisition to track and subsequently correct for the physiological fluctuation [117]. Trajectories crossing the centre of k-space, such as EPI, allow for intrinsic frequency navigation. This can be used to reduce the impact of fluctuations causing global off-resonance shift, like respiration, either prospectively or retrospectively.[118, 119].

Prospective correction using real-time adaptation of the acquisition strategy based on monitored respiratory/cardiac phase has also shown promising results [120]. Finally, multi-echo EPI has been suggested in order to separate potentially relevant "BOLD-like" physiological noise components from "non-BOLD like" components [121].

In addition to correcting for physiological noise, it is useful to try to minimise its contribution by choosing an acquisition strategy, which is less sensitive to it. How physiological noise impacts the temporal SNR of fMRI data is discussed in chapter 4.

3.5.2 Motion in fMRI

Due to the inherently slow sampling process in MRI, subject motion is one of the biggest problems across all MRI modalities. Functional MRI is especially sensitive to motion because it is a dynamic technique, aiming to detect small ($< 5\%$) signal changes. Even a small amount of motion can generate signal changes of larger magnitude than the signal difference we are aiming to detect and can be detrimental to the fMRI experiment [122]. Head motion that is correlated to the task (stimulus-correlated motion) can cause the signal of a brighter voxel to bleed into a darker voxel and vice versa, causing fluctuations which can easily be wrongly interpreted as functional contrast [123]; this is especially pronounced at contrast boundaries. On the other hand, uncorrelated head motion increases the variance in a time series thereby reducing the likelihood of true activation being identified as such.

Types and Effect of Motion

In fMRI it is mostly assumed that head motion is a rigid body process. Non-rigid brain motion occurs as well (mainly pulsatile motion in central and lower brain regions [124]) but for most application it is a reasonable assumption that the motion can be described by rigid body transformation (rotations and translations). This motion can occur between the acquisition of two 3D images (inter-volume) or within the volume (intra-volume). In this thesis we use intra-volume to refer to motion within the acquisition of one k-space (2D or 3D) and not as motion between 2D slices as used in most other publications.

When sampling k-space it is crucial that the object has remained stationary during the acquisition, otherwise k-space becomes inconsistent, resulting in artefacts (see Fig. 3.5). The impact of motion on the image quality depends on the acquisition strategy used and when in the acquisition the motion occurs. Inter-volume motion typically generates

an accurate image of the brain, but moved with respect to some reference volume. To interpret how intra-volume impacts the data, it is important to understand how motion affects k-space: When the head is rotated away from its reference position, the acquired lines in k-space are correspondingly rotated from their nominal orientation. Translations of the head away from the reference position correspond to a phase ramp in k-space. Intra-volume motion can therefore cause phase inconsistencies as well as potentially large un-acquired gaps in k-space. The way this manifests itself in the image is trajectory dependent; it can include blurring, ghosting and drop out artefacts. The likelihood of severe intra-volume motion artefacts increases the longer the sampling period. For this reason 3D sampling strategies are more motion sensitive than 2D ones.

Another problem caused by motion are spin history artefacts. This occurrs for example when saturated spins from slices that were previously imaged (and have not fully recovered yet) move into the space of the current slice prescription thereby reducing the intensity of the acquired images. This is especially problematic in 2D imaging with short TRs and SMS imaging. Spin history artefacts can be addressed for example with the ICA-based approach FIX [125, 126], which uses a classifier to identify "bad" motion-corrupted components and remove them from the data set.

As mentioned in chapter 2, B_0 field inhomogeneity can be compensated by shimming the volume of interest. The shim settings are normally determined in the beginning of an acquisition and then left constant throughout the scan. However, after motion the shim fields no longer adequately compensate for the field inhomogeneity. For example, nodding motion (which is a very common type of subject motion [127]) is expected to have a large impact on the B_0 field as this moves large susceptibility borders relative to B_0. A prospective solution may be dynamic shimming during imaging (for example based on information from field probes [128]).

The transmit B_1 field depends on the position of the head within the coil. Therefore the flip angle and consequently the MR signal can vary with motion due to the changing interaction of B_1 with the head. Motion also causes problems on the receive side,

especially when using multi-channel coils. Due to motion, the coil sensitivities can change relative to the head, which can generate artefacts in the parallel imaging reconstruction [129]. These artefacts are worse the stronger the coil sensitivities vary in space and the higher the under-sampling factors.

Motion Avoidance and Artefact Reduction

Preventing motion by immobilising subjects with physical restraints, for example using foam padding or a bite bar [77], is an effective way to minimise motion artefacts. It can however reduce subject comfort and be difficult to implement for patients or sensitive subjects like children.

Because intra-volume motion is especially detrimental, an obvious method to reduce the effect of motion is to "freeze" it by imaging faster. Furthermore, using trajectories, which are more "motion-tolerant" than Cartesian sequences, can reduce artefacts. These strategies (for example radial and spiral trajectories) typically oversample the centre of k-space, creating motion artefacts that have a more benign appearance in image space [130]. Some of these sequences are also suitable for advanced motion correction approaches based on self-navigation.

Motion Correction

A huge number of motion correction strategies have been developed, which cannot be completely covered here (see [131, 132] for reviews on the topic). In the majority of neuroscience fMRI experiments motion correction is performed as part of the analysis pre-processing pipeline. These image-based motion correction algorithms typically use a rigid body transformation to register motion-corrupted volumes to a reference volume [133]. The transformations are found by minimising image intensity based cost functions. These motion correction approaches perform re-alignment of image volumes and slices, but cannot correct any corruption due to intra-volume motion. Furthermore, the accuracy of the registration can be compromised in presence of large image artefacts. Another simple

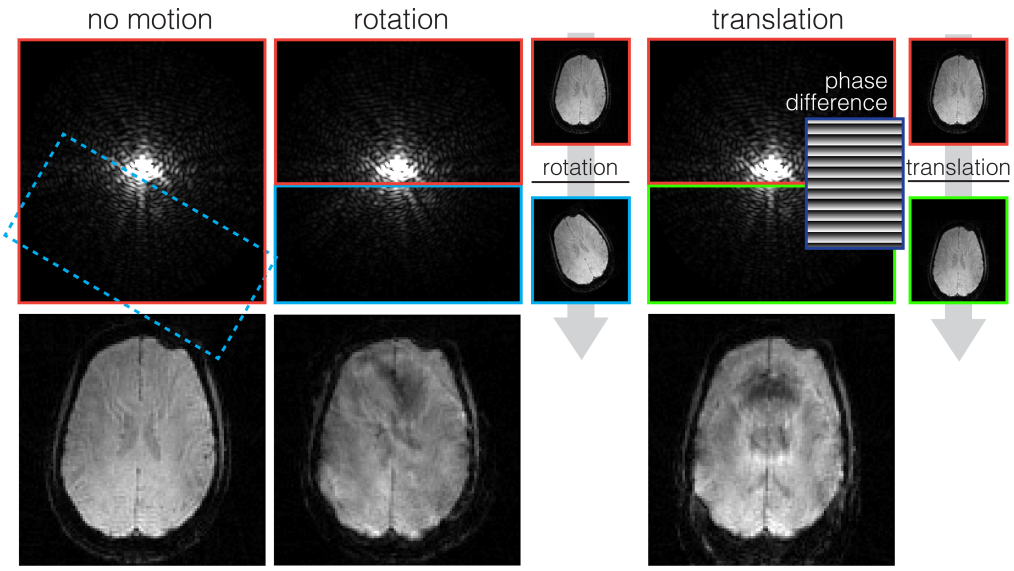


Figure 3.5: Illustration of the effect of intra-volume motion. The left panel shows a motion-free k-space and image. In the centre panel a rotation occurs during the acquisition of k-space causing a part of k-space to be rotated relative to the rest, resulting in image artefacts. In the case of a translation (right panel) a phase ramp is introduced, again causing image artefacts.

way to deal with motion is to discard corrupted volumes, called scrubbing [134]. However, this results in a loss of statistical power. In structural brain imaging selective re-acquisition [135] of volumes where motion occurred has been successful. However, rescanning of data is not used in fMRI, because the measured brain function is specific to every time point.

Because the registration methods described above cannot reduce image artefacts, motion correction approaches beyond a simple image registration step have been developed. These strategies can be classified based on when the correction is done: retrospective or prospective and how the motion is estimated: (a) using external sensors (motion camera), (b) explicit navigation (acquiring additional data specifically to track motion), (c) self-navigation (acquisitions which naturally include navigators) and (d) data driven approaches (e.g autofocusing reconstruction).

In prospective motion correction [136] the subject movement is monitored, using either a motion camera [137] or navigator information [138], and this data is used to dynamically update the imaging volume. The information from motion cameras can also be used to perform retrospective motion correction.

Many different navigator strategies exist and it has for example been suggested that Spherical Navigator Echoes could be applied to fMRI [139]. Navigator correction for fMRI [140] has been done by introducing separate navigators into a 3D PRESTO sequence. However, external navigation is not common in fMRI because it can take a prohibitively long time to acquire a full 3D navigator along with each k-space segment.

Therefore self-navigated sequences are attractive for fMRI, as they do not require the acquisition of any additional data. Examples of such approaches used for fMRI include PROPELLER fMRI [87] and self-navigated spiral fMRI [73]. These methods use the frequently sampled centre of k-space to detect and remove phase corruption.

Motion correction can also be based on intrinsic image properties. For example in autofocus strategies the motion estimation is based on minimizing a suitable image metric, while the raw data are modified according to a motion model [141].

3.6 Summary and Thesis Motivation

Summarising, motion and physiological noise are still a major limitation in fMRI. In this thesis we look at the impact of these fluctuations on 3D acquisition methods, which we identified as one of the most promising future directions for high-resolution fMRI. Specifically we look at the impact of physiological noise in 3D EPI (chapter 4) and present a novel 3D acquisition strategy to mitigate physiological noise and motion (chapters 5 and 6).

4

Temporal SNR Behaviour of 3D EPI at 7 T

Contents

4.1 Introduction 51
4.2 Theory 52
4.2.1 Influence of Slice Acceleration on tSNR in 3D EPI 56
4.3 Methods 58
4.3.1 Sequence Development 58
4.3.2 Sequence Testing and Comparison to 2D EPI 59
4.3.3 Main Experiment: Impact of Acceleration on tSNR 61
4.3.4 Image Reconstruction and Analysis 61
4.4 Results 62
4.4.1 Sequence Testing and Comparison to 2D EPI 62
4.4.2 Main Experiment: Impact of Acceleration on tSNR 65
4.5 Discussion 67
4.6 Conclusion 72
4.A Appendix: Literature Comparison and tSNR Simulation . . 72
4.B Appendix: Investigation of the Temporal Stability of 3D EPI (Phantom Experiments) 76
4.B.1 2D vs. 3D EPI Phantom Comparison 76
4.B.2 λ Measurements 78
4.B.3 Cross-site Comparison 79
4.B.4 Discussion 82

4.1 Introduction

To date the majority of fMRI studies are performed using 2D multi-slice EPI. Recently the use of 3D imaging for fMRI has seen increased interest [81, 112]. One of the main motivations for 3D imaging is a theoretical thermal SNR gain over the 2D imaging. A further advantage is the ability to accelerate (using parallel imaging or Partial Fourier) in more than one direction, allowing substantially faster volume acquisition times. In contrast to 2D acquisitions, the resolution in 3D imaging is not limited by imperfections of the excitation profile. 3D fMRI therefore provides advantages over 2D for ultra-high isotropic resolutions (≤ 1 mm^3), where hardware and SAR limits make it challenging to excite a well-defined sub-mm slice profile. The Ernst angle in 3D imaging is much lower, reducing the specific absorption rate (SAR) of the acquisition. This is especially important for high-field applications, where SAR restrictions can become limiting. 3D sequences are also typically less sensitive to B_1 inhomogeneity due to the saturation effect of short TRs. Another advantage of 3D imaging is the reduced sensitivity to spin history effects due to head motion as well as reduced sensitivity to inflow effects. The latter frequently occur in 2D imaging when unsaturated spins in CSF or vessels flow into a saturated slice or vice-versa.

While it would be desirable to image the entire volume in one shot (EVI), sampling the entire 3D k-space in one shot is sensitive to artifacts, such as distortion, Nyquist ghosting and blurring [88]. Therefore segmented 3D EPI is typically used, where a single EPI plane is read out in every shot and phase encoding is used along the slice direction. Being a multi-shot technique makes segmented 3D EPI inherently sensitive to any temporal fluctuations over the acquisition window, which is typically on the order of a few seconds for a high-resolution scan. The most relevant disadvantage of 3D EPI over 2D EPI is therefore its increased susceptibility to temporal fluctuations, such as physiological noise and motion.

Higher field strength offers new possibilities for studying human brain function, as the sensitivity and specificity of BOLD fMRI is increased relative to lower field strengths [142]. Importantly, higher fields allow the use of smaller voxels, by mitigating the SNR penalty

normally associated with high-resolution imaging. Functional imaging at high spatial resolution is relevant as the cerebral cortex has aspects of its fundamental organization at the sub-millimetre level (cortical lamina and columns). Smaller voxels have the additional benefit of reduced partial volume effect and reduced signal drop-out due to magnetic field inhomogeneities. However, with the signal increase in high field imaging, physiological noise also increases, which can limit the detection sensitivity in fMRI. The additional SNR boost from 3D imaging, the low SAR and the potential for highly accelerated trajectories[85] makes 3D EPI especially suited for high-field applications.

With the siting of Oxford's 7 T scanner, there was interest from neuroscience investigators to explore the benefits of 3D EPI for fMRI, for example for ocular dominance column imaging in the primary visual cortex. After implementing and testing this sequence, our main goal was to investigate the different trade-offs when optimising a 3D EPI acquisition for fMRI, namely the interaction between physiological noise and acceleration.

4.2 Theory

In this section, we provide a theoretical comparison of 2D and 3D GRE sequences in the regime typically used for fMRI. SNR is frequently used to evaluate MRI images and compare competing methods. It is defined as the quotient of the mean signal intensity measured in a region of interest and the standard deviation of a noise measurement (using either a background region ROI containing no signal or a separate zero flip angle scan). Since the main goal of fMRI studies is to detect small signal changes over time, spurious fluctuations over time such as physiological noise need to be taken into account for an accurate measure of the SNR of a time series. Therefore temporal SNR, calculated as the mean signal over time divided by the standard deviation over time, is the most important metric for sensitivity of an fMRI acquisition. SNR and tSNR are equal when only thermal noise is present in the system. In practice this is never the case due to physiological noise and system instabilities.

In many conditions, 3D imaging offers an intrinsic SNR gain over the 2D imaging, because the entire volume is excited at each TR. If N_{k_z} k_x–k_y planes are acquired, N_{k_z} more data points contribute to each point in the 3D Fourier reconstruction, therefore thermal noise scales with $\frac{1}{\sqrt{N_{k_z}}}$, increasing SNR. However, these gains are mitigated by the different signal levels in typical 2D vs 3D imaging protocols. The time between successive excitations in 3D imaging is short compared to the T_1 of the imaged tissues. Therefore the magnetisation does not fully recover and the steady state magnetisation available for signal formation is much lower than typically for the 2D case. For a spoiled gradient echo pulse sequence the signal is proportional to the steady state longitudinal magnetisation:

$$S \propto M_0 \cdot \sin\theta \cdot \frac{1 - e^{-TR/T_1}}{1 - \cos\theta \cdot e^{-TR/T_1}} \cdot e^{-TE/T_2^*} \tag{4.1}$$

$$\theta_{\text{Ernst}} = \cos^{-1}(e^{-TR/T_1}) \tag{4.2}$$

Assuming that the TE is chosen for optimal BOLD contrast, and is therefore the same for 2D and 3D acquisitions, T_2^* decay can be ignored. Evaluating equation 4.1 at the Ernst angle given by equation 4.2 and taking into account the $\sqrt{N_{k_z}}$ noise reduction, we can calculate the theoretical SNR gain expected for 3D imaging. In Fig. 4.1 the resulting SNR is shown for a range of different N_{k_z} for a $T_1 = 1400$ ms (grey matter at 7T) and TR = 100 ms. From Fig. 4.1 we see that the $\sqrt{N_{k_z}}$ noise reduction outweighs the reduced steady state magnetisation available for MR signal formation. Therefore in the absence of physiological noise and motion, we expect the 3D sequence to always outperform the 2D sequence and the SNR benefit is expected to increase as we increase the number of slices. On a side note, 3D EPI has different contrast to 2D EPI because of the rapid re-excitation, which alters the T_1 weighting.

The main disadvantage of 3D imaging is its higher susceptibility to physiological noise due to the longer time window over which the k-space is acquired. Physiological noise is frequently described using the Krüger model [67], where the noise sources are split into two components, thermal and system noise σ_0 and physiological noise σ_p. The intrinsic thermal and system noise variance σ_0 includes thermal noise from the subject

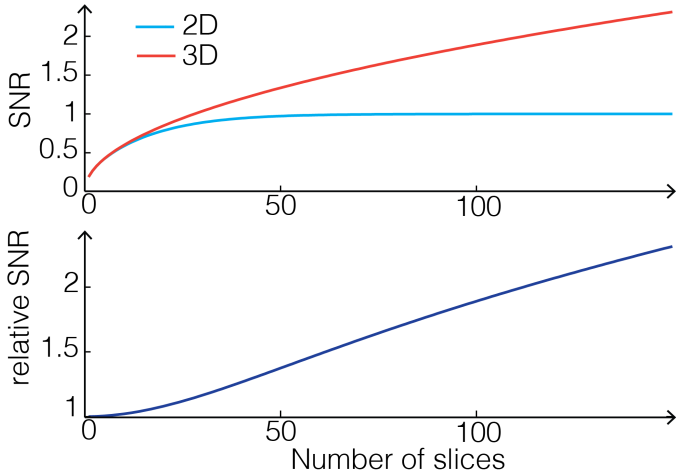


Figure 4.1: Theoretical SNR for 2D and 3D EPI as a function of the number of acquired slices. T_1 = 1400 ms and TR(3D) = 100 ms was assumed in this case. The volume TR increases with the number of acquired slices and is the same for both methods. Plot reproduced as in [81].

and scanner electronics and system noise, for example scanner drift and gradient, RF or shim instabilities.

The physiological noise σ_p stems from subject-related signal fluctuations. This includes variance arising from neural and vascular fluctuations (changes in the basal cerebral metabolism, blood flow, blood volume) as well as fluctuations due to cardiac pulsation and respiration (causing fluctuations in the vascular system, motion from subtle brain pulsatility and magnetic field modulations). The Krüger model assumes that physiological noise scales with signal intensity, while thermal noise does not. The physiological noise σ_p is described as two separate noise components, a "BOLD-like" term $\sigma_B = c_1 \cdot S \cdot \Delta R_2^* \cdot TE$ and a "non-BOLD-like" term $\sigma_{NB} = c_2 \cdot S$. The BOLD-like noise is proportional to the TE and fluctuations in the transverse relaxation rate. The constants c_1 and c_2 govern the strength of the BOLD-like and non-BOLD-like terms respectively.

The physiological and thermal noise sources are statistically independent and therefore the total raw noise is $\sigma_{tot} = \sqrt{\sigma_0^2 + \sigma_p^2}$ and the tSNR be expressed as:

$$tSNR = \frac{S}{\sqrt{\sigma_0^2 + \sigma_p^2}} \tag{4.3}$$

As both physiological noise terms are proportional to S we can express the total physiological noise as $\sigma_p = \lambda \cdot S$, where $\lambda^2 = c_1^2 \cdot \Delta R_2^{*2} \cdot TE^2 + c_2^2$. Using the definition of

thermal SNR, $SNR_0 = \frac{S}{\sigma_0}$, the tSNR in the Krüger model can then be expressed as:

$$tSNR = \frac{S}{\sqrt{\sigma_0^2 + \lambda^2 S^2}} = \frac{SNR_0}{\sqrt{1 + \lambda^2 SNR_0^2}} \tag{4.4}$$

For $\lambda = 0$, the tSNR is equal to SNR_0. Equation 4.4 also allows us to estimate a theoretical limit of tSNR for infinite thermal SNR_0, which is given by $\frac{1}{\lambda}$ (see 4.2). The scaling factor λ can therefore be seen as the measure of the sensitivity of an acquisition to physiological noise. λ is both influenced by the amount of physiological noise and by the transformation the imaging process has on the physiological noise.

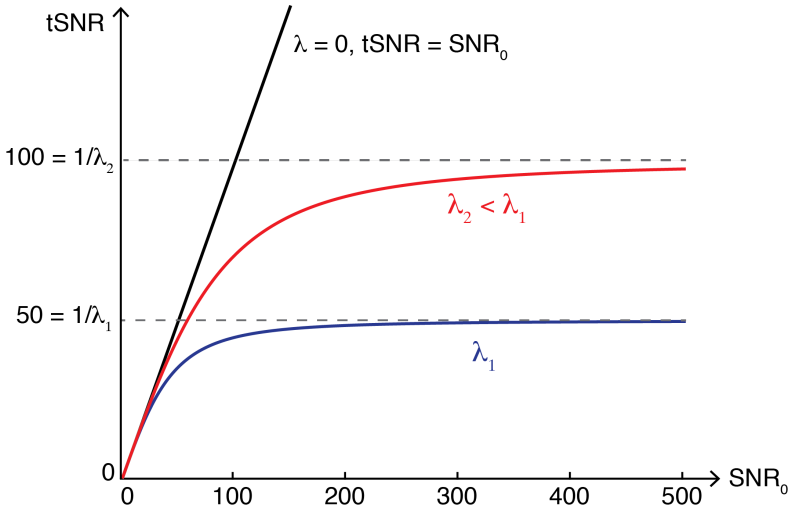


Figure 4.2: Plot to illustrate the relationship between tSNR and SNR as described by the Krüger model. Identity line (no physiological noise) in black and tSNR plots for higher (blue) and lower (red) amounts of physiological noise. As the SNR_0 increases, the tSNR approaches the asymptotic limit for tSNR given by of $\frac{1}{\lambda}$.

We noted above that scanner instabilities contribute to σ_0. More accurately the σ_p in the Krüger model simply lumps together all signal instabilities that scale with S, under the assumption that these are predominantly "physiological". Therefore, σ_0 and σ_p contain some sources of variance from scanner instabilities and determining λ on a phantom gives a measure of the temporal stability of the system. However, generally λ on a phantom is much smaller than in vivo and therefore due to the square summation does not have a big impact on subject data.

Physiological noise is not noise in the traditional (stochastic) sense, but rather a structured confound. If one can measure the source of the fluctuations it is possible that these

confounds can be partially removed. A number of different corrections strategies have been developed to minimize physiological noise, some of which were introduced in section 3.5.1. In this chapter we use the method RETROICOR, which is currently the most commonly used technique for physiological noise mitigation.

4.2.1 Influence of Slice Acceleration on tSNR in 3D EPI

It is generally stated in the literature that 3D EPI is suitable for fMRI when working in the thermal noise dominated regime [81, 143], which is defined as $\frac{\sigma_p}{\sigma_0} < 1$. SNR_0 scales with B_0, therefore extra care should be taken to choose suitable acquisition parameters when working at high field strength. We explored this trade-off in vivo at 7 T in order to get a clearer picture of suitable 3D EPI parameters. Specifically, voxel volume and N_{k_z} impact whether or not the acquisition is in the physiological noise dominated regime.

The voxel volume and coverage (field of view in the slice direction) are often dictated by the scientific question. For example whole brain coverage might be required for resting state analysis and a minimum resolution might be necessary in order to localize activation to anatomical structures of interest. However, these choices in general also affect the achievable temporal resolution, which in turn affects the statistical power (formally, via the degrees of freedom achieved in a given experimental duration). Image acceleration is a parameter we are free to tune that can enable us to achieve a desired resolution or coverage without impacting on temporal resolution. We therefore decided to explore the dependence of tSNR on N_{k_z} for fixed coverage by controlling the degree of k_z acceleration (called the reduction factor R_z). We chose to do this for a set of different spatial resolutions. Our goal was to explore a potential interaction between acceleration (with higher R_z reducing SNR due to thermal noise and g-factor losses) and physiological noise (which increases with the volume acquisition time), such that the tSNR will reflect the balance of these competing processes.

Our hypothesis for the tSNR behaviour is displayed in a sketch in Fig. 4.3. Increasing the acceleration factor will increase the amount of thermal noise in the image due to

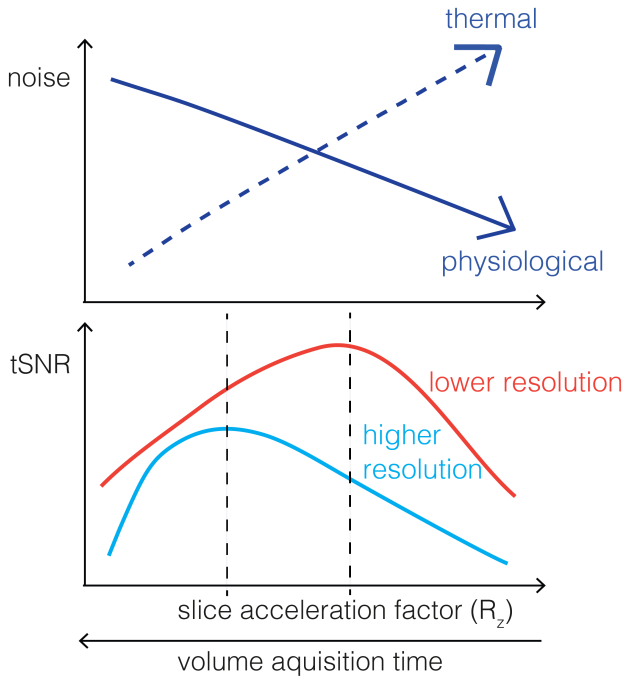


Figure 4.3: Sketch describing our prediction for the tSNR behaviour as a function of acceleration factor or volume acquisition time: Thermal noise (top sketch, dashed line): due to g-factor penalty and because of the reduced number of partitions acquired thermal noise increases with higher acceleration factors. Physiological noise (top sketch, solid line): with shorter acquisition time λ [143] becomes smaller, which means that physiological noise is reduced. The sketch in the bottom panel illustrates the tSNR behaviour we expect as a result of these two processes. The higher the resolution the more will the peak tSNR (indicated by black dashed lines) be shifted towards lower R_z factors.

g-factor noise and because the number of acquired k_z planes is reduced by a factor R_z. At the same time the reduction in volume scan time due to higher R_z reduces the λ, which in turn improves tSNR. We hypothesize that the tSNR will have a maximum when physiological noise contamination balances thermal noise losses due to acceleration (see bottom sketch in Fig. 4.3). The location of the peak tSNR is expected to depend on the voxel volume. For very small voxels we expect to already be in a thermal noise dominated regime, where there will be little tSNR benefit to shorter volume acquisition time. For increasing voxel volumes (i.e., as we enter the physiological noise dominated regime) we expect the peak tSNR to correspond to higher acceleration factors, where the shorter volume acquisition times reduce the amplification of physiological noise. Furthermore we expect that physiological noise correction would also move the peak of the curve to the left, i.e. to lower acceleration factors.

4.3 Methods

This section includes the description of the sequence, general testing of the sequence for fMRI acquisition and the main experiment, which tests the hypothesis described in section 4.2.1.

4.3.1 Sequence Development

The 3D EPI sequence was implemented by modifying the standard vendor-provided 2D EPI sequence, similar to the method used by Poser et al. [81]. This included the addition of phase encoding along k_z (see Fig. 4.4 for a sequence diagram). We chose linear ordering [144] of the slab encoding, starting with the most negative k_z location and ending with the most positive k_z. The 3D EPI sequence has parallel imaging capabilities along both phase encoding directions (skipping readout lines and skipping partitions) and also allows partial Fourier in the in-plane phase-encoding as well as the partition phase-encoding direction.

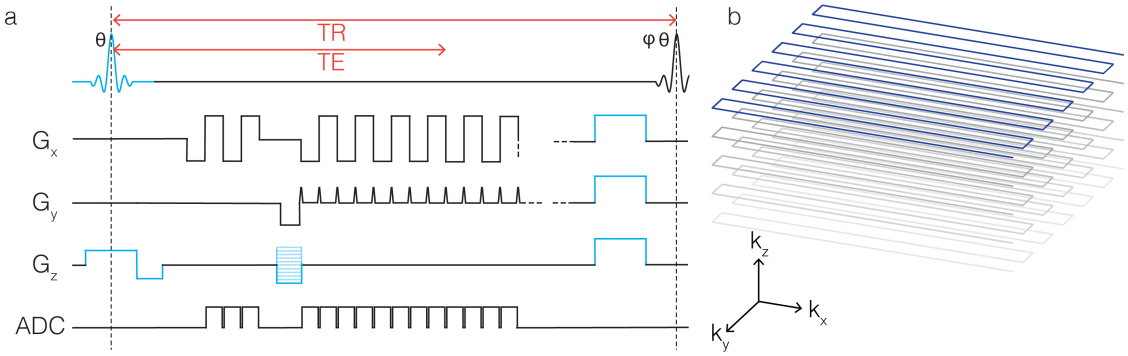


Figure 4.4: Segmented 3D EPI: (a) Pulse sequence diagram with differences to standard 2D EPI sequence indicated in light blue. This includes the excitation of a thick slab (instead of a thinner slice), the addition of phase encoding blips along z, and the inclusion of spoiler gradients (this is also used in 2D but is more important in 3D). (b) 3D k-space trajectory, highlighting the EPI plane collected in a single shot in dark blue.

For excitation we used a binomial pulse (1-2-1 triplet) to selectively excite the water protons [34], unless otherwise specified. This allowed shorter TRs than when using fat suppression, which selectively excites fat and de-phases it with spoiler gradients prior to a standard excitation pulse. It has been shown that spoiling the remaining signal prior to repeated excitation is crucial to signal stability in 3D EPI [145], for

which we used both gradient and RF spoiling. The RF spoiling utilized a quadratic phase increment of 117° [15]. The spoiler gradients were applied simultaneously on all three axes immediately after excitation, with the gradient amplitude chosen to result in a 2π phase dispersion across a voxel.

Data were acquired on a 7 T scanner (Siemens Healthcare, Erlangen, Germany), using a 32 channel head coil for reception and a birdcage head-only coil for transmit (both Nova Medical, Wilmington, MA, USA). All subjects in this study (as well as in all other chapters in this thesis) were healthy volunteers, scanned under standard operating procedures approved by the local ethics committee for technical development research. Details for all acquired data sets can be found in table 4.1. Before evaluating the tSNR behaviour as a function of acceleration factor and resolution (main experiment), we ran a series of experiments to qualitatively assess the sequence performance. This involved acquisition of functional data and tSNR maps, some of which were compared with matched 2D EPI protocols.

4.3.2 Sequence Testing and Comparison to 2D EPI

The 3D EPI sequence was tested on a healthy volunteer using a block paradigm (30 s OFF, 30 s ON; scan duration of approximately 5 minutes) involving visual stimulation with a 10 Hz flashing checkerboard. A data set with a nominal resolution of 0.67 mm isotropic was acquired with a volume TR of 4.8 s. The acquisition used an in-plane acceleration of factor 4, in-plane partial Fourier 6/8 and slice partial Fourier 6/8. It should be noted that all protocols in this chapter used in-plane acceleration because this is required to achieve TEs suitable for BOLD fMRI at 7 T.

We then proceeded to test how the 3D EPI sequence compares to standard 2D EPI. For a qualitative comparison between 2D and 3D EPI the same visual stimulus, as described above, was run on one subject. Both data sets were acquired with 1.5 mm isotropic resolution, a volume TR of 1.6 s and an in-plane acceleration factor of 3. For the 3D acquisition a k_z acceleration factor of $R_z = 3$ was used.

Table 4.1: Imaging Parameters of all in vivo experiments. Remark: For all 3D EPI acquisitions shown in this table an RF spoiling phase of 117° was used and gradient spoiling resulting in a phase dispersion of $0.8 \cdot 2\pi - 1 \cdot 2\pi$ across a voxel. The RF excitation pulse has a duration of 2.56 ms and the TBWP is in between 5.2 and 7.3.

Description		TE/TR/FA [ms]/[ms]/[°]	TR_{vol} [s]	resolution [mm]	matrix	R	PF	BW $[\frac{Hz}{pixel}]$	fat. suppr.	meas.	acq. t. [s]
Visual task 3D (7T)	Fig. 4.5										
3D EPI		33/100/18	4.8	0.67 iso	300 × 300 × 64	4 × 1	6/8 6/8	1042	none	60	339
Visual task 3D/2D (7T)	Fig. 4.6										
3D EPI		25/50/14	1.6	1.5 iso	136 ×136 × 96	3 × 3	8/8	1414	W.E.	180	314
2D sine EPI		65/1600/67	1.6	1.5 iso	136 × 136 × 28	3	8/8	1414	Fat Sat	190	312
tSNR 3D/2D (7T)	Fig. 4.7a										
3D EPI		29/65/16	4.16	2 iso	96 ×96 × 64	2 × 1	8/8	1408	none	50	228
2D sine EPI		29/4160/86	4.16	2 iso	96 × 96 × 64	2	8/8	1408	none	50	220
tSNR 3D/2D (7T)	Fig. 4.7b										
3D EPI		25/51/13	4.16	1.5 iso	150 ×150 × 80	3 × 2	8/8	1388	W.E.	220	472
2D sine EPI		26/2040/69	2.04	1.5 iso	150 × 150 × 64	3	8/8	1388	none	220	459
tSNR res/pat (7T)	Fig. 4.8 and 4.9										
3D EPI		20/41/12	3.9	2 iso	112 × 112 × 96	3 × 1	8/8	1396	W.E.	80	338
”		”	2.0	”	”	3 × 2	”	”	”	”	180
”		”	1.3	”	”	3 × 3	”	”	”	”	128
”		”	1.0	”	”	3 × 4	”	”	”	”	101
”		22/46/13	4.4	1.75 iso	128 × 128 × 96	3 × 1	”	”	”	”	378
”		”	2.2	”	”	3 × 2	”	”	”	”	202
”		”	1.5	”	”	3 × 3	”	”	”	”	143
”		”	1.1	”	”	3 × 4	”	”	”	”	117
”		25/51/14	4.9	1.5 iso	150 × 150 × 96	3 × 1	”	1388	”	”	419
”		”	2.4	”	”	3 × 2	”	”	”	”	223
”		”	1.6	”	”	3 × 3	”	”	”	”	158
”		”	1.2	”	”	3 × 4	”	”	”	”	125
”		29/60/15	5.8	1.25 iso	180 × 180 × 96	3 × 1	”	”	”	60	377
”		”	2.9	”	”	3 × 2	”	”	”	”	205
”		”	1.9	”	”	3 × 3	”	”	”	”	147
”		”	1.4	”	”	3 × 4	”	”	”	”	118
”		43/88/16	8.4	1.0 iso	226 × 226 × 96	3 × 1	”	1302	”	”	552
”		”	4.2	”	”	3 × 2	”	”	”	50	257
”		”	2.8	”	”	3 × 3	”	”	”	”	186
”		”	2.1	”	”	3 × 4	”	”	”	”	151

Finally, two different in vivo tSNR measurements were performed: In a first experiment 50 image volumes of 2 mm isotropic resolution data were acquired with only in-plane parallel imaging. In the second experiment the protocol was modified to 1.5 mm isotropic resolution and k_z acceleration of factor 2 (R_z=2).

4.3.3 Main Experiment: Impact of Acceleration on tSNR

In the main experiment we aimed to test our hypothesis illustrated in Fig. 4.3 regarding the impact of slice acceleration on tSNR, in which we achieved variable thermal SNR using protocols with a range of different spatial resolutions. All acquisitions had an in-plane FOV of 225×225 mm^2 and an in-plane acceleration factor of 3. The TE/TR were minimized for each spatial resolution and the flip angle was chosen as the Ernst angle for grey matter (see table 4.1). The bandwidth was approximately the same for all acquisitions and we acquired between 50 and 80 measurements (with a lower number of measurements for higher resolution to keep the overall experiment time manageable). 96 k_z planes were acquired for all scans, factoring in a slice over-sampling factor of 11.6 % (i.e. the outer 10 partitions were discarded to avoid contamination of the image with aliased signal). Time series were acquired at five different isotropic resolutions (1, 1.25, 1.5, 1.75, 2 mm) and four different slice acceleration factors (R_z = 1, 2, 3, 4). All 20 scans were acquired for one (pilot) subject; for three additional subjects, the resolutions were restricted to 1, 1.5 and 2 mm to reduce the scan time. The protocol was also repeated on a doped agar gel phantom to examine data without any physiological noise.

4.3.4 Image Reconstruction and Analysis

Image reconstruction was performed on the scanner using the default vendor-specified image reconstruction pipeline. This pipeline uses the GRAPPA algorithm [27] implemented by the vendor for parallel imaging and zero-filling for Partial Fourier reconstruction.

Most of the processing was done using the FMRIB software library FSL [146]. Each data set was motion corrected using MCFLIRT [133] and brain extracted using BET [147].

For task experiments, linear regression analysis was carried out using FEAT without any smoothing or statistical clustering. Finally tSNR was obtained voxel-wise as the mean signal over time divided by the temporal standard deviation.

All available volumes were used to calculate the tSNR, but a comparison using 50 volumes for all data sets (including those with 80 volumes) showed no substantial difference. The tSNR maps were registered to the structural space for each subject. Atlas-based cortical grey matter masks were used to calculate average cortical grey matter tSNR. For the phantom the tSNR was calculated using a mask that encompassed most of its volume. Physiological noise correction with RETROICOR [108] (using the PNM tool in FSL [148]) was conducted in two different protocols, using 34 and 12 regressors. We ran a three-factor, fixed-effects ANOVA to examine the interactions between resolution, R_z, and physiological noise correction on the mean tSNR across subjects.

4.4 Results

4.4.1 Sequence Testing and Comparison to 2D EPI

The activation z-score maps for our initial testing experiment are shown in Fig. 4.5. These images highlight the ability for the high-resolution 3D FMRI acquisition to clearly delineate activation of the primary visual cortex with high sensitivity and specificity. This initial validation demonstrated the feasibility of using the 3D EPI sequence for high-resolution FMRI at 7 T.

The results of the fMRI comparison between 2D and 3D EPI are shown in figure 4.6, where the raw z-statistic images (no clustering) are overlaid on an example functional image (see figure 4.6a). Time courses for an example voxel in the primary visual cortex are displayed in figure 4.6b.

Both average and peak z-stats were very similar for the two sequences. However for the same volume acquisition time we achieve whole brain coverage using the 3D EPI, while for the 2D case we only get a slab covering less than half of the cortex. This

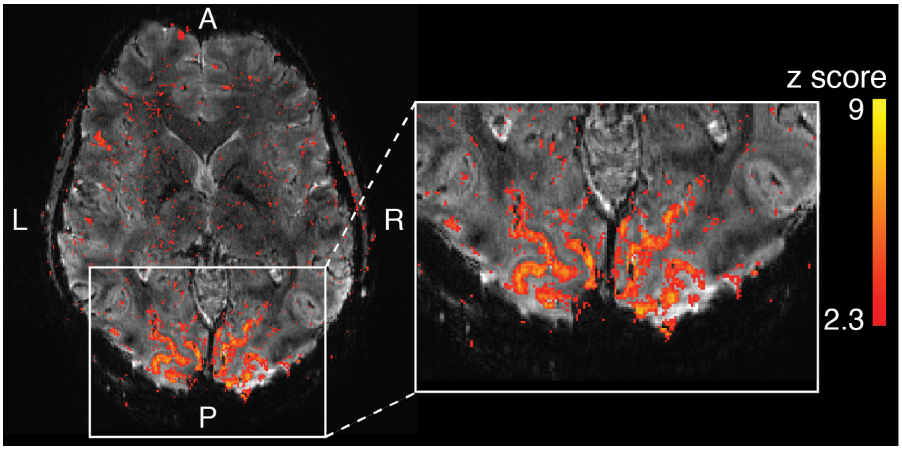


Figure 4.5: High-resolution (0.67 mm isotropic) functional MRI using 3D EPI: z-scores for activation for visual stimulus are overlaid on an example image.

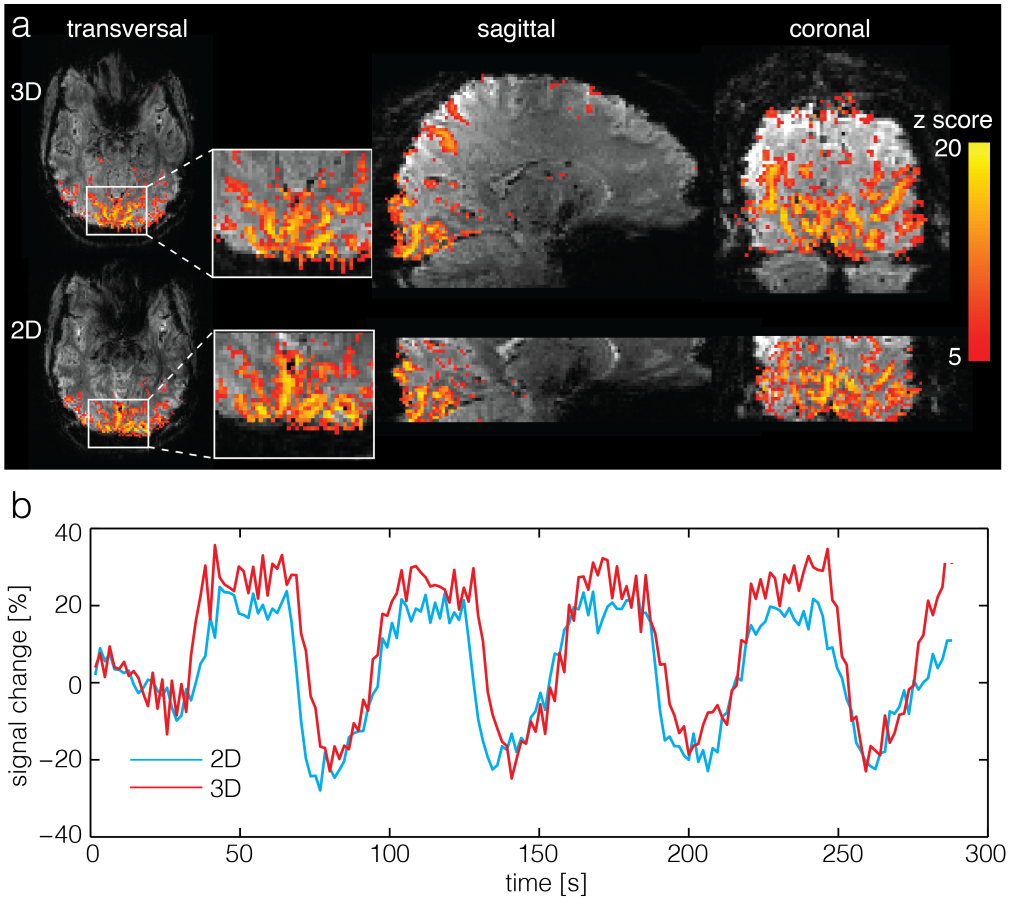


Figure 4.6: Visual task experiment for comparison between 2D and 3D EPI. (a) Activation maps for 3D (top) and 2D (bottom) acquisition in transversal, sagittal and coronal view. (b) Time courses for an example grey matter voxel in the visual cortex.

shows how the ability to accelerate in two dimensions allows us to either image faster or increase the coverage compared to conventional 2D imaging. Note that at the time

of these experiments, simultaneous multi-slice imaging was not yet routinely available, but that this technology now allows similar acceleration of 2D imaging.

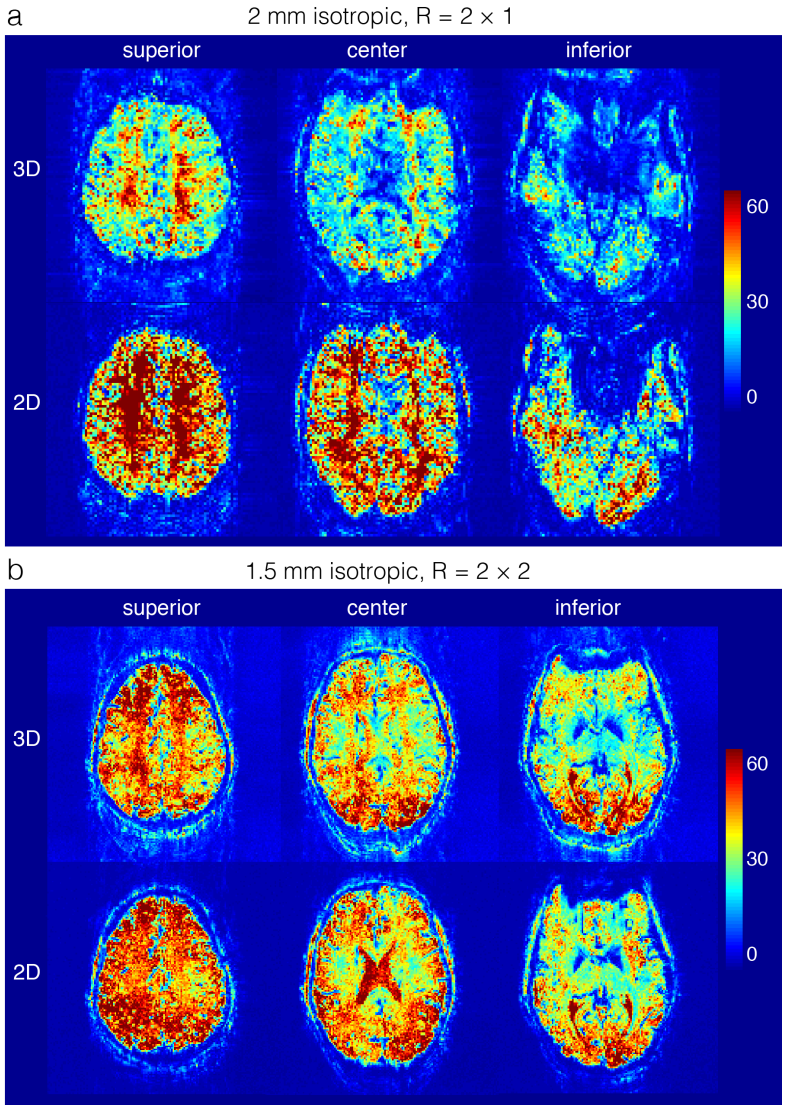


Figure 4.7: tSNR comparison for two matched protocols using the 2D or 3D EPI sequence.

The tSNR comparison between 2D and 3D EPI for two different protocols is shown in figure 4.7. Contrary to predictions based purely on thermal SNR, the temporal SNR was significantly higher (almost two-fold) for the 2D acquisition compared to 3D in the 2 mm experiment (see tSNR maps in Fig. 4.7a). For the 1.5 mm experiment the tSNR for 2D and 3D EPI was comparable (see Fig. 4.7b). The main differences between these two experiments are the smaller voxel size and the parallel imaging factor along the slab phase encode axis. Larger voxels result in higher SNR_0, therefore we expect

a higher amount of physiological noise for this acquisition. Longer volume acquisition times (due to the lack of k_z acceleration) also make the acquisition more susceptible to physiological noise. These two effects seemed a likely explanation for the large discrepancy in the performance of 3D relative to 2D for these two experiments, and these initial results motivated our subsequent experiments to investigate the relationship between k_z acceleration and physiological noise, described in the following section.

4.4.2 Main Experiment: Impact of Acceleration on tSNR

The raw tSNR maps for the pilot subject are displayed in Fig. 4.8, for five different spatial resolutions and R_z=1-4. These data exhibited a complex variation in tSNR across acceleration factors and did not behave as expected. Looking along individual columns, the expected reductions in tSNR with increased spatial resolution are observed. However, looking along the rows, the tSNR can be seen increasing with R_z at low spatial resolutions, but not in a consistent fashion. At some resolutions (e.g., 2 mm), the data conform to the model we suggested with a peak in tSNR occurring at $R_z = 2$. However this trend across acceleration factors is not consistent. For example the 1.75 mm data set has surprisingly low tSNR for the $R_z = 2$ data set, compared to all other R_z factors used. In the data sets with 1.5 mm resolution or higher the tSNR drop due to g-factor losses can be clearly seen in the centre of the brain, while tSNR increases in the periphery for some of the resolutions.

The average tSNR across all experiments is summarized in Fig. 4.9. The phantom results show the pattern predicted by pure thermal effects, with a reduction in tSNR with increasing R_z (left column). The tSNR for all four subjects is shown in the middle column. For the 1 and 1.5 mm data we see the predicted pattern, where tSNR increases with moderate R_z factor, contrary to what would be observed under pure thermal noise. Moreover, the tSNR peak occurs at lower acceleration for the 1 mm data set than for the 1.5 mm, also in agreement with our hypothesis. However, our 2 mm data is not consistent with this picture, displaying a more complicated variation with R_z that is difficult to interpret.

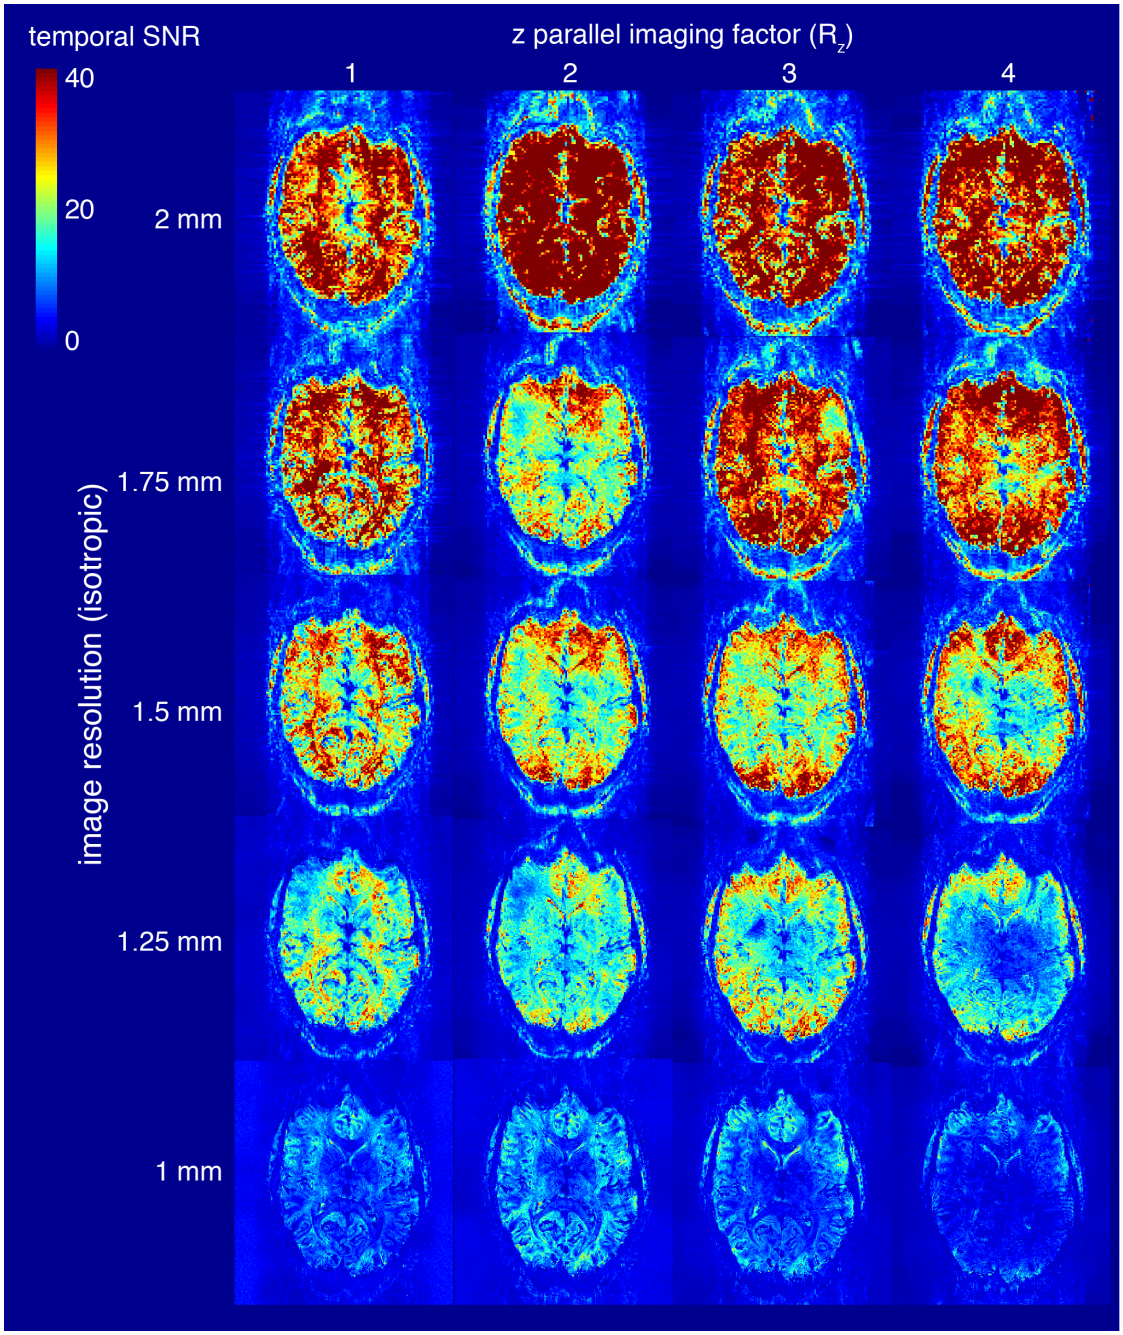


Figure 4.8: tSNR as a function of resolution and z parallel imaging factor (R_z) for an example subject.

Moreover, physiological noise regression did not alter the form of the dependence on R_z factor (right column in Fig. 4.9). We predicted that, in addition to an overall increase in the tSNR, the peak in tSNR would shift to lower R_z factors after physiological noise correction. Instead we see uniformly increased the tSNR across all R_z factors (see figure 4.9). The results of the ANOVA indicated significant main effects of each factor (resolution,

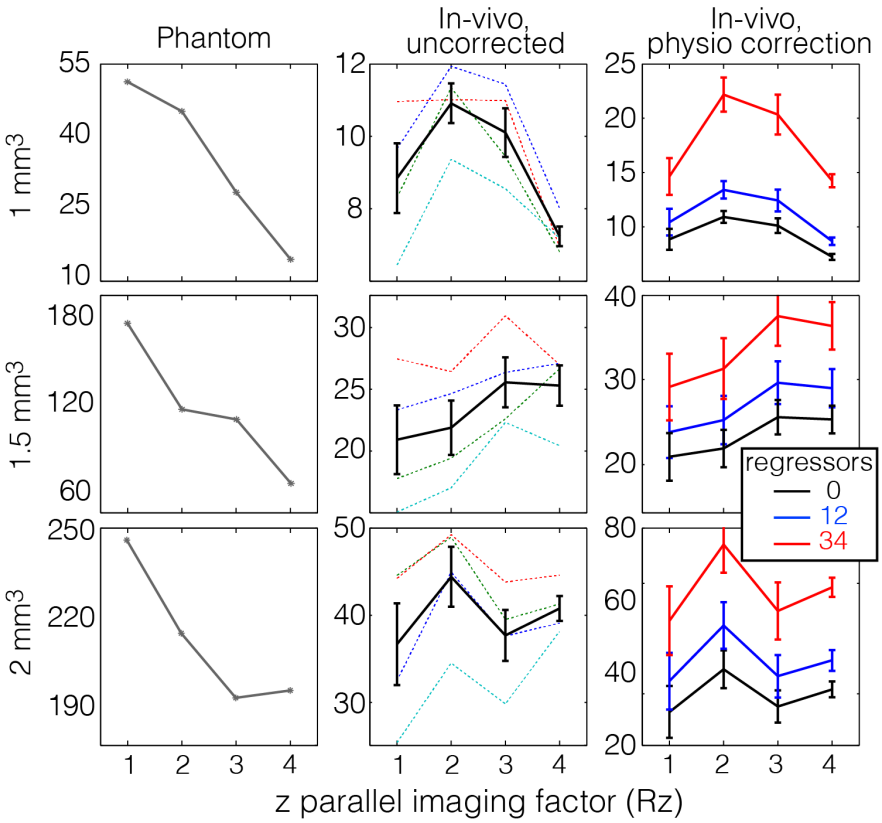


Figure 4.9: tSNR for different R_z factors: tSNR curves observed in phantom (left) and in vivo (middle and right) data. Error bars indicate standard error across subjects (the individual subjects are shown as dashed lines).

R_z and physiological correction), as well as significant interactions of resolution with both R_z and physiological correction (all effects $p < 0.05$).

4.5 Discussion

We successfully implemented and tested an accelerated 3D EPI sequence and acquired fMRI data at 7 T. The high-resolution functional data collected allowed precise localisation of activation to the grey matter architecture. The ability to achieve high spatial resolution along with larger brain coverage and/or higher temporal resolution makes 3D EPI an attractive option for neuroscience applications.

Contrary to our expectations, the functional comparison for 2D and 3D showed comparable results overall. For the visual task experiment, all of our protocols were well above the

detection threshold. In future work it might be useful to compare the performance of different protocols using resting state fMRI or tasks that elicit a more subtle response and therefore more crucially depend on high SNR.

To our knowledge, the effect of R_z acceleration on temporal stability in 3D EPI had not been explored, although alteration of physiological noise sensitivity with segmentation has been reported [143]. While our data are not entirely consistent with our hypothesis, the significant resolution $\times$ R_z interaction suggests that a structured underlying mechanism is driving the complex tSNR behaviour. Furthermore, because no change in the shape of the curves was observed with physiological noise correction, the precise effect of volume-based physiological noise regression on 3D data remains unclear, and warrants further investigation. Because of the very long volume acquisition times it might be beneficial to perform k-space based corrections instead of RETROICOR. However, as described in chapter chapter 3, it is not yet clear what the best correction approach is for 3D EPI [114].

Previously reported values of tSNR for 3D EPI demonstrate that our tSNR values are consistent with the literature (see table 4.2 in appendix 4.A). Assuming physiological noise follows the Krüger model, a simple simulation confirmed the observation that the tSNR of 3D EPI is strongly dependent on the SNR_0 regime. We used equations 4.1-4.4 and the λ values published by van der Zwaag et al. [143] to calculate the expected tSNR for a phantom case (assuming constant λ) and the in vivo case (assuming λ varying linearly with N_{k_z}). For the phantom case 3D EPI theoretically always provides higher tSNR, however the in vivo case (with linearly varying λ for the 3D acquisition) displays much more complex behaviour (see Fig. 4.10). This confirms that 3D EPI is only beneficial over 2D EPI for the low SNR_0 regime and it is highly dependent on the choice of N_{k_z}. This is further illustrated in Fig. 4.11, which demonstrates how the peak tSNR shifts to lower N_{k_z} for higher SNR_0. In Fig. 4.11 we also show the simulated tSNR for same N_{k_z} values as used in the main experiment. This revealed similar trends to what we saw in the in vivo data for 1 mm and 1.5 mm; however, the tSNR behaviour of our 2 mm data shown in Fig. 4.9 remains unexplained. This simulation should be seen as a coarse approximation; its limitations and further details are discussed in appendix 4.A.

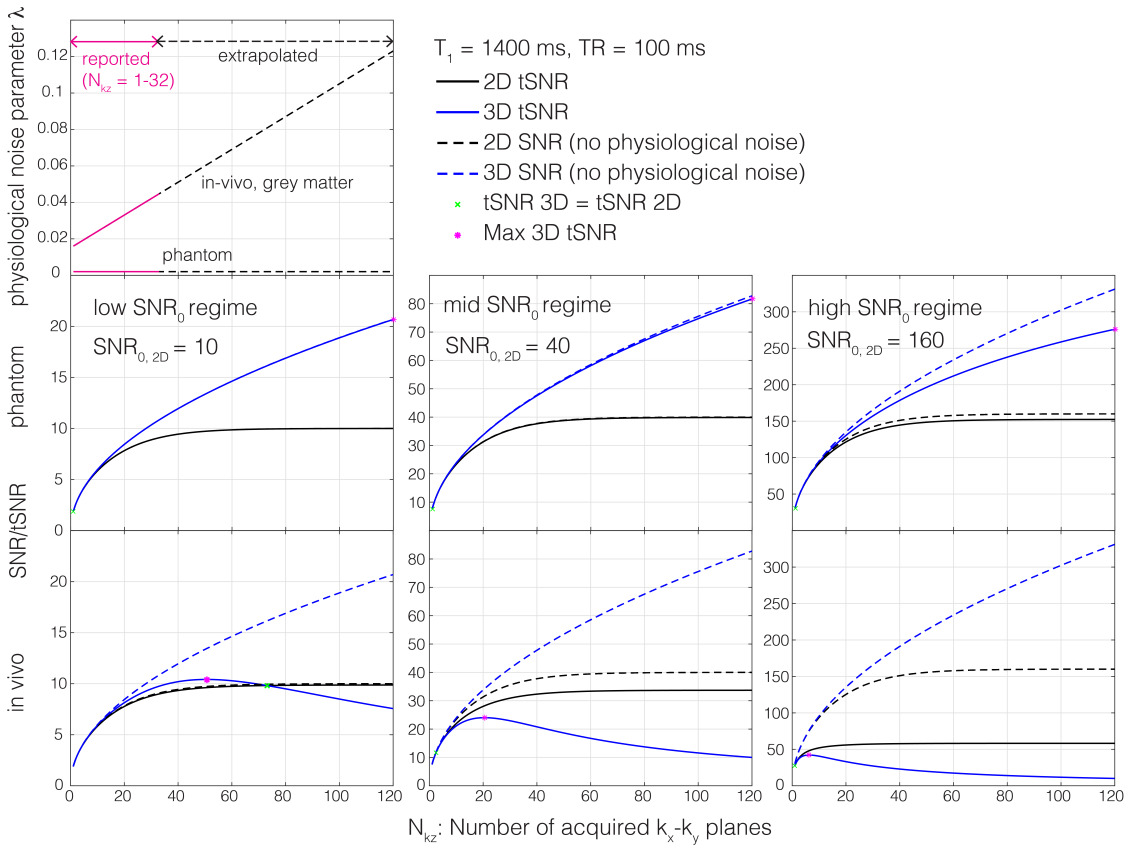


Figure 4.10: Simulation of SNR_0 and tSNR of 2D/3D EPI as a function of the number of acquired slices/partitions. The tSNR calculations take into account λ values as reported by van der Zwaag et al. [143]. The SNR_0 for the fully relaxed 2D case is set to three different values, defining three different SNR_0 regimes.

Due to concerns over potentially low tSNR, we performed a series of phantom experiments, documented in appendix 4.B. The results we observed remain partially unexplained, although based on these phantom experiments it seems likely that much of the unexplained behaviour is likely not unique to in vivo imaging.

The ability to accelerate in a second dimension using 3D EPI is its primary benefit. This should allow higher tSNR efficiency (tSNR divided by the square root of the volume acquisition time, also called sensitivity). Historically tSNR had been viewed as the most relevant factor for fMRI and due to the slow hemodynamic response, accelerating beyond a volume TR of 2-3 seconds was not considered beneficial. However, it has been shown that many functional experiments profit from more temporal degrees of freedom [74, 149]. Therefore, tSNR efficiency should be considered the more relevant parameter to predict BOLD sensitivity.

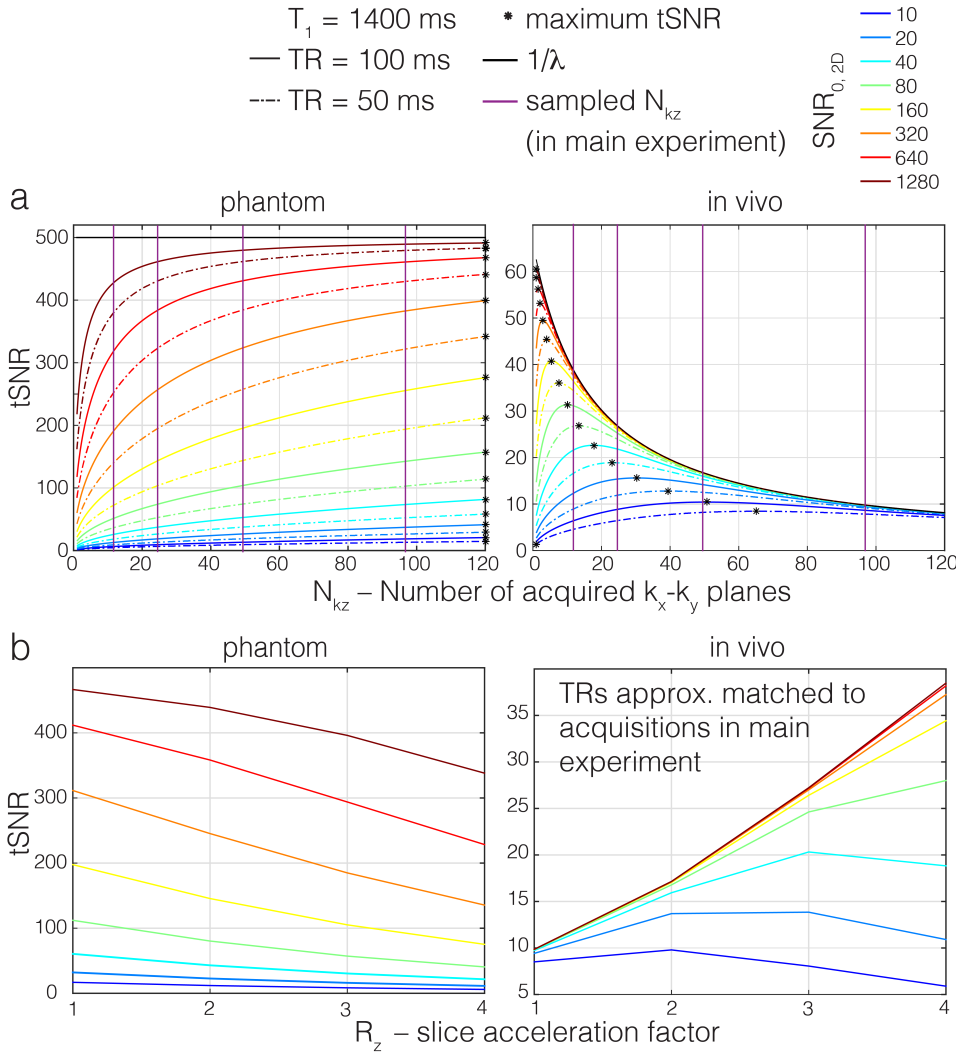


Figure 4.11: Simulation of tSNR for 3D EPI for different SNR_0 regimes. (a) The peak tSNR is indicated with an asterisk. Purple vertical lines indicate the acquisition parameters used for the main experiment. (b) tSNR prediction for the main experiment using a range of SNR_0 and similar TR $\sim$ 35–88 as used for the acquisitions. The curves for the lower SNR_0 values match the acquired curves for the 1 mm and 1.5 mm experiments relatively well (see Fig. 4.9), while the behaviour of the 2 mm data set could not be replicated with this simple simulation.

It is also clear that the benefit of 3D EPI is primarily achieved at ultra high resolution (1 mm to sub-mm). As physiological noise increases proportionally with the signal level, the sensitivity advantages at high field can be better used when going to small voxels and therefore being in the thermal noise dominated regime. Even if high resolution is not required it might be beneficial from an tSNR point of view to acquire at a higher spatial resolution in order to reduce the relative impact of physiological noise (i.e. stay in the thermal noise dominated regime) and then retrospectively smooth down to larger voxels,

which allows regaining SNR [150]. In practice a direct comparison with 2D EPI at these very high resolutions is difficult because it is challenging to achieve very accurate thin slices.

On a side note: Even at more moderate spatial resolutions (1.5 mm) some of our 2D EPI data were visibly blurrier in the slice direction compared to the 3D EPI. When performing tSNR comparisons between different methods and protocols it is very important that the true voxel volume (and thus the effective spatial resolution) is matched and not the nominal one. T_2^* blurring in very long readout trains as well as use of Partial Fourier (when zero-filling as done in the default Siemens reconstruction pipeline) causes substantially lowered effective resolution.

During and after completion of this project the fMRI field made a number of important improvements to 3D EPI. Most notably CAIPI-style [84, 85] sampling was introduced for 3D EPI, called 3D-CAIPI EPI. This more distributed sampling pattern allows higher undersampling factors, and consequently higher tSNR efficiency, compared to the standard 3D EPI. Furthermore, it has been shown that using 2D acceleration (e.g. 2×2) without CAIPI can result in higher g-factor noise than the equivalent 1D acceleration (e.g. 4×1) for certain coil geometries and FOV positions[151]. In 3D RE-EPI [152] two partitions are encoded per shot, exploring the hybrid space between EVI and 3D segmented EPI, with researchers reporting increased temporal resolution and better physiological noise characterization compared to standard 3D EPI. Semi-elliptical sampling has been used to further increase the efficiency of 3D EPI [153] and a new fast fat suppression scheme has been proposed [98].

Furthermore, the importance of the acquisition order of the segmented parallel imaging reference data has been characterized [154]. All the data acquired in this chapter were acquired with the default ordering scheme, which acquires a given segment for all partitions first. We improved this during the sequence development described in the next chapter. We observed relatively high level of image artefacts in the highly accelerated images (using acceleration factors 3×3 and 3×4). The vendor implementation of the GRAPPA reconstruction, which was used for this project, is not optimal and this could lead to unexpected results for 2D acceleration. A new improved vendor reconstruction (icePAT)

has been made available since this study, which should be used for future work on 3D EPI. However, while all these factors can impact the tSNR, it is not clear that they would explain the complex tSNR behaviour we observed.

4.6 Conclusion

We have implemented a 3D EPI sequence with full parallel imaging capabilities. The ability to achieve high spatial resolution and either larger brain coverage or higher temporal resolution makes 3D EPI an attractive option for neuroscience applications. However, the theoretical gains are not always achievable due to physiological noise, which has been established to have a greater relative effect on 3D acquisitions than 2D. We observed some partially inconclusive tSNR behaviour. Overall, we found that it is crucial to use 3D EPI in the thermal noise dominated regime and/or perform physiological noise correction on the 3D EPI data. In addition to physiological noise, motion is a major challenge in 3D EPI, which is the focus of the next two chapters. Due to increased challenges at 7 T (and since the results reported in this chapter could reflect system instabilities), we chose to implement our proposed non-Cartesian EPI sequence (chapter 5) and motion correction approach (chapter 6) at 3 T.

4.A Appendix: Literature Comparison and tSNR Simulation

Over the last decade a range of studies have looked at 3D EPI and the impact of temporal fluctuations on 3D EPI [81, 112, 113, 143, 145]. Poser at al. [81] compared the sensitivity of 3D and 2D EPI at 7 T at resolutions in the range of 1.5-3.0 mm isotropic. They found that the 3D method was of better or comparable tSNR efficiency than 2D EPI for all the tested resolutions; however the tSNR itself was lower. Higher sensitivity for 3D EPI over conventional 2D EPI was also confirmed by Stirnberg et al. [155]. Other studies analysed the impact of physiological noise correction on 2D and 3D EPI. They showed that physiological noise correction provides larger increase in tSNR for 3D EPI than for

2D EPI at 7 T [113] and 3 T [112]. The general consensus is that 3D EPI is promising for high resolution fMRI where the relative physiological noise contribution is lower. The tSNR values for the most important studies are summarized in table 4.2, demonstrating that the tSNR values we report are within the range of values observed.

More recent studies have compared 2D SMS EPI with 3D CAIPI EPI. One comparison showed substantially increased sensitivity for the highly accelerated 3D EPI protocols over SMS [156], while another study showed higher sensitivity for 2D SMS EPI [157]. A recent study also suggested that stimulus correlated physiological noise differences in cognitive tasks could cause a bias between fMRI results from 3D EPI and 2D EPI [158].

We performed a set of simple simulations in order to evaluate whether the tSNR behaviour we saw in our experiment could be predicted using the Krüger model for physiological noise (as described in 4.2). The original Krüger paper [67] reported λ values of $\lambda \sim 0.012$ for cortical grey matter, $\lambda \sim 0.006$ for white matter and $\lambda \sim 0.002$ for a phantom. Triantafyllou et al. [159] found that λ is independent of field strength and reported values of around $\lambda \sim 0.01$ in grey matter and $\lambda \sim 0.00075$ in a phantom. Van der Zwaag et al. [143] studied temporal characteristics of 3D EPI (at 7 T, using TE $\sim$ 28ms) and found that in vivo the temporal instabilities increase with the number of partitions acquired. For example for grey matter they measured $\lambda \sim 0.016 + 0.0009 \cdot N_{k_z}$. In phantom data they found no dependence of tSNR on the number of partitions ($\lambda \sim 0.002$).

For the simulation we used equations 4.1–4.4 (see section 4.2) and the λ values published by van der Zwaag et al. extrapolated to a larger N_{k_z} range. The λ values used for the simulation are shown in the top panel in Fig. 4.10. We simulated tSNR for a range of different SNR_0 regimes, by setting the SNR_0 for the fully relaxed 2D case and scaling the 3D SNR relative to that according to equations 4.1 and 4.2. Fig. 4.10 shows tSNR as a function of N_{k_z} for 2D and 3D EPI with and without physiological noise contribution. In the phantom case the simulated 3D EPI tSNR is always higher than the 2D EPI tSNR. The tSNR for 2D EPI initially rises with more slices, but plateaus as soon as the number of slices is sufficient for the magnetisation to fully relax. The SNR_0 for the 2D

Table 4.2: Example 3D EPI tSNR values (for grey matter ROIs) reported in the literature. If the publication included a comparison to 2D EPI, these values were also reported. We included tSNR or sensitivity (depending on what the study reported). When the study used physiological noise correction (Lutti [112], Jorge [113] and Narsude [85]), the tSNR values using physiological noise correction were included in the respective column. If values are approximate (for example because they had to be read from a figure) they are indicated with a ∼.

Reference	Sequence	tSNR	Sens.	Corr. tSNR	Corr. Sens	B_0 [T]	Resolution [mm^3]	TR_{vol} [s]	Acceleration factor
Poser et al. (2010) [81]	3D EPI	30	17	-	-	7	1.5 × 1.5 × 1.5	1.88	2 × 2
		32	24	-	-		2 × 2 × 2	3.07	2 × 2
	2D EPI	31	13	-	-		1.5 × 1.5 × 1.5	5.96	2
		42	22	-	-		2 × 2 × 2	3.65	2
Van der Zwaag et al. (2012) [143]	3D EPI	∼20-30	-	-	-	7	2 × 2 × 2	1.2-4.8	1 × 1
		∼18-28	-	-	-		1.5 × 1.5 × 2	1.2-4.8	1 × 1
	2D EPI	∼25	-	-	-		2 × 2 × 2	0.15	1
		∼10	-	-	-		1.5 × 1.5 × 2	0.15	1
Lutti et al. (2013) [112]	3D EPI	∼26	-	31	-	3	1.5 × 1.5 × 1.5	3.2	2 × 1
	2D EPI	∼22	-	24	-			3.2	2
Jorge et al. (2013) [113]	3D EPI	-	-	∼30-49	-	7	2 × 2 × 2	3.2	2 × 1
	2D EPI	-	-	∼33-44	-			3.2	2
Narsude et al. (2016) [85]	3D CAIPI EPI	-	∼10	-	∼12	7	2 × 2 × 2	3.2	1 × 1
		-	∼20	-	∼30			1.06	1 × 3
		-	∼30	-	∼42			0.37	1 × 3, PF = 6/8

EPI case is not high enough for the physiological noise reduction to be substantial. In the case of 3D EPI the tSNR increase with N_{k_z} is flattened due to physiological noise and asymptotes at $\frac{1}{\lambda} = 500$. The range of the plot is not sufficient for the asymptotic behaviour to be seen. For 2D EPI the behaviour in the in vivo case is analogous to the phantom case, however the higher λ dictates a maximum possible tSNR value of $\frac{1}{\lambda} \sim 63$. The 3D EPI, with linearly increasing λ, shows much more complex behaviour, which is illustrated in Fig. 4.11. The comparison to 2D EPI in Fig. 4.10 shows, that 3D EPI only has higher tSNR in the low SNR_0 regime and for certain N_{k_z}.

Fig. 4.11 shows tSNR for 3D EPI for a range of SNR_0 values and two different TRs. The peak tSNR shifts to lower N_{k_z} for higher SNR_0. We displayed the same data for only the N_{k_z} acquired in the main experiment in this chapter. The shape of the tSNR curves which we observed for the 1 mm and 1.5 mm data can be explained with this simulation. It should be noted that we were able to reach higher tSNR values in the in vivo experiments than the theoretical upper limit predicted by the simulation.

The simulation does not provide a fair comparison of 2D and 3D EPI if N_{k_z} is changed by altering the slice-acceleration for 3D EPI. In this case 3D EPI maintains the same coverage, while the 2D EPI coverage is reduced. Ideally we should compare protocols with matched coverage and compare tSNR efficiency rather than tSNR. The simulation also does not take into account the g-factor penalty, which alters the SNR_0 regime and thus could potentially change the shape of the observed curves. However, it is unlikely that this would explain the behaviour we observed in the 2 mm data (specifically the dip for $\mathrm{R}_z = 3$).

The basic Krüger model does not hold for multi-channel coils and therefore does not adequately describe the 32 channel data we acquired. An extension for array multi-channel coils has been proposed [160] which should be used for future simulations. We extrapolated λ to a regime, which has not been measured; it is possible that the relationship asymptotes, making our predictions too conservative. Furthermore, the parameter λ is TE-dependent and we applied a literature value measured at TE = 28 ms to simulate our data acquired with TE = 20-43 ms. This results in an over-estimation of the physiological

noise contribution for the lower resolution experiments and an under-estimation it for the high-resolution ones. It is also possible that the dependence of λ on the number of partitions could change with TE.

Likely, the Krüger model does not describe the behaviour of physiological noise in segmented 3D EPI very well. This is supported by the results of van der Zwaag et al. [143], who reported poor quality of fit to the model. It would be useful to expand the model in order to better understand the impact of physiological noise in 3D EPI and more effectively remove it. Despite these limitations the simulation explained the behaviour seen in vivo reasonably well. Both the simulations and the in vivo data suggest that physiological noise has a very detrimental effect on 3D EPI and we can conclude that it is very important to chose suitable imaging parameters and perform physiological noise correction.

4.B Appendix: Investigation of the Temporal Stability of 3D EPI (Phantom Experiments)

We performed a series of phantom experiments to test if SNR and tSNR behave as expected. Firstly, we wanted to confirm the theoretically expected SNR benefit of 3D over 2D for a range of different number of partitions. Secondly, we wanted to confirm the λ value for phantoms, which has been reported in the literature [143, 159].

4.B.1 2D vs. 3D EPI Phantom Comparison

Measurements of the temporal SNR were performed on phantoms and compared to the theoretically predicted tSNR improvement for 3D imaging over 2D imaging. The data was acquired at 7 T with a birdcage volume coil for transmit and receive, using an oil phantom ("Siemens Yellow" phantom, $T_1/T_2 \sim 1420/750$ms at 7 T). The number of k_z planes acquired was varied between 8 and 80 for both the 3D[1] and the 2D EPI sequence. The resulting tSNR maps (Fig. 4.12a) showed the expected behaviour for 2D, with the tSNR

[1]Full imaging parameters 3D EPI: TR/TE/flip angle = 100 ms/42 ms/25°; BW = 1736 Hz/pixel; in plane matrix = 96×96; k_z planes = 1,8,16,32,64,80; 2 mm iso resolution, no acceleration; 2π gradient spoiling, RF spoiling with 117°, no. of meas. = 50

increasing initially and then reaching a plateau when the magnetisation can fully relax. However, for the 3D acquisition we saw spatial heterogeneity (instead of the expected uniform tSNR maps), which becomes more pronounced with larger number of partitions acquired. This spatial pattern makes it difficult to get a quantitative tSNR estimate. It appears that there is temporal instability, which negates any potential benefit from the increased thermal SNR when acquiring more partitions.

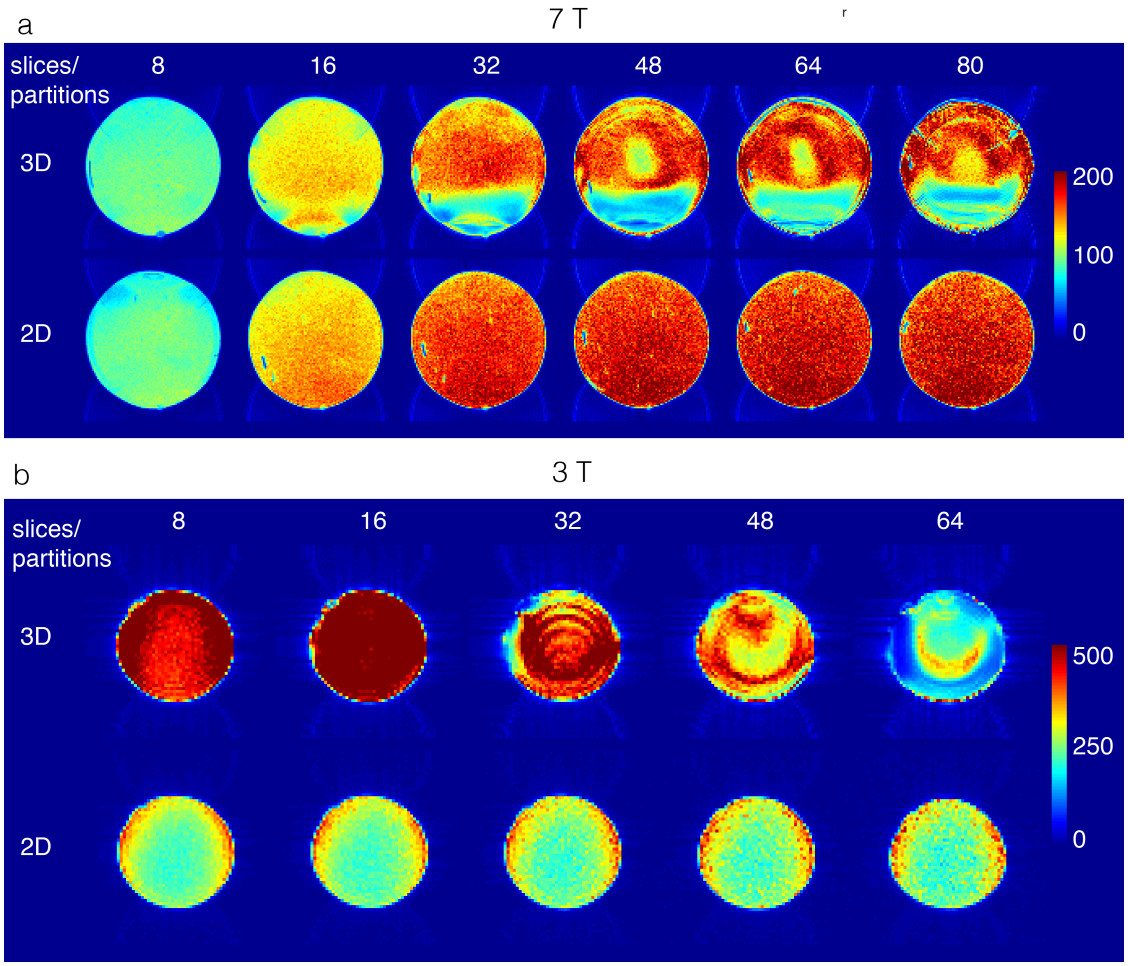


Figure 4.12: Temporal SNR for different numbers of partitions at 7 T and 3 T. Note that different phantoms (with different T_1) were used for the two experiments.

The experiment was repeated on a 3 T Verio scanner using a 32 channel receive coil and a cylindrical doped (very short $T_1 \sim 40$ ms) water phantom. The spatio-temporal patters observed were different, but again did not follow the theoretically predicted trend (see Fig 4.12b).

4.B.2 λ Measurements

In a second 7 T experiment, phantom data was acquired at different resolutions and number of k_z planes, in order to analyse plots of tSNR vs. SNR and to determine the physiological noise parameter λ. This was done using the protocols[2] as described by van der Zwaag et al. [143]. In order to exclude effects due to the use of multi-channel coils, which complicates SNR analysis, the data was acquired using a birdcage volume coil for transmit and receive. We used the same oil phantom as described above. Image reconstruction from raw k-space data was performed offline in Matlab in order to exclude the possibility that any of the problems seen above could be caused by the vendor image reconstruction.

The SNR was evaluated using a signal ROI in the centre of the phantom and a noise ROI in the image background (located far away from the phantom to avoid overlap with residual ghosted signal, which was facilitated by the larger FOV due to readout over-sampling). The mean tSNR values were obtained by averaging over the signal ROI (as used for the SNR calculation) in the centre of the phantom.

For phantoms experiments we expect an increase in SNR to be accompanied by an increase in tSNR, due to the absence of physiological noise. Therefore, any deviation from the line of identity is due to hardware related temporal instabilities or an instability in the phantom (e.g. vibration or movement of the phantom fluid). In the literature values of $\lambda = 0.0008$ have been reported and that λ for a phantom is independent of the number of partitions acquired [143]. Our measurements did not conform to these expectations, and showed substantially larger λ, which may indicate system instabilities. The tSNR maps (see Fig. 4.13a) for a set number of partitions show very consistent spatio-temporal patterns across the different acquisitions (compare the columns in Fig. 4.13b).

We undertook a series of tests that aimed to discover the source of these discrepancies; including a cross-site comparison with research groups in Maastricht and Nijmegen, which is described in the section below. To rule out shot to shot inconsistencies due to

[2]Full imaging parameters 3D EPI: TR/TE/flip angle = 150 ms/41 ms/30°; BW = 2604 Hz/pixel; in plane matrix = 96×96; k_z planes = 1,8,16,32; resolutions = 2×2×2, 3×3×2, 4×4×2, 5×5×2, no acceleration; 2π gradient spoiling, RF spoiling with 117°, no. of meas. = 50

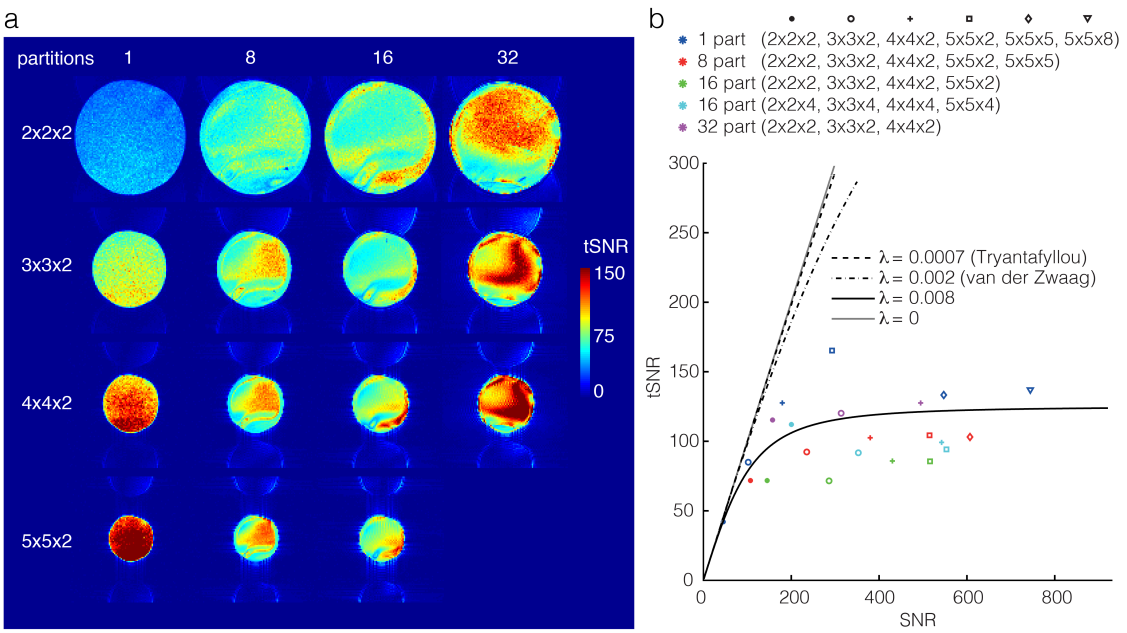


Figure 4.13: Experiment to measure λ: Panel (a) shows tSNR maps for most of the data sets acquired. They vary in voxel size and the number of partitions. Panel (b) shows tSNR plotted over SNR for all data sets acquired. The plots predicted by the Kruger model are also included in the figure.

fluid movement in the oil phantom, the experiment was repeated with agar phantoms. While the resulting spatial pattern varied with phantom type and the exact protocol, the anomalous effects were present in all cases. Extra padding was added to reduce mechanical vibration of the phantom, also to no effect. Multiple RF spoiling phase increments and gradient spoiler moments were tested, but showed no substantial difference. We also varied the time bandwidth product of the RF pulse between 3 and 24 and tested the protocol using non-selective RF pulses, both to no effect. Turning off the cryo-cooler improved tSNR but did not remove the spatio-temporal patterns (see 4.B.3).

4.B.3 Cross-site Comparison

Over the course of this project we compared our implementation of the 3D EPI sequence with the version from Maastricht (Benedikt Poser) and a version from the Bonn/DZNE (Ruediger Stirnberg). In both cases tSNR maps were acquired on the fMRIB 7 T for phantoms and in vivo. No substantial differences between the performance of the

different sequences was detected, ruling out a faulty implementation as the source of the tSNR behaviour.

We also performed a cross-site tSNR comparison in collaboration with research groups from Maastricht and Nijmegen/Donders (Lauren Bains). All sites installed the Maastricht version of the 3D EPI sequence and image reconstruction (as stated above a previous sequence comparison showed no substantial difference). The tests were run with the oil phantom used in 4.B.2, which was available at all three sites. We used the 3D EPI protocol as run in section 4.B.1, acquiring between 8 and 80 k_z planes. The 7 T scanners at FMRIB and at Maastricht had identical gradient hardware (SC72), while the 7 T system in Nijmegen had different hardware. At FMRIB we additionally repeated the experiment with the cryo-cooler/cold-head (helium cooling pump) temporarily switched off.

None of the phantom measurements fit the predicted model of SNR vs. number of partitions. Instead of the expected tSNR increase with partition number, the observed tSNR plateaued and even slightly decreased for increasing partition numbers. There were clear differences between the tSNR maps collected at the three sites (see Fig. 4.14). The tSNR in Donders and Maastricht was substantially higher (up to 100 %) compared to the FMRIB measurements, with the values from Donders in turn being higher than the ones from Maastricht. The large discrepancy between the Maastricht and FMRIB results was especially surprising as these scanners are nominally almost identical. Potential differences could have been caused by different placement of the phantom within the coil and fixation of the phantom (amount of padding). Also contributing to the discrepancy could be the quality of the shim, the transmit adjust settings and how long the phantom was left to settle in the scanner before starting the measurements (potentially resulting in different amount of fluid motion), although as described above we did investigate these possibilities in Oxford. Another potential source of instability could be larger vibrations in the patient bed.

Turning off the cryo-cooler increased the tSNR in the FMRIB data by up to 50 %. The striping artefact in the partition encode direction was reduced when the cold head was turned off. A PCA on the data showed a high variance component at $\sim$1 Hz, which

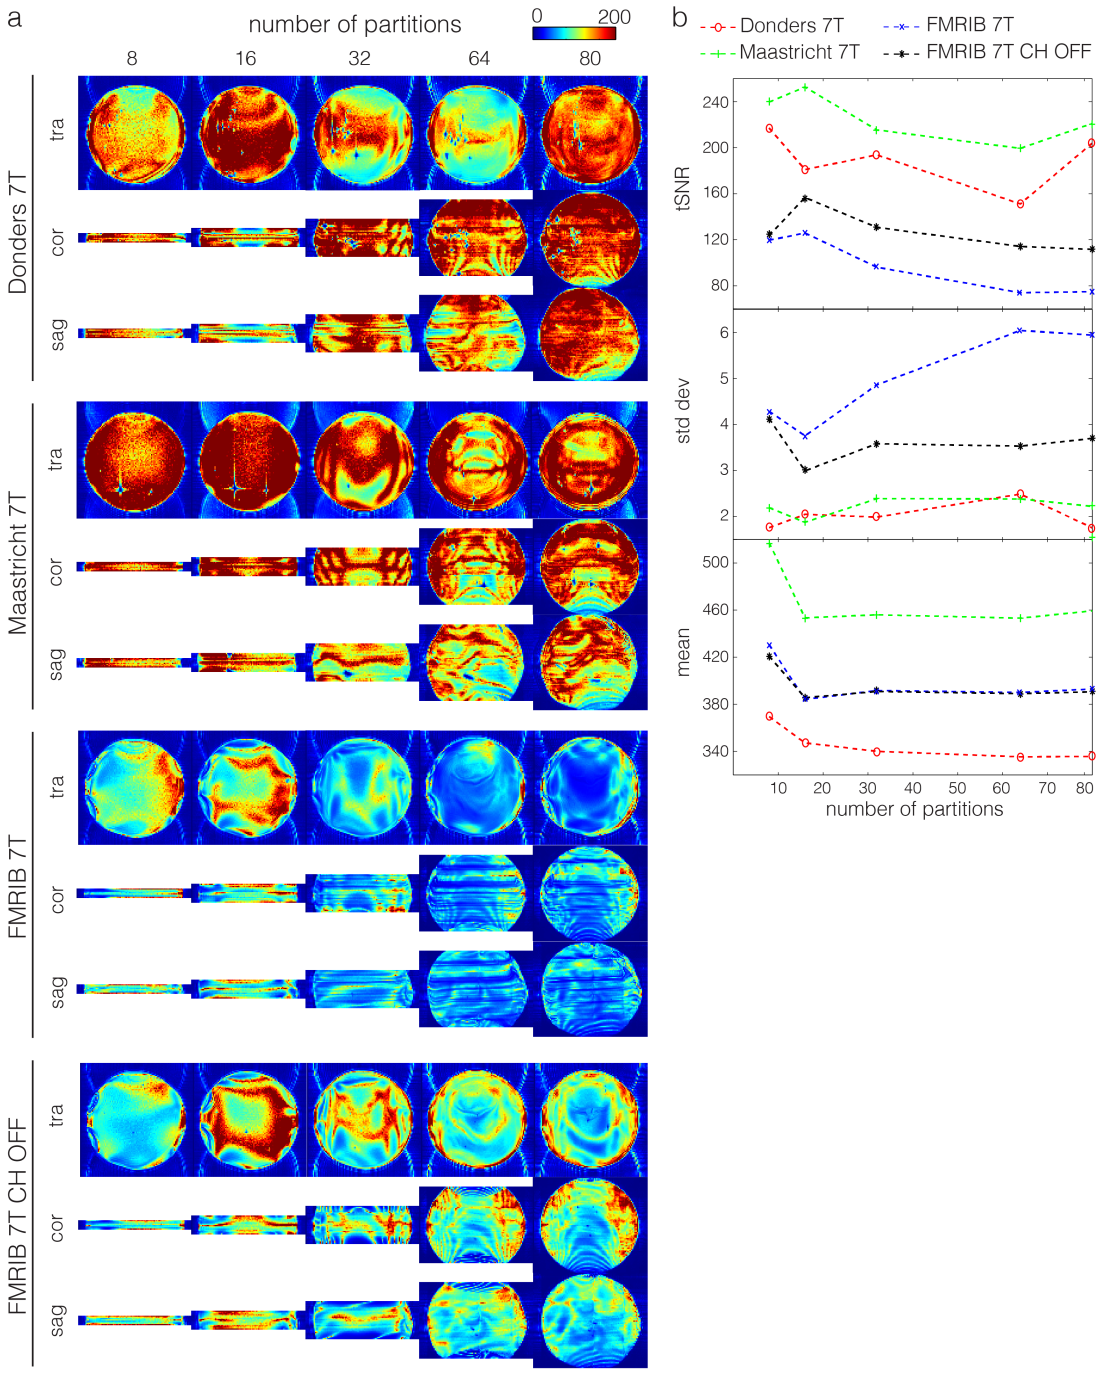


Figure 4.14: Cross-site tSNR comparison: (a) Phantom tSNR maps from three different sites. For the FMRIB measurements a data set acquired with the cold head (CH) turned off is also included. (b) Plots of the average tSNR, temporal standard deviation and temporal mean.

disappeared when the cold head was turned off. The ∼1 Hz fluctuation matches the frequency at which the pump operates and we concluded that the artefacts caused by the helium pump are likely vibration-related artefacts. It would be interesting to test

if temporarily switching off the cryo-cooler has any measurable impact on the quality of the in vivo data as well, but this has not been done to date.

Longitudinal and/or multi-site research studies are becoming increasingly more relevant (for example the recently established UK 7 T consortium). This creates the need to characterize and correct sources of variance related to the hardware, as these could potentially limit the ability to detect subtle functional trends. This experiment showed the challenges of performing a perfectly matched experiment and finding sources of discrepancies.

4.B.4 Discussion

Since these experiments the FMRIB gradient coil was replaced (August 2015) as well as the cold-head (early 2016), but no substantial difference was found when repeating these measurements for the new hardware. We can therefore exclude that these components of the hardware were responsible for the lower tSNR observed at FMRIB.

We concluded that fluid motion is not responsible for the spatio-temporal fluctuation, because it could be reproduced in agar phantoms and because it is unlikely that random fluid motion would produce a spatially consistent temporal fluctuation. We suspected that incomplete spoiling might cause the artefacts, especially because the phantoms used have very long relaxation times (T_2 values, which are more than 10 times larger than typical in vivo values). In the tests we conducted we did not achieve substantial difference by altering the spoiling parameters and experiments with longer TRs displayed similar patterns. However, this should be investigated more thoroughly to completely rule out that the spatio temporal patterns are caused by the build up of stimulated echoes.

To date it is unclear if and how these observations made in phantoms impact our 3D EPI in vivo results. The observed spatial patterns in phantoms were not observed in the human scans. However, the complex shape and different tissue types present in vivo could make it more difficult to detect such spatio-temporal fluctuations, especially if they were more subtle in the subject data.

5

TURBINE Sequence for Functional MRI

Contents

5.1 Introduction 84
5.2 TURBINE Pulse Sequence 85
5.3 Image Reconstruction 90
5.3.1 Gridding and NUFFT 90
5.3.2 Nyquist Ghost Correction 92
5.3.3 In-plane Acceleration Using GRAPPA 95
5.3.4 Radial/3D Acceleration Using SPIRiT 97
5.3.5 SPIRiT Algorithm Parameters 97
5.3.6 Coil Compression and Reconstruction Time Management 103
5.4 tSNR and fMRI Evaluation at 3 T 104
5.4.1 Reconstruction and Analysis 104
5.4.2 Flexible Temporal Resolution for fMRI 105
5.4.3 tSNR and Preliminary Comparison to 2D/3D EPI 106
5.4.4 Preliminary fMRI Comparison to 2D/3D EPI 108
5.5 Preliminary 7 T fMRI 109
5.6 Discussion 111
5.6.1 Image Reconstruction 113
5.6.2 TURBINE Trajectory Properties and Artefacts 114
5.6.3 TURBINE tSNR and fMRI 116
5.7 Conclusion 117

5.1 Introduction

As discussed in chapter 4, 3D EPI is attractive for high-resolution fMRI, due to gains in SNR efficiency as a result of using volumetric excitation and the potential for higher parallel imaging acceleration factors [81, 112, 113, 143]. Cartesian multi-shot 3D EPI [81] is the most commonly used 3D trajectory for fMRI, but other 3D sequences have been proposed including EVI [88, 161], 3D stack of spirals [91] and 3D radial [90] imaging. Radial-based trajectories have a number of benefits, including their lower sensitivity to motion and flow artefacts and frequent acquisition of the k-space centre information. These benefits are interesting for 3D imaging for fMRI, especially the potential for artefact reduction. This is crucial in 3D imaging, where sampling over several seconds can be required to form a complete image.

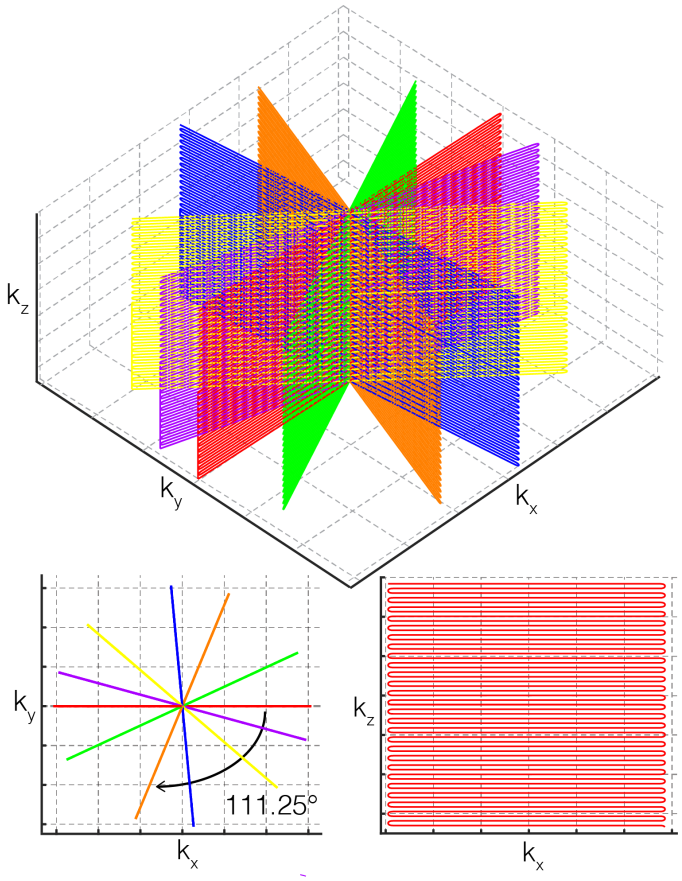


Figure 5.1: Diagram of the 3D hybrid radial-Cartesian sampling scheme: Top schematic shows 3D k-space of 6 shots. Each shot constitutes an EPI readout blade with phase-encoding along k_z and subsequent blades are rotated by approximately 111.25° about the k_z axis. Bottom row shows 2D projections of 6 blades onto k_x-k_y (left) and one example blade onto k_x-k_z (right).

We chose to focus on a novel 3D trajectory using a hybrid radial-Cartesian readout scheme (Fig. 5.1) for functional MRI. The sequence, called **T**rajectory **U**sing **R**adially **B**atched **I**nternal **N**avigator **E**choes (TURBINE), was originally introduced for navigation of diffusion imaging [95]. This combination of EPI with radial imaging has also recently been used for perfusion imaging [162]. Preliminary reports using TURBINE-like trajectories for fMRI have been made by ourselves and other groups, exploring golden angle increments [163–165], use at ultra high-field [166] and in combination with k-t acceleration [166]. This sequence has a number of benefits for fMRI:

1. Using a hybrid of radial and Cartesian sampling enables the potential for motion and physiological noise correction of radial imaging (the centre of k-space is sampled every shot), while maintaining encoding efficiency due to the use of EPI in the third dimension.
2. The radial distribution can be drawn from a golden angle scheme [167]. The golden angle update allows flexible trade-off between spatial and temporal resolution, which can be defined post-acquisition.
3. The increased incoherence compared to purely Cartesian sampling makes this trajectory suitable for combination with k-t acceleration schemes [168, 169].

In this chapter we document the implementation of the sequence, the choice of the reconstruction pipeline as well as typical artefacts seen in TURBINE imaging. We show that the sequence can be used for fMRI at 3 T and 7 T.

5.2 TURBINE Pulse Sequence

The TURBINE pulse sequence is a 3D spoiled gradient echo sequence using a modified EPI readout. TURBINE shares most of its sequence components with the 3D EPI sequence from chapter 4, with modifications to the readout implemented as described below. The original TURBINE pulse sequence for diffusion imaging (written by Jennifer McNab [95]) was available for an older and out-of-use version of the scanner pulse sequence development environment. Given that the sequence needed to be re-implemented, it

made sense to redesign it using the original version as a rough guide, but with new features and extensibility.

Selective excitation is used with the slab direction oriented along the phase-encoding axis, and each TR one EPI readout plane ("blade") is acquired (Fig. 5.1). In contrast to standard multi-shot 3D EPI where readout planes are parallel and offset in a second phase-encoding direction, in TURBINE the EPI blades are intersecting and rotated about the EPI phase-encoding axis. This way, the phase-encoding direction is shared across all shots, but the readout orientation varies from blade-to-blade.

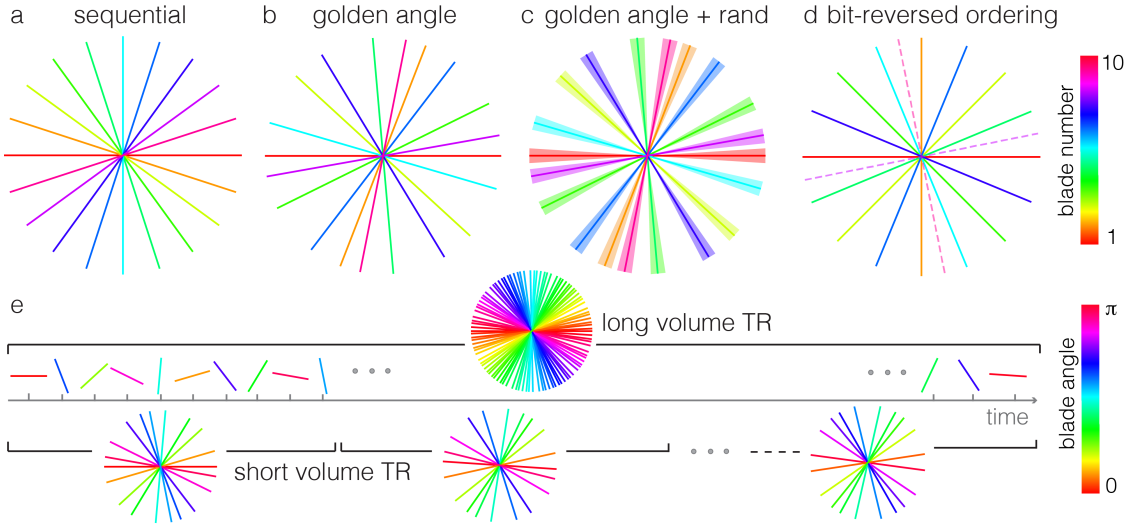


Figure 5.2: Illustration of blade ordering schemes, showing the radial (typically k_x-k_y) projection of TURBINE k-space: (a) sequential ordering; subsequent blades are rotated by $\frac{180°}{N}$, where N corresponds to the number of blades. For this approach the spatio-temporal resolution has to be pre-defined. (b) Golden angle ordering; angle update of 111.25°, used for most of this project. (c) Golden angle ordering with random offset; allows to modulate the spatio-temporal incoherence. (d) Bit-reversed ordering; distribution for 8 blades has uniform blade spacing, blades 9–10 are indicated with dashed lines. (e) Illustration of the flexible spatio-temporal resolution property of golden angle sampling.

The sequence allows arbitrary angle updates between subsequent blades (example blade ordering schemes are illustrated in Fig. 5.2). In the simplest ordering scheme the readout starting positions are evenly distributed from 0 to π, and are acquired in ascending order. For most acquisitions, we chose an update angle based on the golden ratio, where two temporally succeeding blades are rotated by the angle $\frac{180°}{(\sqrt{5}+1)/2} \sim 111.25°$. In this scheme each added blade divides one of the largest azimuthal gaps in the golden

ratio. This ensures near-uniform angular coverage of k-space for an arbitrary number of consecutive blades [167].

Conventionally, 3D k-space is sampled over a pre-defined acquisition window and changing the temporal resolution requires a new scan. For the golden-angle sampling scheme the temporal resolution (how many blades are used to reconstruct a volume) can be defined post-acquisition, therefore no a-priori knowledge of the optimal temporal resolution or the amount of under-sampling tolerated is required. This enables flexible post-acquisition trade-off between temporal and spatial resolution based on the Nyquist criterion, or alternatively variable temporal resolution at a fixed spatial resolution by varying the degree of under-sampling in parallel imaging reconstruction.

If the sequence is used for methods which benefit from increased incoherence in the k-t domain, such as compressed sensing [170, 171] or low-rank matrix [169] completion, the angle update scheme can be modified by adding a random perturbation to each increment [172]. For some acquisitions we set the angle scheme to repeat after a certain number of blades (called "repeating golden angle"), but most scans the angle updates followed the golden angle scheme over the entire time series acquisition (called "continuous golden angle").

Another angle scheme which has been suggested for flexible temporal resolution acquisitions is bit-reversed ordering. This approach maximizes spacing between the blades as long as the number of blades used is a power of 2 and it has been used for radial EPI [90] and spiral imaging [173]. We chose golden angle ordering over bit-reversed ordering because the golden angle is more optimal for arbitrary numbers of blades.

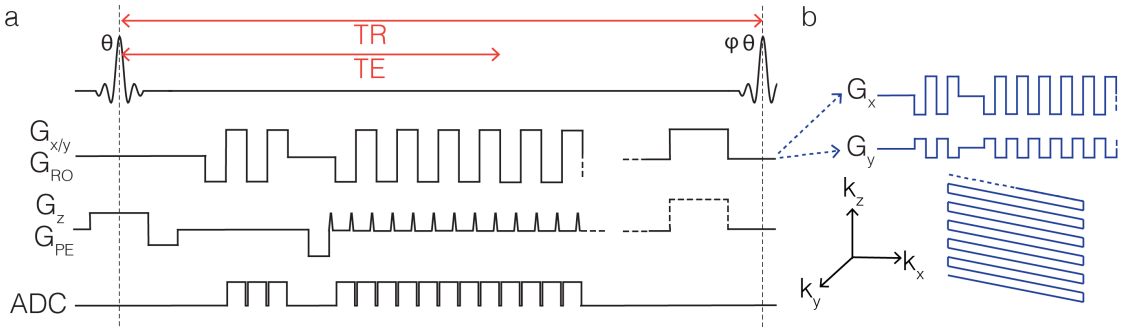


Figure 5.3: TURBINE pulse sequence diagram: (a) Illustration of the TURBINE pulse sequence and (b) gradients waveforms and k-space for an example shot.

A pulse sequence diagram is shown in Figure 5.3. For most acquisitions the blade phase-encoding direction was chosen along the scanner head-foot direction (corresponding to the z-gradient axis), which results in a Cartesian sampling distribution along k_z, and radial distribution in k_x and k_y. It is of course possible to acquire an angulated slab and to have the phase-encoding axis oriented along any oblique direction.

The same sampling pattern could be achieved having the axis of rotation along the readout direction. This would likely cause artefacts due to non-constant distortion, because the phase encoding direction rotates, instead of the readout direction. The trajectory samples a cylindrical k-space volume, similar to a 3D stack-of-stars acquisition, with the difference being that the Cartesian axis is sampled in a single shot. This makes the sequence suitable for fMRI acquisition, due to the long TE required for optimal BOLD contrast and volume TRs on the order of seconds for whole brain imaging. Another advantage of TURBINE over a 3D stack-of-stars trajectory is that TURBINE samples the centre of k-space in every shot.

The sequence has conventional parallel imaging capabilities along the blade phase-encoding direction (skipping k_z lines within the EPI readout) while the radial acceleration factor is defined post-acquisition, depending on the number of shots binned into each 3D k-space. Reference data for the calibration of the parallel imaging reconstruction can be acquired in the same scan prior to the first under-sampled measurement.

Typically, we chose the parallel imaging calibration scan to have a sufficient number of blades to create a fully sampled volume at the target resolution. The calibration scan is segmented in plane, with each blade containing a sufficient number of EPI segments to match the effective phase-encoding bandwidth of the under-sampled readout train. Our initial implementation first acquired one segment for all blades before proceeding to the next segment. This meant that the different segments for one blade were typically acquired more than 10 s apart, which caused inconsistencies between the segments and led to artefacts in the parallel imaging reconstruction. This is consistent with what has been reported in the literature [154, 174]. We changed the sequence to acquire all segments of a given blade in immediate succession, which substantially improved

the in-plane parallel imaging reconstruction. The observed inconsistencies however are still present in the calibration data; we effectively traded within blade inconsistencies for inconsistencies between blades.

Immediately after excitation three non-phase encoded "navigator" readout lines are acquired for Nyquist ghost correction. For efficient sampling, data was acquired in the rising and falling stages of the gradient waveform (ramp sampling) unless otherwise specified.

For excitation we mostly used a binomial pulse, as described in chapter 4. In the final optimized TURBINE protocols a TBWP of 10 or larger was used to ensure a sufficiently sharp slab profile. Generally, the slab thickness was set to <100 % (typically 80-90 %) of the phase-encoding FOV. Encoding a larger volume than was excited avoided contamination of the outer slices with aliased signal.

The readout train was balanced by playing rewinder gradients compensating for the gradient blips played during the EPI readout train. Every shot the residual transverse magnetisation was spoiled using gradient and RF spoiling with a quadratic phase increment of 117° [15]. It has been indicated that other phase increments (for example 45° instead of 117°) can be less sensitive to motion [175] and it would be useful to investigate the impact of spoiling in more detail in future work. The magnitude of the spoiler gradients was chosen to generate 2π phase dispersion across a voxel, as using smaller gradient moments resulted in visible artefacts. Originally, the sequence used gradient spoiling along all three axes. Later, this was changed to only the x- and y-axes as in empirical testing this showed reduced artefacts compared to using all three gradient axes. Note that the spoiling gradients are of equal magnitude for every shot, i.e. the same gradient moment is accrued for every blade, which has been shown to be important to reduce artefacts from incomplete spoiling [176].

5.3 Image Reconstruction

Modifications to existing methods were required to provide a suitable pipeline for the processing of TURBINE data, which is documented in this section. To evaluate the reconstruction quality both phantom (a water-filled structural phantom) and in vivo data sets were used. The data used in section 5.3 were acquired on 3 T Siemens Verio and Prisma systems, using a 32 channel head coil for reception. The details of the acquisition protocols are listed in table 5.1. All the image reconstruction and processing steps described here were performed offline in MATLAB (MathWorks, Natick, Massachusetts, USA). Section 5.3 follows the steps required for reconstruction of TURBINE data and illustrates the flexible temporal sampling as well as typical artefacts. Different reconstruction pipelines were used as the TURBINE project progressed. To simplify the discussion of the different approaches (pipelines A-E) an overview is provided in Figure 5.4.

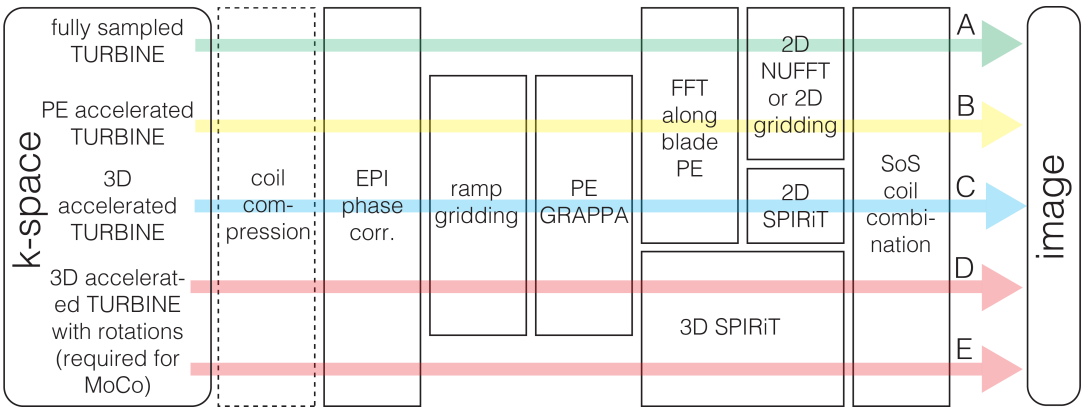


Figure 5.4: Reconstruction pipelines: An overview over the different reconstruction pipelines used throughout the project. Pipeline A was used for the reconstruction of fully sampled TURBINE data, while B was applied to acquisitions accelerated only within the EPI blade. When using radial acceleration pipelines C-E were used. D and E were specifically created to allow motion correction with TURBINE, where sample locations can be arbitrary. In practice pipeline D provided the best results and was used for all reconstructions in chapter 6.

5.3.1 Gridding and NUFFT

The TURBINE sampling scheme is a hybrid radial-Cartesian sampling pattern; i.e. it produces equally spaced samples in one dimension and non-equidistant in the others. Applying a FFT in the Cartesian (usually k_z) direction, transforms the data into a hybrid

Table 5.1: Imaging parameters of all experiments in chapter 5: For the TURBINE protocols the volume TR, matrix, acceleration factor and number of measurements depend on the temporal bin size chosen for the reconstruction. For data sets, which were reconstructed at multiple temporal resolutions this table lists an example value (these values are marked with an asterisk). For the excitation the slab thickness is listed as a percentage of the phase encoding FOV in the slab direction. Abbreviations: res = resolution, iso = isotropic, GA = golden angle, R = acceleration factor, RF = type of RF pulses used, no. meas. = number of measurements, and acq. t. = acquisition time

Description		TE/TR/FA [ms]/[ms]/[°]	TR_{vol} [s]	res [mm^3]	matrix blade order	R	BW $[\frac{Hz}{pixel}]$	RF duration; TBWP; slab th.	no. meas. acq. t. [s]
TURBINE - 3 T Verio Image quality test	Fig. 5.5	45/123/20	18.45	2.3 iso	100 × 100 × 150 GA	none	1852	Regular pulse + FatSat 2.56 ms; TBWP 5.2; 50 %	1 18.45
TURBINE - 3 T Prisma Phantom image test	Fig. 5.6, 5.7	30/65/15	16.64*	1.5 iso	134 × 108 × 256* GA	2 × 1*	1244	WE 3.84 ms; TBWP 10; 90 %	1* 16.64
TURBINE - 3 T Prisma fMRI + image qual.	Fig. 5.6, 5.8, 5.11, 5.12	29/50/15	2.5*	2 iso	100 × 76 × 50* GA (continuous)	2 × 3*	1852	WE 3.84 ms; TBWP 10; 90 %	120* 300
TURBINE - 3 T Verio fMRI	Fig. 5.10	25/56/15	7.17*	2 iso	100 × 74 × 128* GA (repeating)	2 × 1.2*	1724	WE 2.56 ms; TBWP 5.2; 80 %	40* 300
TURBINE - 3 T Verio fMRI	Fig. 5.11, 5.12, 5.13	25/50/15	2.5	2 iso	100 × 76 × 50 GA (repeating)	2 × 3	1724	WE 2.56 ms; TBWP 5.2; 80 %	120 300
3D EPI - 3 T Verio fMRI	Fig. 5.12, 5.13	25/50/15	3.6	2 iso	100 × 100 × 72	2 × 1	1724	WE 2.56 ms; TBWP 5.2 %	84 302
2D EPI - 3 T Verio fMRI	Fig. 5.12, 5.13	25/4590/90	4.59	2 iso	100 × 100 × 76	2	1666	Regular pulse + FatSat 2.56 ms; TBWP 5.2; 100 %	66 303
TURBINE - 7 T Image quality test	Fig. 5.14	20/40/10	4.0	2 iso	100 × 50 × 100 GA (repeating)	1 × 1.6	1924	WE 2.56 ms; TBWP 12; 90 %	1 8
TURBINE - 7 T Image quality test	Fig. 5.14	21/44/12	4.4	1.5 iso	133 × 102 × 100 GA (repeating)	2 × 2	1776	WE 2.56 ms; TBWP 25; 90 %	1 8.8
TURBINE - 7 T fMRI + image quality test	Fig. 5.14, 5.15	20/42/12	2.69	1.25 iso	160 × 83 × 64 GA (continuous)	3 × 4	1144	WE 2.56 ms; TBWP 25; 90 %	112 300
TURBINE - 7 T fMRI + image quality test	Fig. 5.14, 5.15	22/47/13	6.02	1 iso	200 × 100 × 128 GA (continuous)	3 × 2.5	1190	WE 2.56 ms; TBWP 25; 90 %	50 300

(k_x ,k_y ,z) space, turning the reconstruction from a general 3D non-Cartesian problem into a stack of independent 2D non-Cartesian problems. Initially, we implemented 2D gridding (see chapter 2) for the non-Cartesian reconstruction. Later, this step was replaced by the equivalent, but more efficiently implemented NUFFT, using the Fessler image reconstruction toolbox [22]. Both of these methods require compensation for irregular sampling density. For regularly spaced blades, for example when using the sequential blade ordering scheme (with no or corrected ramp-sampling), density pre-compensation was performed using the analytic density compensation function for radial imaging (Ram-Lak filter [177]). In all other cases, for example when using the golden angle update, the sampling density was estimated by computing a 2D Voronoi diagram [178] (using the built-in Matlab function) and using the size of the individual Voronoi cells as a measure of the local sample density. For all reconstructions the final magnitude images were generated as a sum-of-squares combination of the individual reconstructed coil images.

5.3.2 Nyquist Ghost Correction

Gradient induced eddy currents and other hardware imperfections can cause trajectory errors. In EPI this typically causes inconsistencies between the forward and reverse lines, which results in the well-known Nyquist ghost. To correct for this artefact we performed the standard one dimensional EPI Nyquist ghost correction [179] on every blade using three non-phase encoded navigator k-space lines acquired immediately after excitation for every readout.

The primary difference between TURBINE and conventional EPI is that the readout direction is shot-dependent. As described in [162], EPI phase correction serves two purposes in TURBINE: Firstly, minimizing the Nyquist ghost by removing inconsistencies between odd and even lines in the same readout. Secondly, reducing constant and linear image phase differences between different blades, which serves as a basic k-space trajectory correction (correcting only the constant parallel component of the trajectory error) [180, 181].

The k-space shift was estimated in image-space by fitting a linear phase ramp to the phase difference between the inverse Fourier transformed forward and reverse navigator lines (using the high SNR centre of navigator lines). The resulting constant ϕ_0 and linear ϕ_1 phase terms were averaged over all coils (discarding outliers) and used to correct all the acquired lines (by extrapolating over the whole field of view and removing the phase ramp from all the individual coil images in image space). In order for the k-space lines to meet in the centre of the TURBINE cylinder, the forward and reverse lines were shifted to meet in the centre by applying half the phase ramp to the forward lines and the negative of that to the reverse lines. The resulting projections were Fourier transformed to yield the corrected k-space lines.

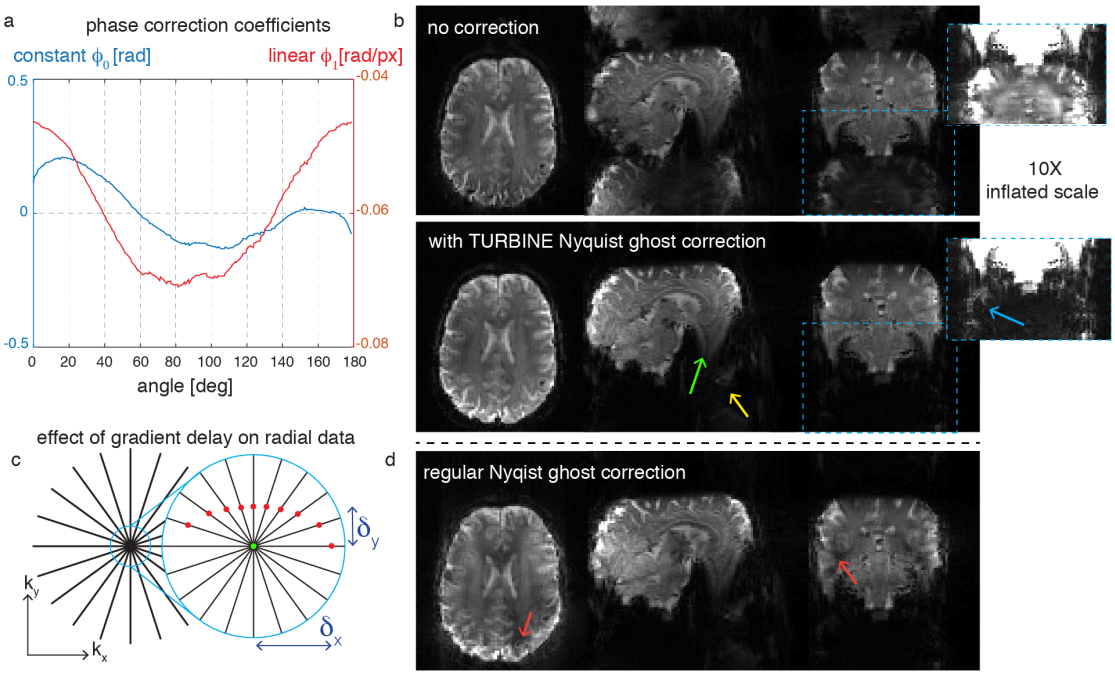


Figure 5.5: EPI phase correction: (a) Measured constant and linear phase offsets (in image-space) for a 3 T TURBINE acquisition. (b) The image before (top) and after (bottom) correction using the phase coefficients in (a). (c) Illustration of the effect of gradient delays on radial acquisitions. The true k-space centre is displayed in green, the actual ones in red. (d) Image reconstructed using standard Nyquist ghost correction without the adaptation to TURBINE data. The arrows indicate different artefacts: distortion (green), aliasing (yellow), residual Nyquist ghost (blue) and dropout due to phase cancellations (red).

This phase correction was tested on an in vivo 3 T TURBINE data set using no parallel imaging (TE/TR = 45/123 ms, 2.3 mm isotropic resolution). 150 blades were acquired with the golden angle scheme and the reconstruction was performed according to pipeline A

(using gridding). The measured phase terms ϕ_0 and ϕ_1 are displayed as a function of the blade angle in Fig. 5.5a. Fig. 5.5c illustrates how for radial sampling, a gradient delay causes the spoke-centre to not fall in the centre of k-space. When using a combination of several gradients, we expect the resulting k-space shift to be a linear combination of the individual gradient terms. For radial MRI this should look like a sinusoidal function, which was observed in our data (Fig. 5.5a).

As shown in Fig. 5.5b the phase correction removes most of the Nyquist ghost artefact (on a normal scale no residual ghost was visible in the images). A small amount of artefact was still present in the corrected data (see blue arrow in the inset image with scale inflated to 10 %). Figure 5.5d was included to highlight the importance of centring the EPI blades in the TURINE cylinder. In this reconstruction the odd and even lines were aligned shifting the odd lines only. While this removed most of the Nyquist ghost, it introduced more severe radial trajectory mismatch, which can be clearly seen by the increase in image artefact (see red arrows in Fig. 5.5d). The original image in Figure 5.5a also contains some of these artefacts, which are reduced in the corrected image.

Being one of the first TURBINE acquisitions, this scan also brought to light some other important considerations when acquiring TURBINE data. The volume had not been adequately shimmed (resulting in field inhomogeneity in the z-direction), which caused substantial distortion (see green arrow in sagittal view in Fig. 5.5b). We can also see aliased signal from the neck in the lower part of the image (yellow arrow). For most future acquisitions a TBWP of 10 or larger (compared to 5.2 used in this acquisition) was used to improve the slab profile and every acquisition was carefully shimmed.

The EPI phase correction we used relies on the assumption that the navigator lines lie perfectly on top of each other (i.e. we are neglecting any trajectory shifts perpendicular to the blades) and that the odd-even offset is constant along the readout. Gradient imperfections beyond a simple delay and B_0 inhomogeneity can both violate these assumptions. These effects likely explain the slight asymmetry of the constant phase term, which we observed (Fig. 5.5a). This effect was reproducible for different acquisitions on

the same scanner (both in vivo and phantom experiments as well as using different blade ordering schemes), which indicates that it is likely due to reproducible gradient effects, potentially gradient cross terms or non-linear gradient responses.

The ghosting and other artefacts could be further reduced by using one of the existing more advanced methods for EPI phase correction; most of which rely on the collection of 2D reference data [182]. Improved trajectory correction without any additional calibration or acquisition could also be performed using the GROG [183] method.

On the 7 T scanner we performed trajectory monitoring and gradient system characterization [40] using a dynamic field camera [184, 185]. Preliminary results are presented in chapter 7 and in future work we hope to integrate information from the gradient coil characterization into the image reconstruction to further reduce artefacts related to trajectory imperfections [42].

5.3.3 In-plane Acceleration Using GRAPPA

In-plane parallel imaging (with in-plane referring to within the EPI blade) is required in order to achieve whole brain coverage with echo times for optimal BOLD contrast. The increased effective bandwidth also reduces distortion. We used GRAPPA [27] to fill in the missing k-space locations. The GRAPPA kernels were calibrated separately for each blade (using a kernel size of [4 3] in [RO PE]). For the data presented here the closest blade in the calibration data set was used to calibrate the GRAPPA kernel. It is also possible to generate data reference for each blade angle by inverse gridding the reconstructed calibration data set. Ramp sampling was used on most TURBINE acquisitions, so readout lines were first gridded to ensure the equidistant samples required for GRAPPA reconstruction. Figure 5.6 shows two example images acquired with the TURBINE sequence, reconstructed from k-space data, which were fully sampled in the radial direction and accelerated by a factor 2 in the EPI plane. The distortion is reduced compared to previous acquisitions without parallel imaging (cf. Fig. 5.5), but a small amount of distortion can still be seen in the frontal lobe in the sagittal view of the brain.

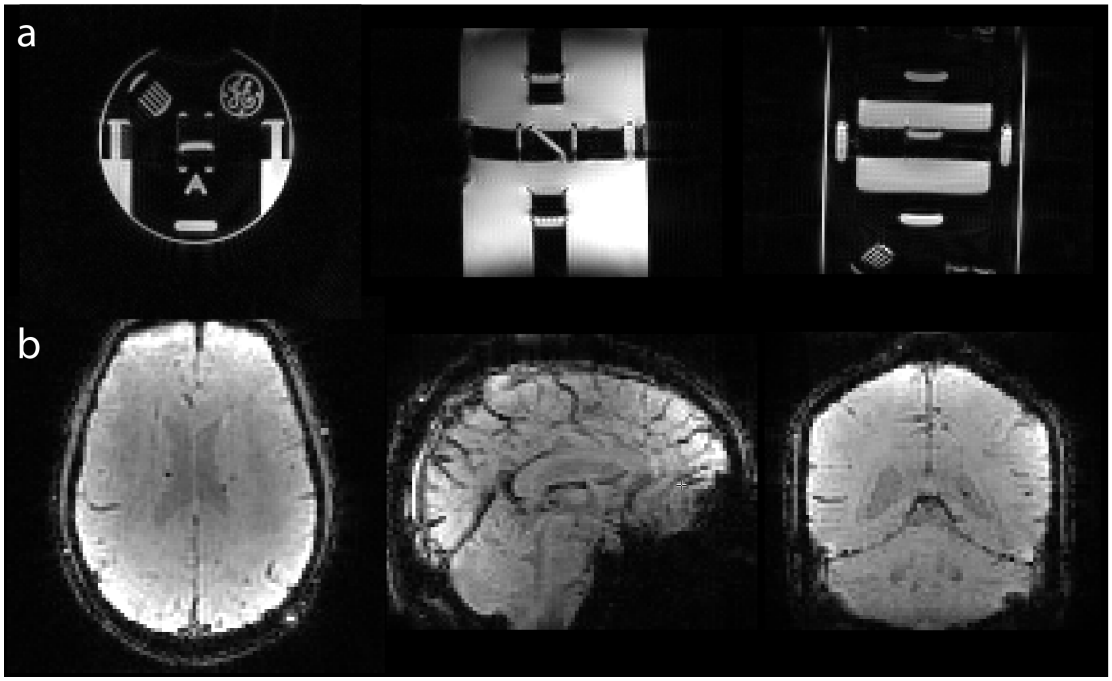


Figure 5.6: TURBINE images: Acquisition of (a) a structure phantom (1.5 mm isotropic resolution) and (b) in vivo data (2 mm isotropic resolution) acquired with TURBINE at 3 T. An in-plane acceleration factor of 2 was used and no radial acceleration.

To satisfy the Nyquist criterion in the radial direction at the edge of k-space the number of blades must be by a factor $\frac{\pi}{2}$ larger compared to the number of EPI planes in an equivalent Cartesian acquisition (assuming uniform angles; for the golden angle scheme the gaps can be slightly larger). The total acceleration factor is always stated with respect to a fully sampled radial volume as $R_{tot} \sim R_r \cdot R_{\mathrm{EPI}}$, where R_r corresponds to the radial undersampling factor and R_{EPI} to the amount of in-plane acceleration.

For the in vivo 2 mm data set in Fig. 5.6 157 blades are required to meet the Nyquist criterion, which results in a volume TR of almost 8 seconds (an equivalent Cartesian acquisition would take $\sim$ 5 s). For a functional MRI acquisition, where the volume TR for a typical task fMRI experiment should ideally not exceed a few seconds and faster sampling provides statistical advantages, we need to further accelerate the acquisition. This can be done by using parallel imaging in the radial plane additional to within-blade parallel imaging.

5.3.4 Radial/3D Acceleration Using SPIRiT

We used the SPIRiT [178] parallel imaging algorithm for non-Cartesian parallel imaging. Details on the choice of the reconstruction parameters are provided in section 5.3.5 below.

Figures 5.7 and 5.8 illustrate the trade-off between temporal resolution and image quality offered by the radial golden angle sampling scheme. The under-sampled zero-filled reconstructions using the 3D NUFFT adjoint operator (i.e. not using any parallel imaging) for different numbers of blades contain aliasing artefacts characteristic of radial imaging (Figs. 5.7b and 5.8b). Radial trajectories naturally tolerate higher amounts of under-sampling due to the more diffuse radial aliasing, therefore reconstructions at fairly high under-sampling factors appear reasonable (see for example NUFFT-reconstructed $R_{tot} \sim 6$ image in Fig. 5.8b). However, larger uncovered regions of k-space cause visible streaking artefacts, making it essential to perform parallel imaging to fill in the missing k-space information. Reconstructing using the 3D SPIRiT (pipeline D) parallel imaging algorithm (Figs. 5.7c/e and 5.8c) reduced the aliasing artefacts, with no visible artefacts up to total acceleration factors of $R_{tot} \sim 6$. For higher acceleration factors, an increase in image artefacts and noise can be observed. At very large acceleration factors detailed structure is lost (see zoomed plot in Fig. 5.7b), but the overall shape of the object or brain can still be reliably determined (see Fig. 5.7c). This is important for the motion correction using high temporal resolution navigators presented in chapter 6.

5.3.5 SPIRiT Algorithm Parameters

The early TURBINE reconstructions performed a FFT along the phase-encoding direction (pipelines A-C) and used 2D SPIRiT for the reconstruction of the under-sampled radial planes. This was done using freely available 2D SPIRiT reconstruction software [32], which uses the above mentioned NUFFT toolbox [22] to accommodate arbitrary k-space sampling locations. The SPIRiT code is implemented in Matlab and was downloaded from `https://people.eecs.berkeley.edu/~mlustig/Software.html` (September, 2013). We used

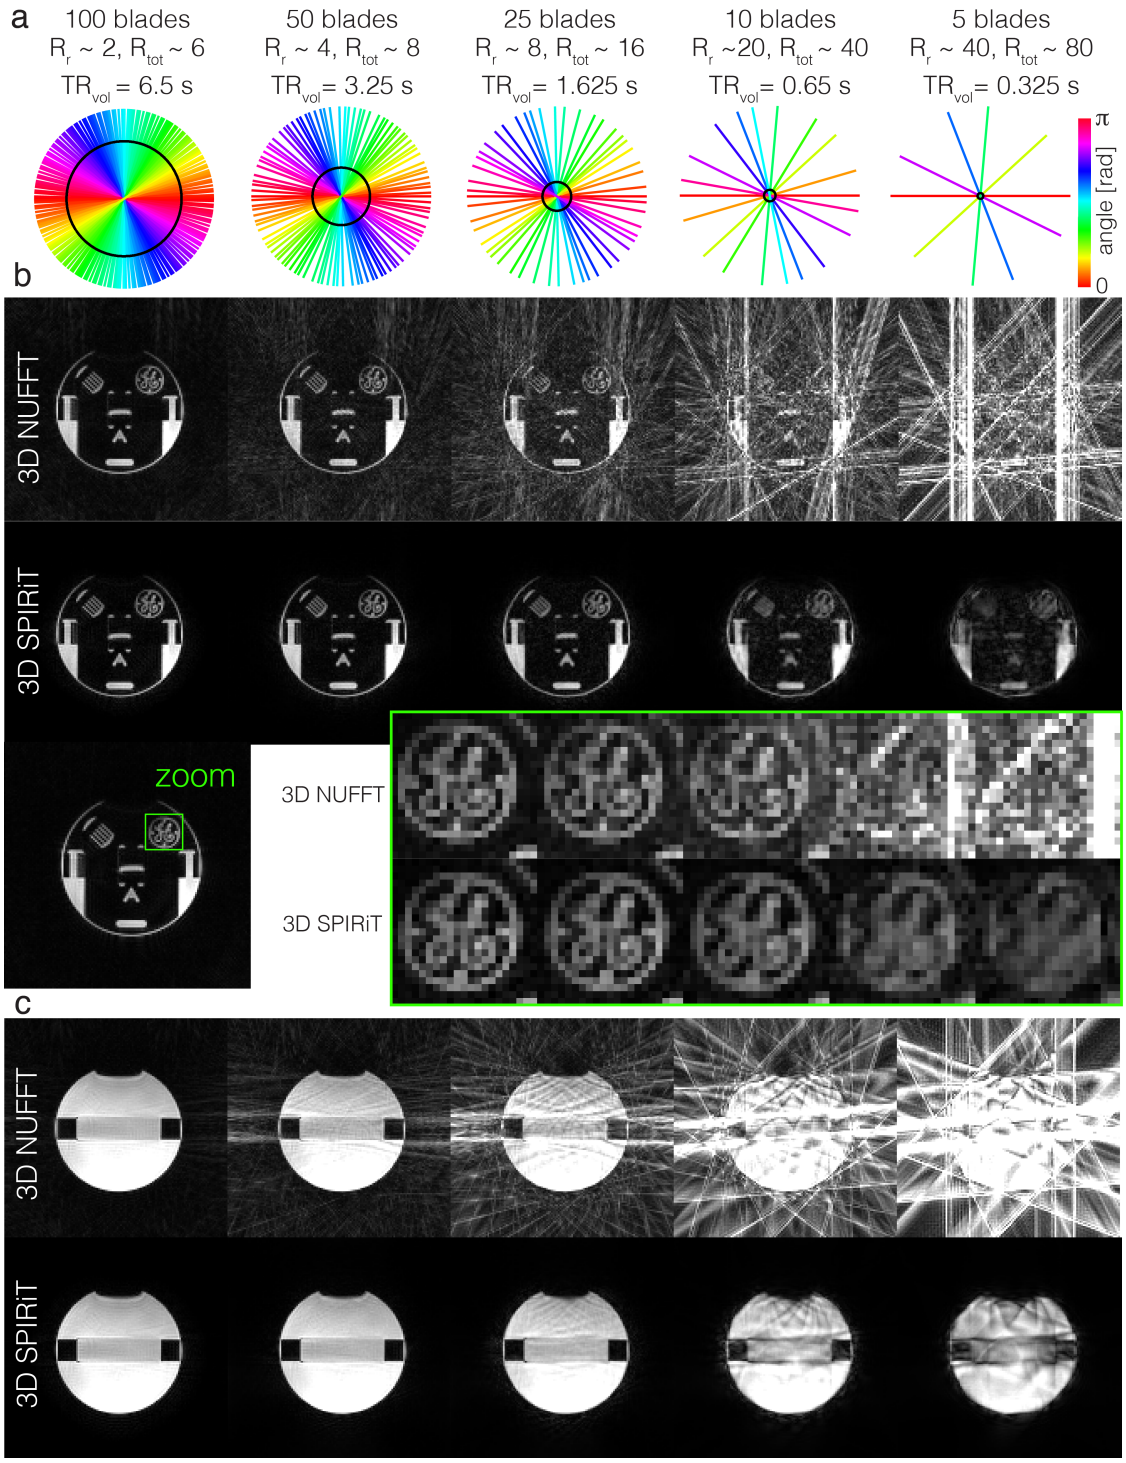


Figure 5.7: Flexible temporal resolution of hybrid radial-Cartesian EPI: (a) K-space plot of blades used in the reconstructions below with the volume acquisition time and the corresponding under-sampling factor relative to Nyquist criterion; the black circle indicates the point in k-space at which the Nyquist criterion is still met. (b) Zero-filled 3D NUFFT reconstruction, with aliasing artefacts introduced due to under-sampling and reconstruction using the 3D SPIRiT parallel imaging algorithm. The zoomed inset below shows detail of the structure phantom. (c) An inferior slice illustrates how the shape of the object can still be reliably identified in the SPIRiT reconstruction at very high acceleration factors.

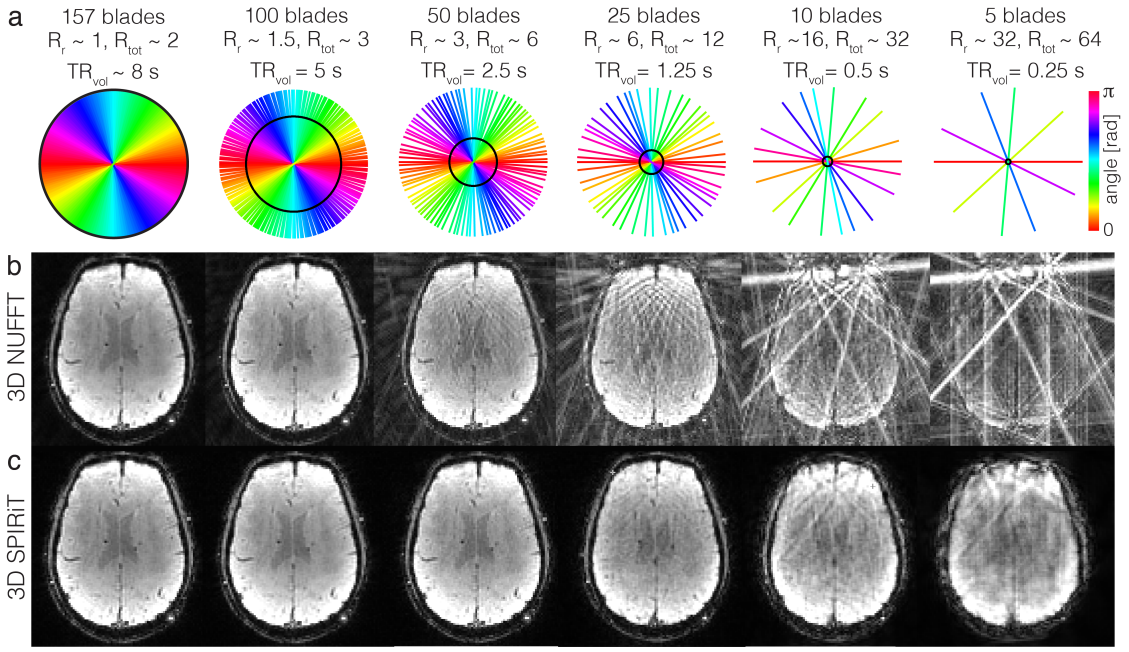


Figure 5.8: Flexible temporal resolution of hybrid radial-Cartesian EPI: (a) K-space plot of blades used in the reconstructions below with the volume acquisition time and the corresponding under-sampling factor relative to Nyquist criterion; the black circle indicates the point in k-space at which the Nyquist criterion is still met. (b) Zero-filled 3D NUFFT reconstruction of in vivo data, with aliasing artefacts introduced due to under-sampling and (c) reconstruction using the 3D SPIRiT parallel imaging algorithm.

the image-based variant of the SPIRiT algorithm using a conjugate gradient method to solve the least squares problem.

The general reconstruction pipeline was set up for the motion correction study presented in chapter 6. As rotations in image space correspond to rotations in k-space, the samples are no longer hybrid radial-Cartesian after head rotation has occurred. For the purpose of motion correction a 3D SPIRiT reconstruction was implemented by generalizing the 2D SPIRiT code to 3D. A number of optimisations had to be made because of the extra memory and computational burden of 3D versus 2D reconstruction.

Iterative reconstruction methods like SPIRiT do not require explicit density compensation, as they converge towards the image that best matches the acquired data. The SPIRiT algorithm converges faster when used with density pre-compensation; however, especially for highly under-sampled data sets, we achieved better results when using no density compensation, especially in regions with high static field inhomogeneity. Large field offsets,

invalidating the density weights, could be responsible for the residual artefacts when using density compensation in the reconstruction. Therefore, spatially varying density compensation (using "local k-space formalism" [186]) would be required to improve the quality of the image reconstruction [187, 188], but further investigation would be required to confirm this. The results for moderate under-sampling factors were comparable with and without density compensation as long as an adequate number of iterations were used for both approaches. Another advantage of not using density compensation is that we do not have to worry about accurately calculating the density weights (which in the general 3D case can be challenging and would require the implementation of a 3D Voronoi or equivalent method). We ran all the final reconstructions in chapter 6 without using density compensation. The consequence of this approach is the increased number of iterations and thus longer reconstruction time.

There are a number of parameters in the SPIRiT algorithm which can be tuned to achieve the best possible reconstruction (listed in table 5.1). The most relevant parameters are the weight between the data consistency and the kernel consistency λ, the SPIRiT kernel size and the number of conjugate gradient iterations. The SPIRiT reconstruction also includes several steps using the NUFFT, for which the type of kernel and the oversampling factor need to be chosen. We chose most of the parameters close to the recommended values in the original SPIRiT paper [32], but performed additional empirical testing to verify that these choices also provide good results for the 3D SPIRiT reconstructions.

The 3D SPIRiT kernel calibration was performed on the fully-sampled calibration data. Analogous to GRAPPA a larger SPIRiT kernel can capture more complex spatial variation in the coil sensitivities, therefore theoretically a larger kernel should always be better (with GRAPPA becoming equivalent to SENSE when kernel size = matrix size). However, for large kernels the reconstruction problem becomes ill-conditioned [189]. We found a kernel size of between [3,3] and [7,7] in 2D and [3,3,3] and [7,7,7] in 3D to provide good results. Larger kernels increase the memory and processing demand of the reconstructions; therefore we chose a kernel size of [5,5] or [5,5,5] for our final reconstructions. The kernel calibration was performed on central k-space region of [20,20] or [20,20,20] size, which was

Table 5.2: Overview of the SPIRiT reconstruction parameters for 2D and 3D SPIRiT with and without density compensation. We also included the parameters optimised for a highly undersampled data set. The parameters in the two last columns are used for reconstruction in chapter 6.

	2D SPIRiT		3D SPIRiT		3D SPIRiT highly undersampled
Density compensation	ON	OFF	ON	OFF	OFF
λ: Trade-off between data consistency and calibration consistency	1	3	1	2	5
SPIRiT kernel size	[5,5]	[5,5]	[5,5,5]	[5,5,5]	[5,5,5]
Size of calibration region	[20,20]	[20,20]	[20,20,20]	[20,20,20]	[20,20,20]
Tykhonov regularization parameter for kernel calibration	0.01	0.01	0.01	0.02	0.05
Number of CG reconstruction iterations	15	100	15	60	50
NUFFT kernel size	[6,6]	[6,6]	[6,6,6]	[6,6,6]	[4,4,4]
NUFFT oversampling factor	1.5	1.5	1.5	1.5	1.2
NUFFT kernel	kaiser	kaiser	kaiser	kaiser	kaiser

resampled from the calibration data onto a Cartesian grid (using the NUFFT). Tikhonov regularisation was used to stabilise the kernel calibration. We found regularisation parameters between 0.01 and 0.05 to provide good results. However, the amount of regularisation required depends on the amount of noise in the calibration data, lower noise requiring stronger regularisation [189]. Ideally the regularisation should be optimised for every data set. For simplicity, however, we chose the same value for all of our final reconstructions (including the 7 T data, but it should be noted that these results are preliminary and we plan to optimise the reconstruction parameters separately for the 7 T data in future work).

As described in chapter 2, the parameter λ in the SPIRiT reconstruction controls the trade-off between data consistency and kernel consistency. For reconstructions using density compensation we used $\lambda = 1$. When using no density compensations it is beneficial to choose a larger λ, i.e. putting more weight on the kernel consistency; we used $\lambda = 2$ for most of the final reconstructions (see table 5.1).

The SPIRiT algorithm exhibits a "semi-convergence" behaviour, which has an initial convergence phase toward the exact solution and later divergence. This is documented in the original SPIRiT paper [32] and has been shown to occur for other iterative

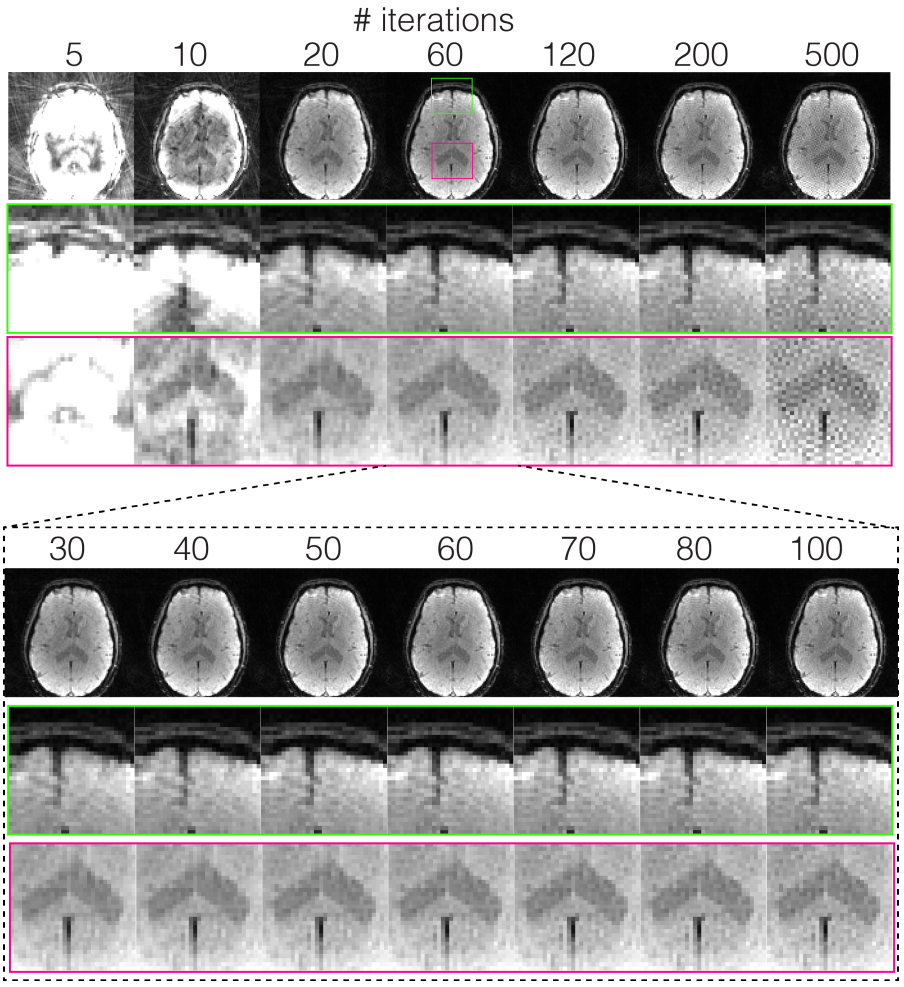


Figure 5.9: Images reconstructed using different numbers of CG SPIRiT iterations illustrating the importance of the iteration number as a regularization parameter (controlling the trade-off between residual artefact and SNR).

reconstruction strategies as well [190]. The semi-convergence behaviour is caused by the ill-conditioning of the encoding matrix of the reconstruction problem and the intrinsic regularization property of the conjugate gradient iterations. The iteration count therefore acts as a regularization parameter and controls the compromise between SNR and residual aliasing artefact: with very few iterations the reconstructed images are dominated by aliasing artefacts (for the TURBINE trajectory this consists of typical radial streaking artefacts in-plane), when using many iterations the images become noisier. We demonstrate this for an example data set in Fig. 5.9. There is a transition region, which presents a compromise between artefact level and noise level. Because the conditioning of the encoding matrix depends on the under-sampling factor (worse

for more under-sampling), the transition zone between artefact and noise dominated regime occurs at a different point depending on the amount of under-sampling [191]. In this case we chose an iteration number based on comparing a range of reconstructions for a typical data set and target under-sampling factor. For very under-sampled data sets we chose a different parameter set (see table 5.1), but for future work it would be useful to implement an automated stopping criterion.

We applied GRAPPA on each blade prior to reconstructing the entire volume with 3D SPIRiT (pipeline D), because this removed aliasing (in the blade phase encoding direction) more robustly than 3D SPIRiT alone (pipeline E). This might be due to inconsistencies in the reference data or trajectory imperfection, which lead to better reconstruction kernels for a separate GRAPPA step when estimated over each blade independently. However, a more thorough investigation is required to make a definitive conclusion.

5.3.6 Coil Compression and Reconstruction Time Management

When moving from direct to iterative reconstruction schemes, reconstruction times can become prohibitively long. Therefore coil compression, which exploits redundancy in coil sensitivity, was used to accelerate the image reconstruction by reducing the dimensionality of the problem. We performed a singular value decomposition [192] and selected virtual coils explaining $\geq$90 % of the variance. This yielded between 8 and 15 virtual coils depending on the data (subject or phantom) and resolution, with in vivo data and higher-resolution data requiring more coils. The coil compression matrix was determined on the reference data and applied to the entire time series. The time required for coil compression is negligible compared to the total reconstruction time, and it sped up the reconstruction with no noticeable degradation in quality compared to reconstructions with the full 32 channels.

One of the challenges of the current 3D SPIRiT implementation is the large memory usage and long processing time, which can become limiting for ultra-high resolution data sets (such as the high-resolution 7 T data sets in section 5.5, which can take several

hours). As mentioned above, the 3D reconstruction can be decoupled into independent 2D reconstructions for the regular TURBINE data set (when no motion correction is performed) and reconstruction can be sped up by using density compensation. It has been shown that 3D SPIRiT (or the related ESPIRiT) algorithm can be accelerated using a scalable parallel implementation to allow reconstruction within minutes [193]. For the proof-of-principle TURBINE studies performed so far reconstruction time was not a priority but for future work we plan to test this more efficient implementation (included in the BART toolbox [194]) or explore other options.

5.4 tSNR and fMRI Evaluation at 3 T

The main goal of the following sections is to test the use of TURBINE for functional MRI, demonstrating its flexible temporal resolution properties. To evaluate TURBINE tSNR and functional sensitivity, we acquired several sets of TURBINE fMRI data using either a visual stimulus or a visual stimulus with simultaneous motor task. The visual stimulus consisted of a 30 s OFF/ON 8 Hz flashing checkerboard stimulus. When combined with a motor task the subjects were asked to perform bilateral finger-tapping during the ON blocks. All data sets were acquired using TURBINE with the continuous or the repeating golden angle scheme. For two subjects, matched acquisitions (in coverage, TE/TR and in-plane acceleration) using 2D and 3D EPI were acquired in the same session, with acquisition parameters detailed in table 5.1.

5.4.1 Reconstruction and Analysis

The data set used for the demonstration of flexible temporal resolution (Fig. 5.10) was reconstructed using 2D SPIRiT (pipeline B), all other data sets used the standard 3D SPIRiT pipeline (pipeline D). Task fMRI analysis was performed in FEAT using no spatial smoothing, temporal filtering, or pre-whitening of the voxel time-series. To account for residual autocorrelation the resulting z-statistic maps were corrected

using mixture modeling [195] to ensure correct null distributions with zero mean and unity standard deviation.

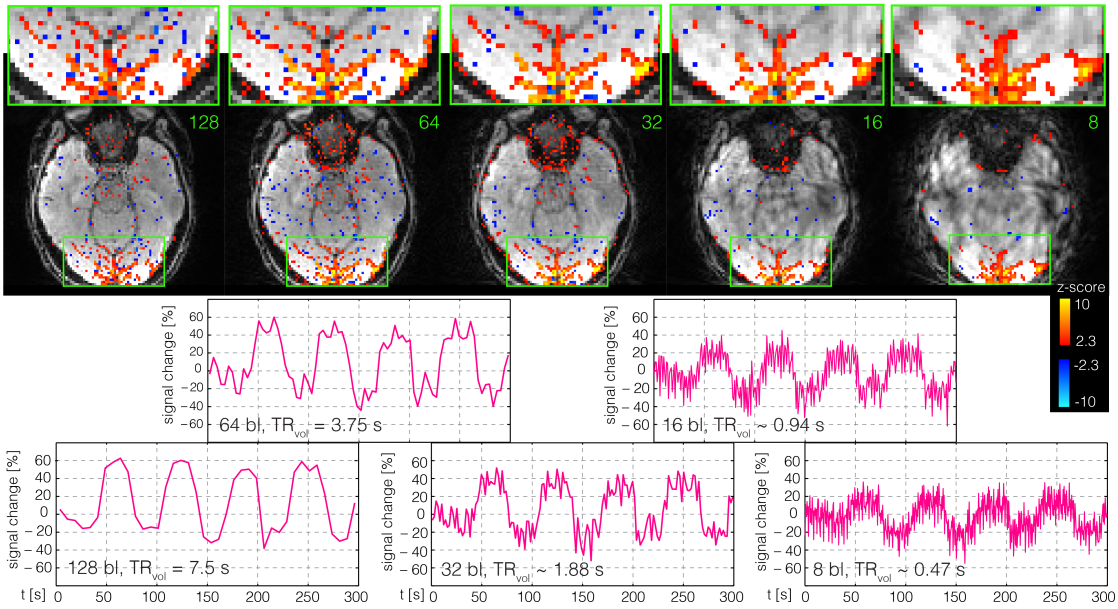


Figure 5.10: fMRI with flexible temporal resolutions. Activation maps and time courses for five different temporal resolutions.

5.4.2 Flexible Temporal Resolution for fMRI

A visual task fMRI data set with 2 mm isotropic resolution was reconstructed using 128, 64, 32, 18 and 8 blades per volume, resulting in TR_{vol} between 7.17 and 0.45 s (TR = 56 ms). Activation maps overlaid on example TURBINE images for the different temporal resolutions are shown in Figure 5.10. The spatial resolution of the images and the specificity of the z-statistic maps deteriorate as the number of blades is reduced, due to increased artefacts and the impact of regularization in the reconstruction. The deterioration in resolution was quantified using a spatial smoothness metric calculated as part of the FEAT analysis, using the FSL "smoothest" tool, which reports spatial smoothness by estimating the size of Gaussian smoothing kernel needed to filter perfectly white noise to match the data residual smoothness (see table 5.3).

With increasing temporal resolution, the temporal response to the stimulus can be more accurately sampled (compare for example the 64 blade time course showing to the 128

Table 5.3: Estimate of spatial resolutions for different reconstructions in Fig. 5.10. The smoothness parameter in unit voxel volume (top row) was obtained using the FSL "smoothest" tool, the same parameter is also expressed below in mm^3 and as approximate 1D smoothness (assuming isotropic smoothing) for easier interpretation.

number of blades	128	64	32	16	8
smoothness parameter (3D) [unit voxel volume]	0.725	4.07	4.83	5.05	12.88
smoothness parameter (3D) [mm^3]	5.8	37.6	38.64	40.4	103.04
1D approx. smoothness [mm]	1.8	3.2	3.4	3.44	4.68

blade time course); however, the noise in the time courses increases. This is for example visible in the 8 blade reconstruction, which has a high-frequency "jitter" in the time course. This behaviour is likely driven largely by increased time-varying aliasing, although it is impossible to differentiate actual signal variance from artefact without a ground truth. One draw back of the golden angle scheme is therefore the introduction of temporally variable aliasing. The main advantage is that the flexible temporal resolution property enables a post-acquisition choice of how much under-sampling is tolerable, and what the best temporal resolution is to analyse a given fMRI experiment.

5.4.3 tSNR and Preliminary Comparison to 2D/3D EPI

In order to get an estimate of the average TURBINE tSNR, we analysed the tSNR values from the task fMRI data sets collected on seven subjects over the course of the project. Three subjects were acquired on the 3 T Verio and 4 subjects acquired on the 3 T Prisma scanner. The parameters used on the two scanners were nearly identical; the main differences were a change in gradient spoiling (from using all gradient axes to only x/y) and the use of the continuous golden angle scheme (versus the golden angle acquisition resetting every 256 blades for the Verio protocol). For two of the subjects acquired on the Verio scanner, we also collected 2D and 3D EPI fMRI data in the same session. Note that this experiment was not set up as a systematic tSNR comparison between TURBINE, 3D and 2D EPI. It is included here to illustrate potential challenges of TURBINE and an extensive comparison is planned for future work.

The tSNR was calculated as the temporal mean image divided by the standard deviation of the residuals of the GLM fit performed from FEAT. Figure 5.11 shows tSNR histograms

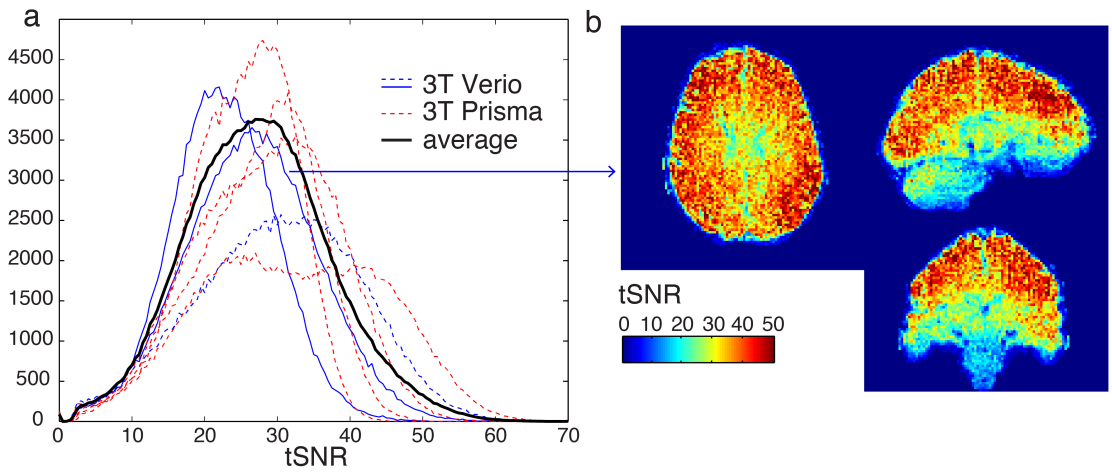


Figure 5.11: Temporal SNR of TURBINE: (a) Histograms of the tSNR of TURBINE. The data sets acquired on the 3 T Verio are shown in blue, the 3 T Prisma data sets in red. The solid lines indicate the two data sets for which comparison 2D/3D EPI data sets are available (see Fig. 5.12). The average tSNR (solid black line) over all subjects is ∼27. (b) tSNR maps for one example 3 T Verio subject.

for all subjects (blue lines corresponding to the Verio data, red lines to data acquired on the Prisma scanner) and the average tSNR for all subjects (bold black line). The mean and standard deviation tSNR over all subjects is 27 ± 9.7 and there was no substantial tSNR difference between the data acquired on two scanners. The blue solid lines in the tSNR histogram correspond to the two subjects for which 2D and 3D EPI data was also acquired. One of these two TURBINE data sets had below average tSNR, while the other one had approximately average tSNR. The differences between the different subjects could be due to different quality of shim and receive coil loading or differing amount of physiological noise as well as subtle intra-volume motion (none of the subjects used in the comparison had substantial amount of bulk motion). As can be seen in the sagittal and coronal views in Fig. 5.11b, the maps generally exhibit higher tSNR in the superior and lower tSNR in the inferior parts of the brain (likely due to higher physiological noise [196] and larger field inhomogeneity in the lower parts of the brain).

We compared the tSNR and tSNR efficiency (to account for differences in volume TR) of TURBINE with 2D and 3D EPI (see Fig. 5.12). The tSNR efficiency was defined as the tSNR normalized by the square-root of the volume acquisition time (TR_{vol} = 2.5/3.6/4.59 s for TURBINE, 3D EPI and 2D EPI respectively). For both subjects the tSNR values

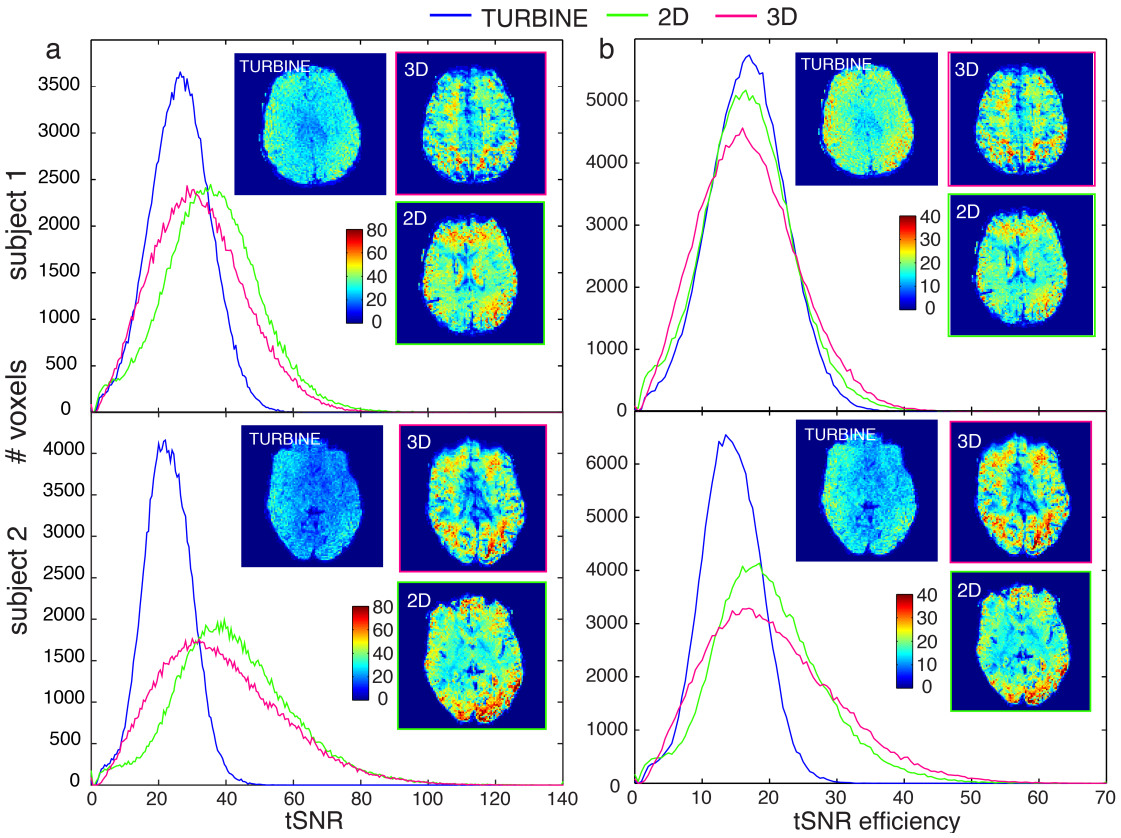


Figure 5.12: Comparison of tSNR and tSNR efficiency of TURBINE with 2D and 3D EPI for two subjects. The tSNR of subject 1 was very similar to the average tSNR over all subjects (see Fig. 5.11), the tSNR of subject 2 was lower.

were lower in TURBINE compared to 2D and 3D EPI. For subject 1 the tSNR efficiency was comparable to 2D and 3D EPI, for subject 2 it was lower. Given that subject 1 has an approximately average tSNR distribution (see Fig. 5.11a) we hypothesize that TURBINE can have comparable tSNR efficiency to 2D and 3D EPI. However, a more thorough comparison would be required for a definitive conclusion.

5.4.4 Preliminary fMRI Comparison to 2D/3D EPI

Two TURBINE task fMRI data sets were compared to 2D and 3D EPI data acquired in the same session. The activation maps are shown in Fig. 5.13, overlaid on example functional images. The image quality of TURBINE is comparable to 3D EPI; however, for subject 2 the TURBINE data set has some ringing (see transversal view in Fig. 5.13). The quality of activation maps is difficult to compare quantitatively in the absence of

a ground truth. Generally the z-statistic values are lower for 2D EPI, which is likely due to the long volume acquisition time required to achieve whole-brain coverage. For 3D EPI and TURBINE the activation maps look comparable for subject 1: 3D EPI appears to have more false positive voxels, but also more highly active voxels in the motor cortex. In subject 2, where TURBINE had substantially lower tSNR, the number of active voxels in the TURBINE activation map is reduced compared to 3D EPI. For future comparisons more realistic and optimised protocols for all three sequences should be chosen as discussed below. Although preliminary, these results indicate that TURBINE fMRI can produce comparable fMRI results to 2D and 3D EPI.

5.5 Preliminary 7 T fMRI

The (ultra) high field data was acquired on a 7 T Siemens system, using a 32 channel head coil for reception and a single channel head transmit coil (as used in chapter 4). TURBINE acquisition protocols for four different spatial resolutions (2, 1.5, 1.25 and 1 mm) were tested (see table 5.2 for experimental parameters). For the two highest resolutions we collected functional data using the same visual-motor task and processing pipeline as described for the 3 T data above.

Example images at 2, 1.5, 1.25 and 1 mm isotropic resolutions are shown in Fig. 5.14. There are larger amounts of distortion, streaking artefacts and variations in image intensity in the 7 T TURBINE images compared to the 3 T images (compare the 2 mm data to the images in Fig. 5.6). The lower resolution scans exhibit Gibbs ringing, substantial amount of streaking artefacts and distortion (most prominently visible in the 1.5 mm data set).

Increased B_0 inhomogeneity is likely, at least partially, responsible for the higher amount of artefacts. A larger amount of trajectory error as well as the higher physiological noise might also contribute to the image deterioration. Finally, the increased B_1 inhomogeneity can cause spatially varying contrast across the brain.

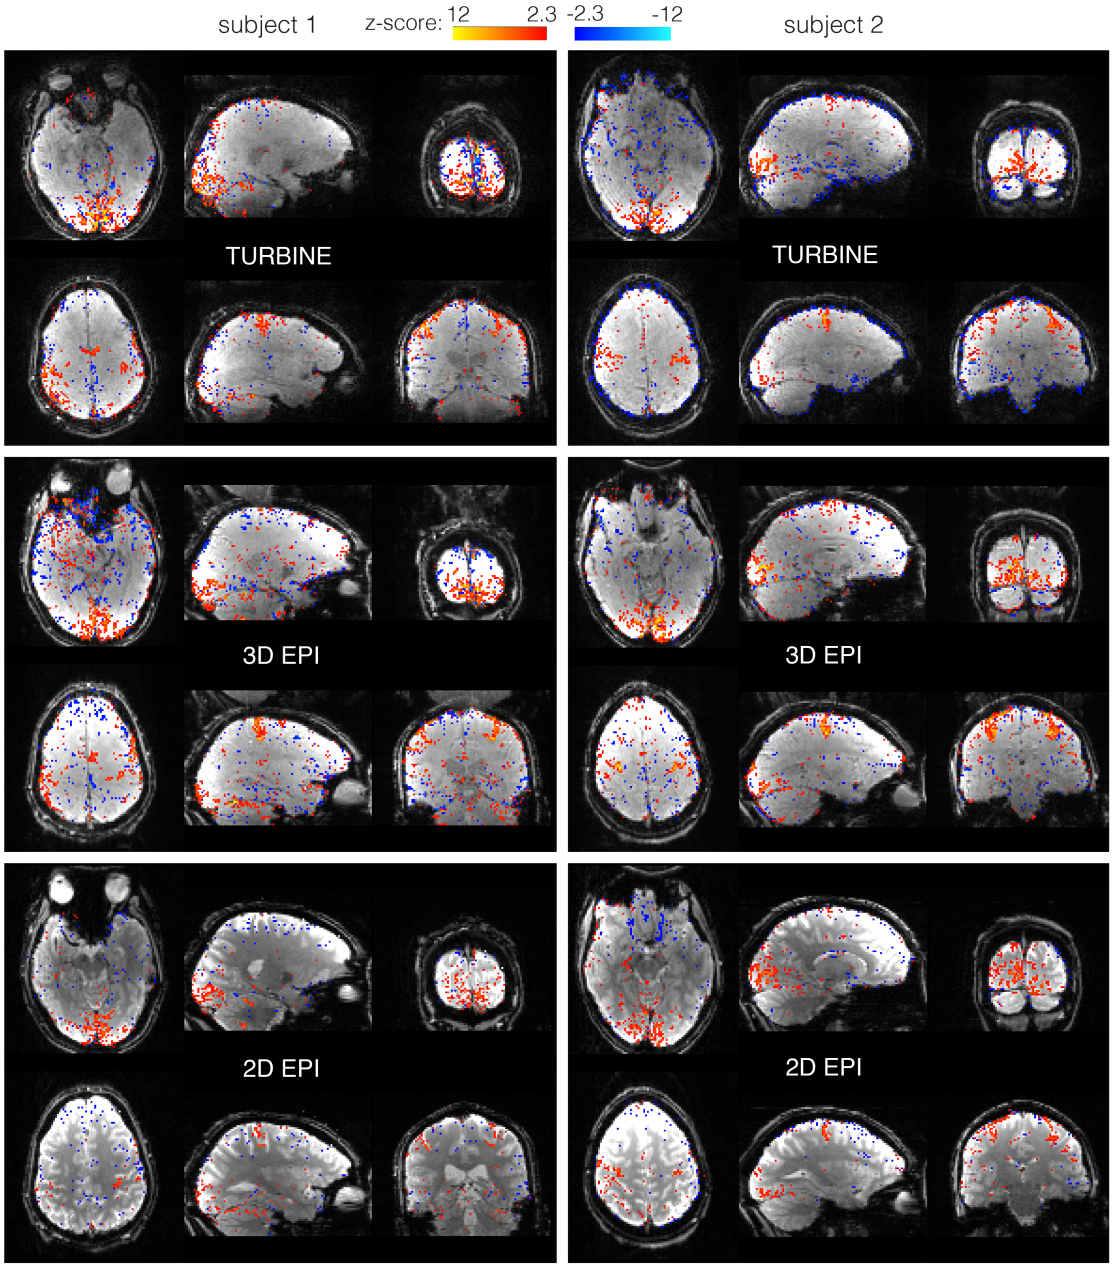


Figure 5.13: Comparison of activation maps for TURBINE, 2D and 3D EPI fMRI for two subjects. The z-statistic maps are overlaid on example functional images.

Despite these challenges, the preliminary functional results for one subject at 1.25 mm and 1 mm resolution are promising. The high-resolution z-statistic maps show functional activation following the grey matter architecture very accurately. For example the motor activation from the bilateral finger-tapping can be accurately localized to the hand motor area on the pre-central gyrus (see omega [197] feature in zoomed image in Fig. 5.15).

Whole-brain coverage at the higher resolutions or sub-mm resolution would require more

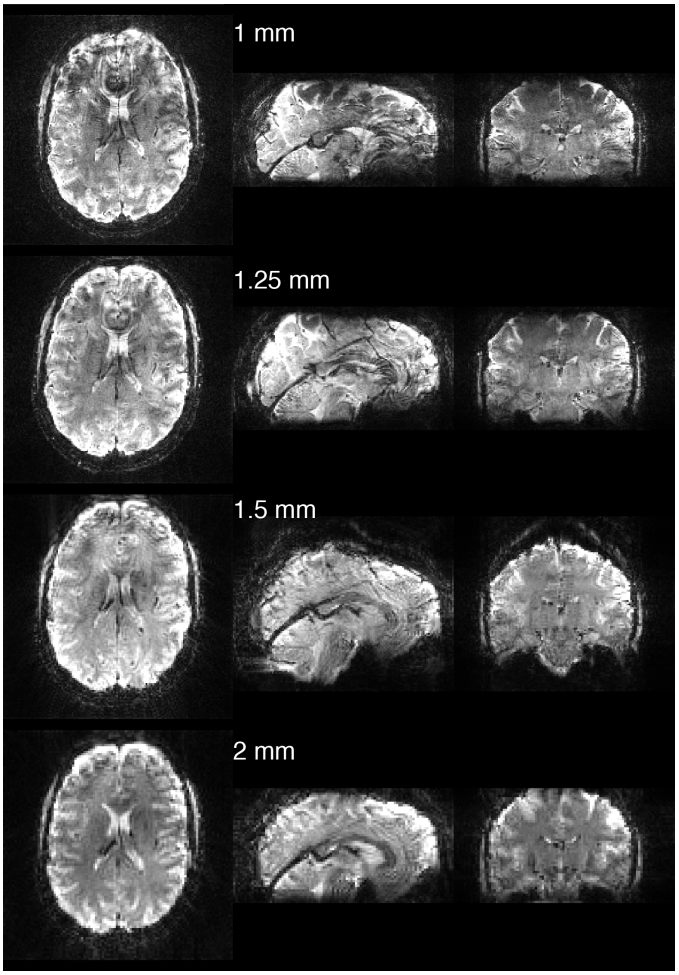


Figure 5.14: Example 7 T TURBINE images for four different TURBINE data sets.

acceleration (for example using also Partial Fourier and more distributed under-sampling patterns). In chapter 7 we discuss some of the changes we are planning to implement to improve high-field TURBINE imaging.

5.6 Discussion

We have shown that our TURBINE pulse sequence generates good quality images and fMRI data at both 3 T and 7 T. We set up a reconstruction pipeline, which can adequately process TURBINE data and remove artefacts due to gradient delays. We demonstrated the flexible spatio-temporal resolution property of TURBINE acquisitions using a golden angle blade update.

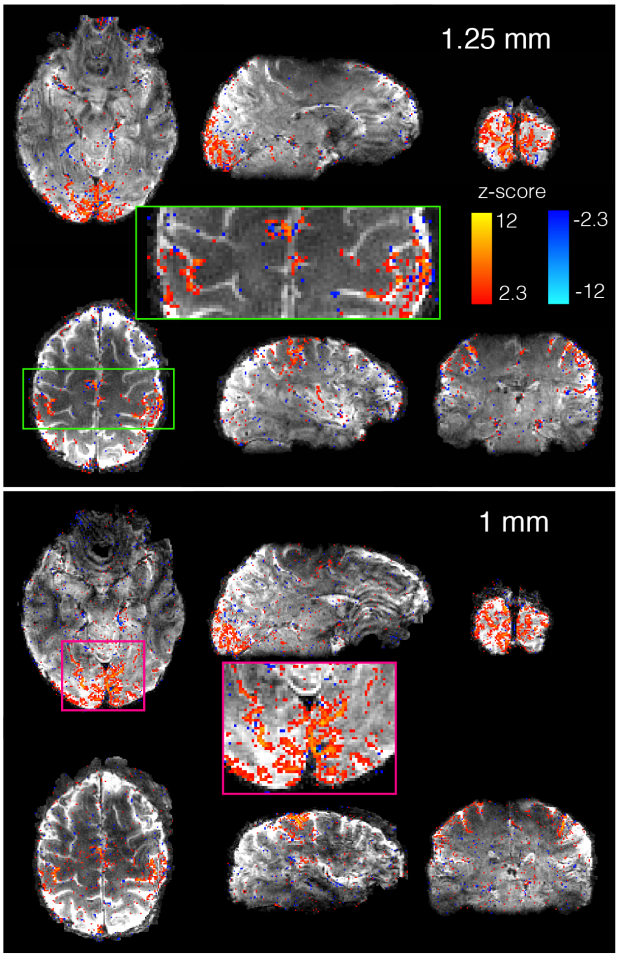


Figure 5.15: 7 T fMRI results: Activation maps for 1.25 mm and 1 mm resolution TURBINE fMRI data sets. Example zoomed regions show good correspondence of activation with the grey matter architecture.

TURBINE acquisitions were observed to induce low levels of PNS when minimal TRs and TEs were used. The PNS was within the scanner safety limits and of similar magnitude as for standard EPI using the y-gradient (AP direction readout) for the readout. Stimulation could be reduced by changing the gradient axes used in the readout direction or by increasing the echo spacing (resulting in a longer acquisition).

The rotating readout direction in TURBINE results in changing eddy currents, compared to 3D EPI, where the readout gradients stay constant. For this reason eddy current artifacts are more severe and more difficult to remove for TURBINE imaging. It has been reported that for fast balanced steady state sequences the rapidly changing eddy currents

in the golden angle scheme are responsible for image artifacts [198]. Using a 'tiny golden angles' scheme, which reduces the jump between consecutive blades but has very similar properties to the golden angle, they were able to reduce artefacts. Likely, it would provide less improvement for our sequence but it could account for a small amount of artefact.

The aliasing properties of the trajectory could be further improved by ensuring more distributed k-space samples. For in-plane acceleration, the skipped lines could be alternated from shot-to-shot (e.g. for R = 2, every second blade acquires even lines, the others odd lines). It may also be beneficial to include CAIPI [199, 200] gradient blips (shifting k-space samples perpendicular to the blade orientation), as 3D EPI-CAIPI [84, 85, 156] has been shown to have superior under-sampling properties and high tSNR. Potential sequence improvements as well as alternative designs of the sequence are discussed in more detail in chapter 7.

5.6.1 Image Reconstruction

Possible improvements to the reconstruction pipeline include adding an additional regularization term in the objective function of the SPIRiT reconstruction, for example a total variation or wavelet penalty (L_1-SPIRiT) [32]. We would expect this to help prevent the semi-convergence behaviour. Incorporating off-resonance correction into the iterative reconstruction (by including information from a field map into the data consistency term) would likely improve the image quality, especially at 7 T.

Any reconstruction framework that can handle arbitrary non-uniform trajectories with sensitivity encoding/parallel imaging features would be suitable to reconstruct TURBINE data. SPIRiT was chosen because it could handle non-Cartesian data, and was straightforward to extend to three dimensions. However, other non-Cartesian parallel imaging methods, such as CG SENSE [191] or ESPIRiT [28] could be considered instead of SPIRiT. For example, we have already demonstrated the use of rank-constrained k-t accelerated reconstructions based on a sensitivity encoded forward model using this trajectory [168].

5.6.2 TURBINE Trajectory Properties and Artefacts

TURBINE combines radial imaging with Cartesian sampling, therefore the sequence inherits some of the benefits and challenges of both. Aliasing caused by under-sampling is characterized by the point-spread function of the sampling, resulting in radial artefacts in one plane and in the EPI-like artefacts in the Cartesian direction.

Compared to a fully Cartesian sampling scheme, TURBINE imaging is inherently less efficient. To satisfy the Nyquist criterion at the edges of k-space (in the angular direction) $\sim$1.6 times more EPI blades need to be acquired than EPI planes for an equivalent Cartesian acquisition. However, a higher amount of under-sampling is tolerated due to the more incoherent radial sampling artefacts. Noise amplification in parallel imaging reconstruction (characterized by the g-factor) depends on the coil geometry as well as the distribution of the under-sampling in k-space. For TURBINE imaging (and non-Cartesian approaches in general), the acceleration is performed in multiple directions, which enables the coil sensitivities to be used more effectively [82]. Due to the over-sampling of central regions of k-space, high acceleration factors mainly affect the periphery of k-space, while the image contrast and geometry are still captured accurately.

Radial imaging is associated with more complicated artefacts due to off-resonance – blurring rather than shifts, which are harder to correct for. In TURBINE, each shots comprises an EPI blade (rather than a plane of radial spokes); therefore the off-resonance causes shifted signal and looks similar to 3D EPI using sagittal or coronal partition orientation.

Trajectory errors are generally more problematic in non-Cartesian imaging. In Cartesian data, with the readout direction constant, to a first approximation (assuming only a simple delay) all samples are shifted by a similar amount. This results in a constant phase ramp across the image and the magnitude image is not affected. In radial imaging a simple delay causes misalignment of the k-space centre (see section 5.3.2). This adds directionally dependent phase modulations in image space, which can result in constructive and destructive interferences. After reconstruction the object will exhibit signal cancellation, streaking and blurring.

Robustness to motion is typically cited as an advantage of radial trajectories. This could also be rephrased as Cartesian acquisitions being especially sensitive to motion. Due to their focal point spread functions, the phase errors caused by motion result in ghosts, which can overlap with the main image. These overlaps can easily be confused with structure in anatomical imaging and can cause misplaced activation in fMRI. In pure radial imaging the energy of the artefacts are spread more homogeneously over the object and phase inconsistencies between spokes cause streaking and blurring, which is generally considered less confounding when interpreting the image. In 3D EPI (and TURBINE) every EPI plane is acquired very quickly (tens of milliseconds). Thus, the majority of the motion artefacts in 3D EPI and TURBINE stems from inconsistencies between the different planes. TURBINE motion artefacts mainly result in radial streaking artefacts for rotations and streaking, blurring and dropout artefacts for translations. In Cartesian 3D EPI a head rotation can leave a large gap in k-space, while in TURBINE using a golden angle blade update this is very unlikely to occur.

Another specific feature of the TURBINE trajectory is the overlapping of all blades in the k-space centre, which creates both opportunities and challenges. Under ideal conditions the signal phase in the k-space centre should remain static over the acquisition. Changes in B_0 caused for example by subject motion can violate this assumption. Self-intersecting trajectories generally are very sensitive to phase errors because incoherent phase can result in destructive interference. This can manifest itself as a signal dropout in the reconstructed image, particularly when the intersections occur near the centre of k-space. Global phase offsets between the blades are straightforward to remove as mentioned above. However, this only removes the spatially averaged effect of the field changes, and does not take into account any localized phase evolution across the acquisition. It is well known that the B_0 field changes caused by respiration are not homogeneous across the brain and changing the position of the head within the B_0 field will change the field distribution inside the head. The phase errors could be reduced using shorter echo times but this is not compatible with fMRI.

While TURBINE motion artefacts mostly are more benign than in 3D EPI, TURBINE can have more severe dropout artefact in the presence of large phase changes. However, in contrast to 3D EPI, TURBINE is well suited for motion correction schemes building on its self-navigating properties; this will be explored in chapter 6.

5.6.3 TURBINE tSNR and fMRI

Due to the different sampling strategy in TURBINE, changes in B_0 (creating phase inconsistencies between blades) and time-varying aliasing (compared to using fixed sampling) could reduce tSNR. The tSNR values of TURBINE are within the range reported for 2D and 3D EPI in the literature [67, 81, 85, 112, 113, 143, 159, 160], see table 4.2. Our direct comparison for two subjects showed reduced tSNR for TURBINE and comparable tSNR efficiency for one subject and lower for the other one. Overall we think that likely the TURBINE tSNR is slightly lower than the tSNR of 3D EPI due to the presence of residual time-varying aliasing. However, further investigation is required to confirm this.

It would also be informative to compare the tSNR values calculated here to SNR to determine the signal stability. Unfortunately, the SNR is difficult to assess, and calculation would need to account for receive coil inhomogeneity, parallel imaging noise amplification, spatial noise auto-correlation due to the non-Cartesian trajectory, and the presence of residual streaking artefacts. For these reasons, SNR calculated with an ROI-based method is unlikely to provide meaningful results, and we focus our analyses to tSNR, which is the more relevant quantity for fMRI.

A number of approaches could be tested to improve the TURBINE tSNR. Firstly, TURBINE could be acquired with a fixed angle scheme to have the same aliasing pattern for all imaging volumes. This version would lose full flexibility of the temporal resolution. Our decision to use the continuous golden angle approach was motivated by the desire to demonstrate flexible temporal resolution, enabled by the low levels of observed aliasing in well-behaved subjects. Secondly, other blade ordering schemes should be tested in order to see how the tSNR compares to the golden angle acquisition. Thirdly,

since the spatio-temporal aliasing pattern is well-known we could attempt to regress out the fluctuations prior to image reconstruction.

The functional comparison between 2D EPI, 3D EPI and TURBINE indicates that TURBINE provides comparable fMRI results to the Cartesian sampling methods. However, for this preliminary comparison, none of the three sequences were used to their best potential. As a state-of-the-art 2D protocol, an optimized 2D simultaneous multi-slice sequence would be more appropriate, which would allow substantially faster volume TRs. For the 3D EPI parallel imaging along the second phase encoding direction should be employed, using CAIPI blips. For TURBINE the specific properties of the sequence should be exploited by performing physiological noise and motion correction (see chapter 6). Furthermore, we could make use of the increased incoherence using a reconstruction relying on sparsity or low-rank properties.

In future work it would be useful to compare fMRI results for a larger subject population and higher level tasks. The visual and visual-motor fMRI tasks are very robust and provide acceptable results, even in data sets with relatively low tSNR. It would therefore be useful to compare TURBINE to 2D and 3D EPI for resting state fMRI, where increased statistical power from faster sampling is more important. We have previously showed that TURBINE produces good results for resting-state fMRI in combination with fixed-rank k-t reconstruction methods [168].

5.7 Conclusion

We implemented a 3D hybrid radial-Cartesian EPI sequence and showed that it is suitable for the acquisition of fMRI data at 3 T and 7 T. We were able to achieve reasonable reconstructions from only a low number of blades with significant under-sampling factor. These high temporal resolution images could be used as navigator images for motion correction, which will be explored in chapter 6.

6

Motion Correction for Functional MRI with TURBINE

Contents

6.1 Introduction 118
6.2 Motion Correction Concept 120
6.2.1 Correction for z-Translations Only 121
6.2.2 Image-Space Based Motion Estimation 124
6.2.3 Estimation of Motion Parameters in k-Space 128
6.3 In Vivo Validation of TURBINE Motion Correction 130
6.3.1 Methods 130
6.3.2 Results 134
6.4 Discussion 141
6.5 Conclusion 146

6.1 Introduction

Bulk motion and other temporal fluctuations due to hardware instabilities and physiological noise have a detrimental effect on the ability to detect functional activity. Corruption by spurious fluctuations can require the exclusion of affected imaging volumes from analysis (scrubbing) or even entire subjects from studies, which reduces statistical power. Motion corruption also makes clinical application of fMRI (for example pre-surgical planning,

where robust single subject fMRI is required) challenging, as many patients, especially paediatric patients, are unable to lie still for extended periods of times.

The standard motion correction approach in fMRI is optimized for rapid 2D imaging protocols, where it is assumed that subject motion is slow relative to the sampling window (tens of milliseconds for a single slice). In this case the motion can be considered as inter-volume only, which can be effectively dealt with by using image-based registration in post-processing [133]. In 3D whole brain imaging, intra-volume motion is non-negligible as the sampling window is much longer (typically more than a second), which makes shot-to-shot k-space inconsistencies more likely [145]. Uncorrected, these k-space inconsistencies can result in significant image artefacts. This includes dropout due to phase cancellations and ghosting artefacts as well as blurring, which reduces the effective spatial resolution of the acquisition. Therefore, in order to fully benefit from 3D imaging, these k-space errors must be minimized or effectively removed.

As discussed in section 3.5, motion correction is a vast topic with many different approaches. Motion correction specifically for 3D functional MRI has not yet been studied extensively, mainly due to the newly found popularity of 3D EPI. Motion cameras have shown promising results for motion correction in 3D EPI [137, 201], but they are costly and not yet commonly available. External full 3D navigation can be problematic due to the lack of dead time in an optimized fMRI pulse sequence. Therefore self-navigated sequences such as TURBINE are particularly attractive for fMRI. TURBINE has self-navigating properties while maintaining encoding efficiency due to the use of EPI. In this chapter we explore the potential of TURBINE in addressing the motion sensitivity of multi-shot 3D acquisitions. We will use TURBINE to perform fMRI experiments in the presence of subject motion, exploiting the trajectory's self-navigating and flexible spatio-temporal resolution properties for retrospective motion correction.

6.2 Motion Correction Concept

Motion correction with the proposed trajectory takes advantage of basic Fourier properties. In this work we correct for rigid body motion only. Any rigid body transformation can be expressed as a rotation by the linear transform $\boldsymbol{R}$ followed by a shift by the translation vector $\boldsymbol{T} = (\Delta x, \Delta y, \Delta z)$. Consider the object $I_0(x, y, z)$ translated by $\boldsymbol{T}$:

$$I(x, y, z) = I_0(x - \Delta x, y - \Delta y, z - \Delta z) \tag{6.1}$$

The Fourier shift theorem states that in k-space this image-space shift corresponds to a linear phase accumulation between the Fourier pairs $S(k_x, k_y, k_z) = \mathcal{F}\{I(x, y, z)\}$ and $S_0(k_x, k_y, k_z) = \mathcal{F}\{I_0(x, y, z)\}$:

$$S(k_x, k_y, k_z) = S_0(k_x, k_y, k_z)e^{-2\pi i(k_x\Delta x + k_y\Delta y + k_z\Delta z)} \tag{6.2}$$

Since rotation is preserved under the Fourier transform, rotating the image I_0 by the rotation $\boldsymbol{R}$, $I(x, y, z) = \boldsymbol{R}I_0(x, y, z)$ corresponds to the same rotation in k-space. Expressing the nominal k-space locations as $\boldsymbol{k} = (k_x, k_y, k_z)^T$, the rotated k-space is:

$$S(\boldsymbol{k}) = \boldsymbol{R}S_0(\boldsymbol{k}) = S_0(\boldsymbol{R}\boldsymbol{k}) \tag{6.3}$$

An object moved by a general rigid body transformation consisting of the rotation $\boldsymbol{R}$ and a shift by the translation vector $\boldsymbol{T}$ can be written as:

$$S(\boldsymbol{k}) = S_0(\boldsymbol{R}\boldsymbol{k})e^{-2\pi i(\boldsymbol{T}\cdot(\boldsymbol{R}\cdot\boldsymbol{k}))} \tag{6.4}$$

B_0 fluctuations, for example, caused by susceptibility change due to respiration, result in zero-order phase offsets Δk_0. Taking all of this into account, the k-space S of the moved object can be expressed with respect to the k-space S_0 acquired of the object at the reference position:

$$S(\boldsymbol{k}) = S_0(\boldsymbol{R}\boldsymbol{k})e^{-2\pi i(\boldsymbol{T}\cdot(\boldsymbol{R}\cdot\boldsymbol{k})+\Delta k_0)} \tag{6.5}$$

Ideally the rotation matrix, translation vector and global (zero order) phase offset would be known for all time points. In practice the temporal resolution of the motion estimates are determined by how we derive them using the self-navigating properties of the sequence.

There are two main stages in our motion correction scheme: estimating the motion parameters and correcting for the motion. In our approach the correction is performed in k-space (using equation 6.5), by adjusting the k-space sampling locations according to the rotations given and by removing the phase ramps corresponding to the translations, as well as the global phase offset.

The motion parameters can be estimated by making use of self-navigating properties of the TURBINE trajectory. The blades of TURBINE cross the centre of k-space for every shot, allowing estimation of global phase fluctuations (caused by motion and respiration) on a shot-by-shot basis. The central column (in the phase-encoding direction) of k-space is acquired for every blade. This enables the detection of shifts in the blade phase-encoding direction (typically along the z-gradient axis) on each blade. However, if simultaneous rotations occur this is no longer possible because the acquired central column of k-space does not have a constant orientation (as it rotates with the head rotation).

The golden angle sampling approach of TURBINE ensures that k-space is covered nearly optimally for an arbitrary number of blades. Although each blade alone only provides information about the projection onto the plane in which it lies, combining information from several blades allows determination of all six rigid body motion estimates (three translations, three rotations). This can be done either in k-space directly or in image-space, by reconstructing a high temporal resolution navigator time-series. The naming conventions for rigid body head motion are introduced in Fig. 6.1.

We first examined the special case of z-translations only, and subsequently investigated image-space based estimation of motion with all six degrees of freedom. Lastly we briefly discuss the potential for full motion estimation in k-space.

6.2.1 Correction for z-Translations Only

For a preliminary motion experiment, fMRI data on 3 healthy volunteers was acquired on a 3 T Siemens Verio scanner (32 channel receive coil) with a visual stimulus (30 s ON/OFF 8 Hz checkerboard). During a 200 s fMRI scan, subjects were asked (using auditory cues)

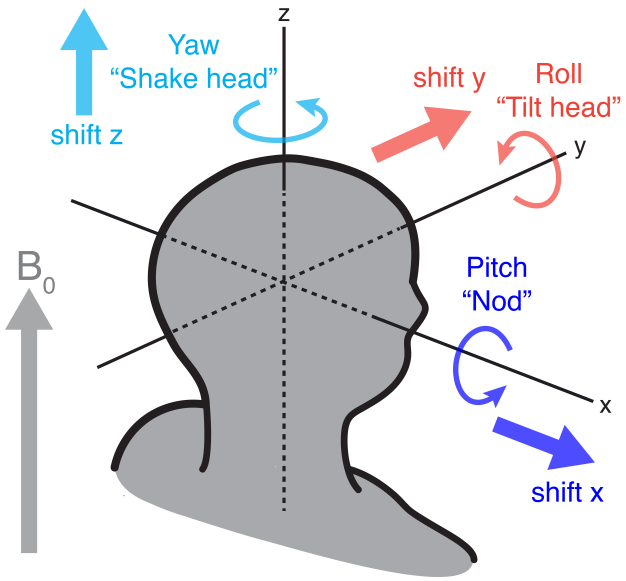


Figure 6.1: Illustration oft the six types of rigid body motion (three translations, three rotations) and the naming convention used.

to deliberately shift in the z-direction on a defined schedule (green arrows in Fig. 6.2). The acquisition used TR/TE=50/25 ms, blade matrix 100×76, 2 mm isotropic and R=2 along the blade phase-encoding direction. Nyquist ghost correction was performed as described in subsection 5.3.2. Note that the EPI phase correction lines were used only to correct Nyquist ghosting, and were not included in any of the motion or B_0 estimation (the same applies to all subsequent experiments in this chapter. GRAPPA reconstruction was applied along k_z prior to re-gridding every k_z plane. The deliberate motion data sets were reconstructed using 100 blades/volume (TR_{vol} = 5 s).

The coarse temporal resolution (and subsequently low under-sampling factor) enabled reconstruction with no visible aliasing artefacts, despite not using any parallel imaging in the radial direction. Linear phase ramps and zero-order phase offsets were fitted to the readout of each blade to determine the corresponding shifts along z and a B_0 offsets. The fit was performed on a high-SNR central k-space region consisting of a column of k-space at the nominal k-space centre (N/2 + 1 readout point) with a radius of 10 in the z direction (resulting in a 1x21 vector). Prior to the fit of the phase data the information from all coils was combined using a complex sum. Temporal phase unwrapping was used and the resulting phase estimates were low-pass filtered. The

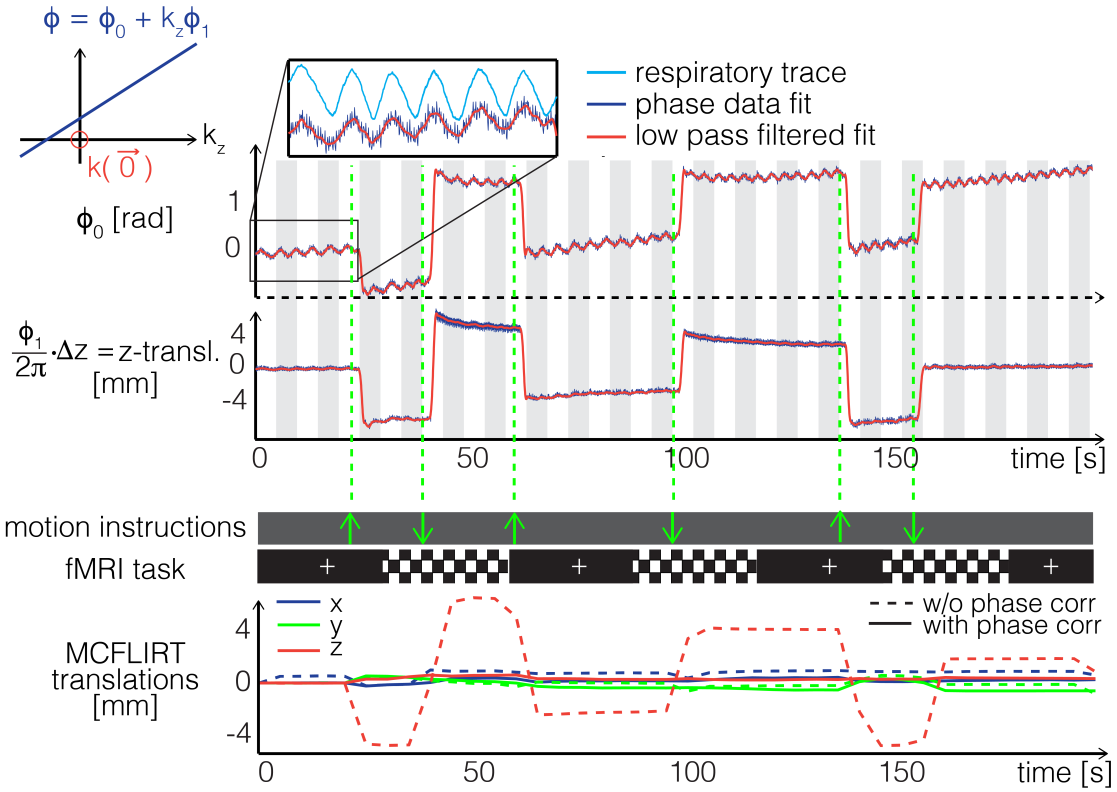


Figure 6.2: Pilot study with z-translations only: The motion cues (green arrows) and the fMRI visual stimulus timing are displayed in the middle of the figure. The top row shows the fitted phase coefficients. We can detect both the subject motion as well as the subject respiration (insert, overlaid with directly measured respiratory trace). The bottom panel shows the motion estimates determined from the reconstructed images before and after applying the phase correction.

determined phase effects were then removed in k-space from each blade prior to image reconstruction. To evaluate z-shift correction the data was reconstructed both with and without the correction step. The fMRI data was analysed in FEAT using no smoothing or pre-whitening. Rigid body motion correction using MCFLIRT was performed across all datasets. This final registration step was performed because the main goal was to evaluate the reduction in artefacts caused by intra-volume motion and this estimate could be confounded by small residual shifts or rotations.

Estimated phase correction parameters are shown for a representative subject in Fig. 6.2. The linear phase parameters show good correspondence to motion correction z-shift estimates from uncorrected data, as well as the intended head motion. The phase estimates capture the motion on a blade-by-blade basis and are therefore able to characterise the rapid motion transitions in these data sets. The zero-order phase offset, reflecting

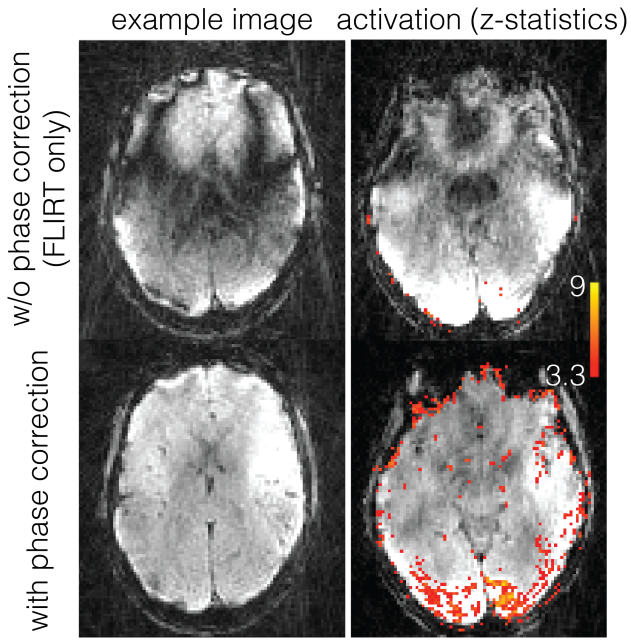


Figure 6.3: Result of z-translation and global phase correction: Example image at t = 46 s (left) and z-statistics maps with and without phase correction (right). MCFLIRT was applied for all data sets.

variations in B_0, tracks the measured respiration and slow B_0 drift. Under deliberate motion, the uncorrected data sets contain volumes with severe artefacts (Fig. 6.3) and show no above-threshold visual activation. After correction, corrupted volumes and visual activation can be recovered, although there is some artificial activation at image edges due to stimulus-correlated motion. Overall, the voxel-wise mean temporal variance was reduced by 45/35/21 % and the mean z-translation was reduced by 93/83/66 % for subjects 1-3 respectively.

6.2.2 Image-Space Based Motion Estimation

The approach to derive motion parameters from navigator images has been used for a number of different applications [202, 203]. It has the benefit that it is straightforward to implement as it essentially consists of two image reconstruction stages with an image registration step in-between. The proposed TURBINE motion correction scheme consists of the following steps (Fig. 6.4):

1. **Navigator reconstruction:** Reconstruct temporal subsets of the data creating a high temporal resolution navigator image time series. Whenever we use the term "navigators" in this chapter we refer to these image-space navigators.
2. **Motion estimation:** Register all navigators to a chosen reference volume (e.g. the first one in the time series).
3. **Motion correction:** Correct the raw k-space data corresponding to each subset, based on the motion estimates determined for each navigator.
4. **Final reconstruction:** Reconstruct the motion corrected image at the target temporal resolution using the corrected k-space data.

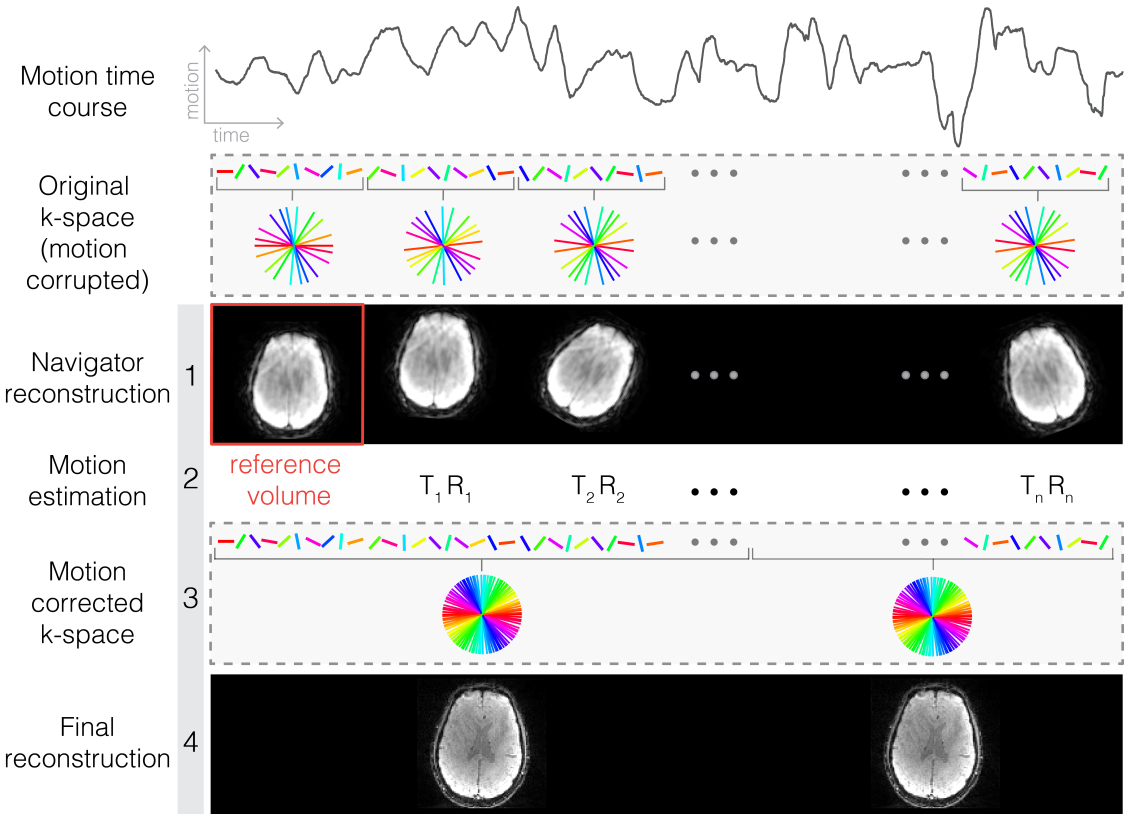


Figure 6.4: Schematic illustrating the use of the flexible temporal resolution of TURBINE to reduce intra-volume motion corruption.

For this approach accurate estimation of the motion is crucial and was studied on simulated data (6.2.2.1) before testing the full motion correction pipeline on in vivo data (6.3).

6.2.2.1 Motion Estimation Using Simulated Motion Data

The goal of the simulation was to study the choice of the temporal bin size for the navigator reconstruction (step 1 in the list above), i.e. how many blades to use to reconstruct each volume of the navigator time series. We hypothesise that there is an optimal number of blades expressing the trade-off between estimation of motion parameters (favouring more blades due to reduced artefacts) and the temporal accuracy of motion estimates (favouring fewer blades). We investigated this using simulated motion time courses.

The motion simulations used an 8-coil whole brain image with 2 mm isotropic resolution as the reference image. The coil sensitivity information was not manipulated in the motion simulation, i.e. the coil sensitivities were assumed to move with the head. We used multiple coils because we wanted the simulations to mimic the real-life reconstruction pipeline, including the parallel imaging reconstruction using SPIRiT. Only inter-blade motion was simulated (i.e. brain is stationary during the acquisition of each blade) and for each blade time point the reference image was rotated and translated into the prescribed location. The corresponding k-space data for this blade was generated using the NUFFT. We generated three artificial motion data sets allowing to study different types of motion: (a) a simple motion block design containing shifts and rotations; (b) motion based on a motion trace acquired with a motion camera for prospective motion correction, where the subject had been instructed to move the head slowly; (c) a four-fold accelerated version of the motion time-course in (b). We added the third simulation because we observed much faster motion in our data sets with free motion compared to the data from the motion camera. We termed these three simulations: "step motion", "slow motion" and "fast motion".

All simulated motion data sets were reconstructed at seven different temporal resolutions (100, 50, 25, 20, 15, 10 and 5 blades per image, using no sliding window) using image reconstruction pipeline D described in chapter 5. The motion parameters were determined by registering the navigator time-series to the first volume using rigid body motion registration with MCFLIRT. The extracted motion parameters were compared to the ground truth motion estimates, which had been used to generate the motion time courses.

The difference to the ground truth position was characterized using the Euclidean distance for translations and rotations (for simplicity the difference in rotation was calculated using the Euclidean distance of the Euler angles rather than a more appropriate distance metric,for example the geodesic on the unit sphere.

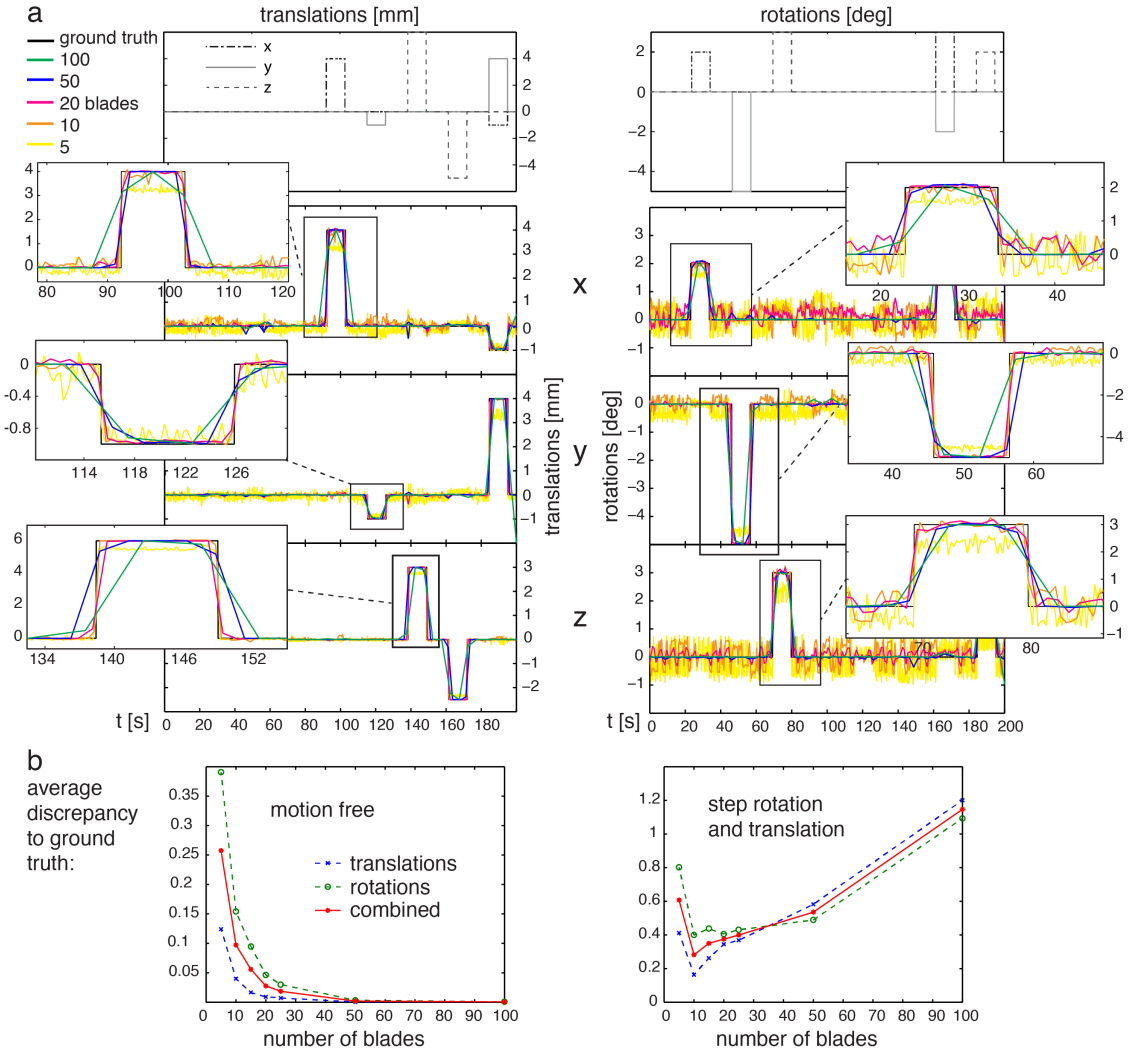


Figure 6.5: Motion estimation testing on simulated step motion data: (a) The top row shows the ground truth motion time course, which was used to generate the simulated motion. The six panels below show the motion estimates obtained for different navigator time series (using bin sizes between 100 and 5 blades). (b) The graphs show the differences to ground truth in a motion free time window as well as for a time window containing a step rotation and translation.

The simulations (see Figs. 6.5 and 6.6) show the expected trade-off between temporal resolution and precision of the motion estimates. When using a large number of blades (see estimates for 50-100 blades) we are unable to capture fast motion dynamics, while motion estimates based on the reconstructions of very few blades (see 5 blade estimates)

contain "jitter". This is especially pronounced for the rotation estimates and occurs due to the presence of residual aliasing in the highly under-sampled images. This aliasing varies from frame-to-frame, which can cause the cost function of the registration to be minimized when "fitting" to some of the streaking artefact. Motion occurring within the navigator acquisition window can further increase artefacts, making the image registration less reliable. In the step motion simulation we calculated the average distance from the ground truth for a motion-free window and for a window containing a rotation and translation step (see Fig. 6.5c). This illustrates how in a motion free time-series the error increases with fewer blades used. For a step motion change it shows that the estimate of motion requires high temporal resolution, in this specific example the 10 blade time-series (500 ms temporal resolution) provided the best result. This indicates that the optimum number of blades used for motion correction will depend on the temporal dynamics of the motion. This is confirmed in the simulation using the realistic motion trace. For "slow motion" the motion is captured accurately even for very coarse temporal bins, while for the "fast motion" higher temporal accuracy is required (Fig. 6.6a-b). The average difference over the entire simulations revealed minimum error for 20 blade bins for the slow motion and 10 blade bins for the fast motion (see Fig. 6.6c).

These simulations confirmed that there is a trade-off between temporal accuracy (given by the temporal resolution of the navigator time series) and spatial accuracy (driven both by resolution and artefact level of the navigator images). The best navigator choice depends on the image resolution, the amount and type of motion as well as the under-sampling pattern.

6.2.3 Estimation of Motion Parameters in k-Space

A disadvantage of the image-based motion detection approach is its reliance on registration and interpolation. An alternative approach is to determine the motion parameters in k-space. A direct estimation of the motion parameters in k-space would be more elegant than the image-space based approach and has been shown to work well in 2D radial sequences [204]. Above, we successfully performed single-blade estimates of z-translation directly in k-space, when motion only along this direction is present [165]. This could

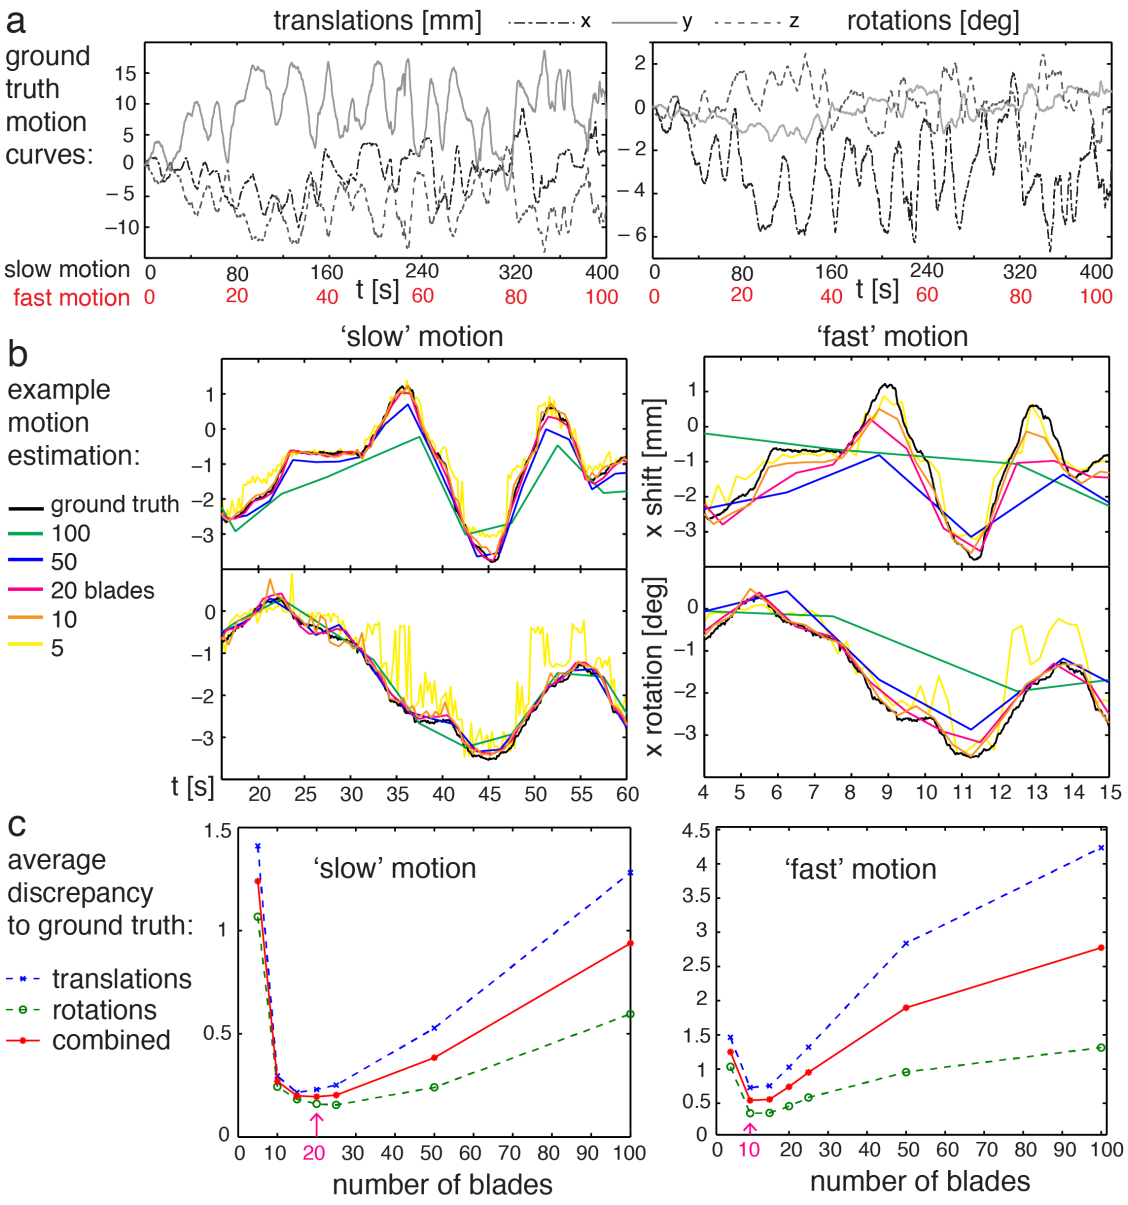


Figure 6.6: Motion estimation testing on natural motion simulation: (a) The top row shows the ground truth motion time course obtaining from a motion camera. It was used in its original timing as accelerated four fold to obtain the "slow motion" and "fast motion" data sets. (b) Example motion estimates for x translations and rotations obtained for different navigator time series (using between 100 and 5 blades). (c) The graphs with the difference to ground truth indicate the trade-off between temporal resolution and accuracy of the motion estimates. The temporal resolution resulting in the minimum amount of residual error is indicated with a pink arrow.

be expanded to determine all three translations by combining two or more blades for one motion estimation time point and fitting the phase ramps the same way as we had done for the z-translation correction. However, if an unknown rotation has occurred this estimate becomes invalid due to the rotation affecting both the k-space magnitude

and phase. The 3D rotations parameters could be estimated by registering the k-space magnitude and phase to a reference k-space [205]. However, this might exhibit similar trade-offs as the image-space approach: the fewer the number of blades the higher the temporal resolution but the less accurate the fit becomes. It is more promising to perform motion correction directly within the iterative reconstruction algorithm by incorporating the motion into the reconstruction model [206, 207]. This is discussed in more detail in the future works section in chapter 7.

6.3 In Vivo Validation of TURBINE Motion Correction

The goal of this study was to test the motion correction approach described above for retrospective motion correction in fMRI. The focus of the investigation was the motion correction but we also exploited the frequent crossing of the centre of k-space to remove global phase fluctuations driven by subject motion and respiration.

6.3.1 Methods

6.3.1.1 Data Acquisition

All data were acquired on a 3 T Siemens Prisma system using a 32 channel head coil for reception. The objective of the in vivo experiments was to evaluate the extent to which the proposed k-space trajectory and motion correction algorithm can correct for the whole range of expected subject motions. All five subjects were compliant healthy volunteers, who moved very little unless instructed to do so. All acquisitions used TR/TE = 50/29 ms with a blade EPI matrix of 100×76, 2 mm isotropic resolution with R = 2 acceleration along the blade phase-encoding direction. The readout bandwidth was 1852 Hz/pixel. For excitation we used a binomial water-excitation pulse with a flip angle of 15° (Ernst angle for grey matter), 3.84 ms duration and a time-bandwidth product of 10. The excitation slab thickness was set to 90 % of the phase-encoding FOV. At the end of every shot the residual transverse magnetisation was spoiled using 2π gradient spoiling along the x- and y-axes in addition to RF spoiling with a quadratic phase increment of 117° [15].

Three experiments were performed: (1) Cued motions, (2) Task fMRI with natural (deliberate but uncued) motion, (3) Task fMRI without explicit motion. Experiments two and three were performed in the same session with the order of the scans randomized.

1. **Cued motion:** 5 subjects were visually cued to perform 6 different motions (x/y/z translations and rotations) in turn at specified intervals during a 130 s scan. The experiment was designed to evaluate the proposed correction in the full range of isolated rotations/translations.
2. **Task fMRI with natural motion:** Two subjects were instructed to be "fidgety" during a 300 s task fMRI experiment with visual stimulus and simultaneous motor task. Specifically subjects were told to for example cross and uncross their legs and slightly reposition themselves as if they were uncomfortable. The visual stimulus consisted of a 30 s OFF/ON 8 Hz flashing checkerboard stimulus and subjects were also asked to perform bilateral finger-tapping during the ON blocks.
3. **Task fMRI without explicit motion:** Repeat of experiment 2 with subjects instructed not to move.

We also performed a temporal SNR measurement for two subjects, consisting of two 10 minute resting scans each, one with natural motion (motion instructions as for experiment 2) and one with no motion (motion instructions as for experiment 3).

All subjects were instructed not to move in the first 20 seconds of all scans to avoid motion contamination of the parallel imaging calibration data. Subject breathing was measured using a respiratory bellows.

6.3.1.2 Image Reconstruction and Motion Correction

The images were reconstructed as described in chapter 5. In short, the data were preprocessed using EPI Nyquist ghost correction and coil compression was used to reduce the reconstruction time (with virtual coils explaining ≥ 90 % of the variance, typically yielding 10-12 coils). The 2D GRAPPA kernels were calibrated for each blade angle using the closest blade in the calibration data and GRAPPA was run on every blade prior to

the 3D SPIRiT step. Note that we applied the 3D SPIRiT reconstruction (pipeline D in figure 5.4) to all data sets, including the ones without motion correction. These data sets could have been reconstructed used 2D SPIRiT (pipeline C), but we chose to keep the reconstruction consistent between data sets in order to make sure that differences in reconstruction would not be a confound when comparing data sets with and without motion correction. The final magnitude images were generated as a sum-of-squares combination of the individual reconstructed coil images.

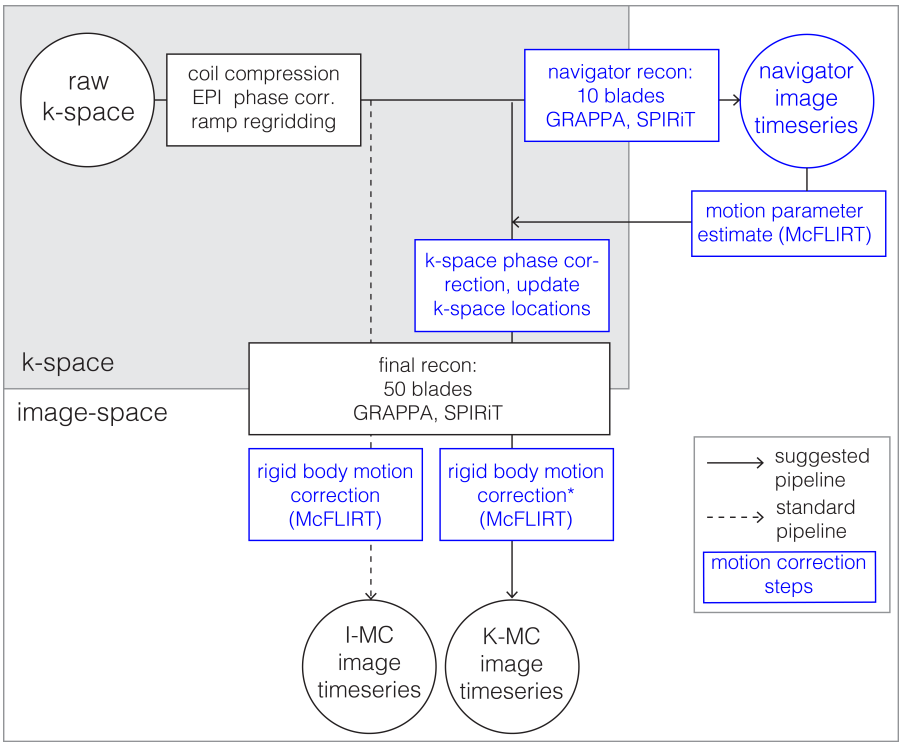


Figure 6.7: Overview of the motion correction and reconstruction process: "K-MC" pipeline presented in this project (solid arrows) and standard approach "I-MC", which was performed for comparison (dashed arrows). The grey shaded area indicates the steps of the pipeline taking place in k-space, all other steps operate in image space. Boxes and text in blue denote operations related to motion correction. The step marked with an asterisk is only performed for functional and image quality analysis.

An overview of the motion correction pipeline is presented in Figure 6.7.The entire motion correction pipeline was implemented in Matlab, with the execption of the navigator image registration step for motion detection, which was performed using MCFLIRT. The high temporal-resolution navigator time-series was reconstructed using 10 EPI blades per navigator image (TR_{vol} = 0.5 s, using no sliding window), 2D Hanning filtered to reduce high spatial frequency artefacts. The rigid body motion was estimated in image-space

using MCFLIRT, relative to a reference volume which we chose as the first volume of the time-series. The motion parameters estimated from this registration were then fed back into the reconstruction to correct each blade individually in k-space. To generate motion estimates for each blade from the 10-blade estimates, linear interpolation and spherical linear interpolation [208] were used for the translations and rotations respectively. The zero-order phase changes were estimated on a blade-by-blade basis (TR = 0.05 s) using the phase at the centre of k-space, and motion was corrected according to equation 6.5. To correct for rotations, k-space sampling coordinates for each blade were rotated accordingly. The shifts and global phase offsets were removed by multiplication with the conjugated phase ramps. The final reconstructions using the corrected k-space data were performed using 50 blades per image (TR_{vol} = 2.5 s) corresponding to a total under-sampling factor of ∼6 ($2 \cdot 3.14$), based on the Nyquist criterion for radial sampling (assuming regularly spaced blades). The correction pipeline was applied to all data sets, including those with no deliberate motion.

6.3.1.3 Evaluation and FMRI Analysis

We compared the proposed motion correction (MC) pipeline to the standard approach using only post-hoc image realignment (see Fig. 6.7). All datasets were processed with MCFLIRT to remove motion from the reconstructed image time-series. When this was the only correction, we refer to this as I-MC. Where we applied motion correction within the k-space image reconstruction to remove reconstruction errors, we call this K-MC; in this case, the final MCFLIRT alignments are expected to only remove small residual motions.

To quantify global inter-volume motion in subjects before and after K-MC in a single metric we used a motion score [135] derived from MCFLIRT motion estimates, which corresponds approximately to a maximum possible displacement. Image quality, reflecting the impact of the intra-volume motion, was quantified with an artefact score defined as the root mean-square of the difference image to a motion free reference volume.

Task fMRI analysis was performed in FEAT using no temporal filtering, spatial smoothing or pre-whitening. To account for residual autocorrelation the resulting z-statistic maps were corrected using mixture modelling [195] to ensure correct null distributions with zero mean and unity standard deviation.

The functional data from the experiment without deliberate motion (experiment 3) was used to derive masks for the functional evaluation of experiment 2. As this data was acquired in the same session, it should provide an estimate of the activation expected in a minimally motion corrupted acquisition to derive masks that neither favour K-MC nor I-MC. The masks were generated by intersecting thresholded ($z \geq 2.3$) z-stat maps from experiment 3 with hand-drawn visual and motor cortex maps. All activation maps are displayed on a previously acquired T_1-weighted structural scan of the same subject.

6.3.2 Results

The global (zero-order) phase offset k_0 reflects variations in B_0 (Fig. 6.8). It tracks the measured subject respiration (red lines), which can be seen through comparison to the externally measured respiratory trace (black lines), and captures the slow B_0 drift over the course of the acquisition (trend in red and green lines). In the presence of large motion, jumps in the global phase term can be observed.

6.3.2.1 Cued Motion Experiment

The results of the cued motion experiment are shown in Figures 6.9 and 6.10. The motion estimates derived from the navigator time-series (red lines) and the motion estimates from the final reconstruction using K-MC (blue lines) are shown for a representative subject in Figure 6.9a. The subjects performed translations in x (left-right), y (anterior-posterior) and z (head-foot), followed by rotations about x ("nod"), y ("tilt head") and z ("shake head"). All subjects executed the motions correctly, but most contained a combination of movements (e.g. x translations accompany rotations about y and z in Fig. 6.9a), as some of the motions are difficult to execute in an isolated fashion. The low motion parameters

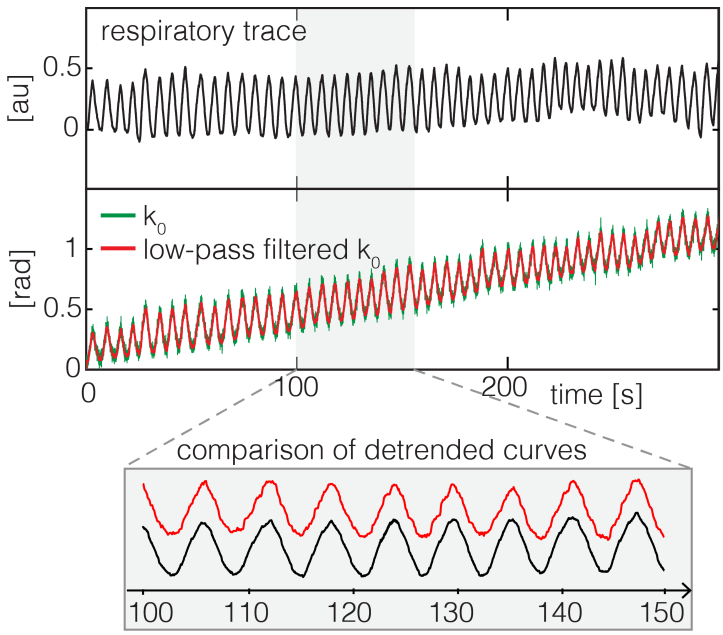


Figure 6.8: Global phase correction: Zero-order phase term for example data set with no motion displayed in green and low-pass filtered version in red. Zoomed plot below shows comparison of detrended global phase estimate with the corresponding respiratory trace, measured with a respiratory bellows.

and score after K-MC (blue lines) in Fig. 6.9a-b indicate very low residual motion. These values demonstrate the efficacy of the K-MC recon, which is expected to remove bulk motions as well as I-MC does. The artefact scores as well as visual inspection of the images were used to assess how effectively K-MC reduces artefacts compared to I-MC, which cannot correct for within-volume data inconsistencies. The artefact score (Fig. 6.9b) was reduced substantially following K-MC (blue line), compared to I-MC (green line). Many volumes of the uncorrected datasets contain severe artefacts (Fig. 6.9c), such as radial streaking artefacts and dropout artefacts, caused by phase inconsistencies from blade to blade at the k-space centre where the trajectories intersect. These artefacts are removed or reduced in the K-MC time-series. Some volumes, especially those where large rotations occurred, display residual artefacts (e.g. timepoint t_5), although in all cases these artefacts are reduced in K-MC compared to I-MC.

The results of the cued motion experiments for all five subjects are summarized in Figure 6.10, with Figures 6.10a-b showing plots of motion scores over time and Figures 6.10c-d showing box-whisker plots of motion and artefact score for all time points for

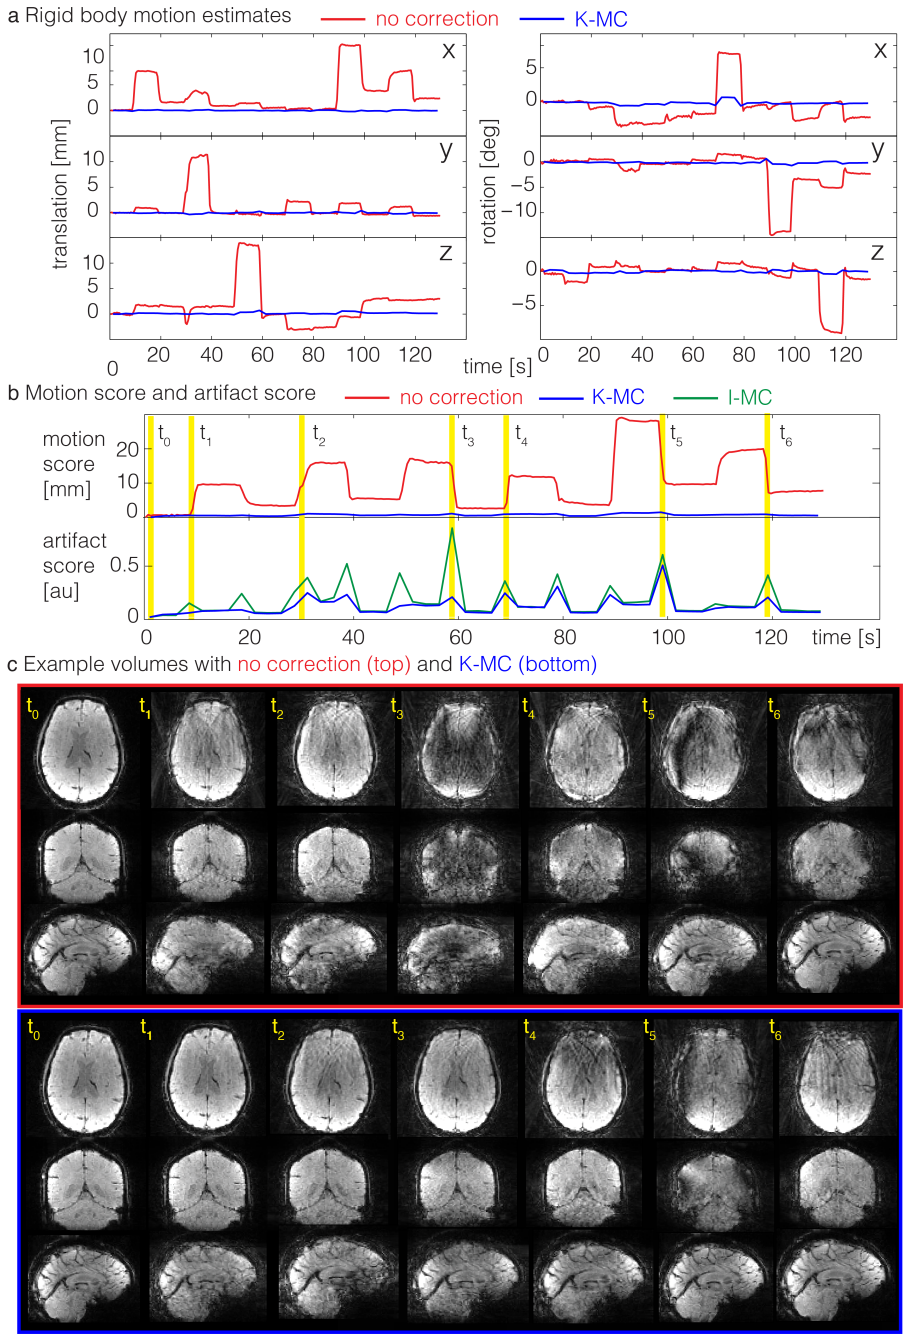


Figure 6.9: Cued motions (experiment 1): (a) Translation and rotation estimates prior to correction for a representative subject displayed in red, residual motion after K-MC (detected using MCFLIRT) in blue. The motion estimates before correction are calculated from the navigator time-series and used for correction, while the residual motion is measured on the final image time series. (b) Motion score summarizing motion prior to correction (red) and residual motion after K-MC (blue) in top panel; artefact score on final images using K-MC (blue) or I-MC (green) in bottom panel. (c) Example images, showing transversal, coronal and sagittal slices of uncorrected (top) and K-MC (bottom) volumes for seven time points. Time points marked with yellow lines in (b); t_0 is the first volume containing no motion, t_1-t_6 are example volumes at the different motion transitions.

each individual subject. The plot in Figure 6.10a displays the average (bold line) and minimum and maximum (shaded area) motion scores over all subjects with no correction (red) and using K-MC (blue). The average motion score for the motion estimates before correction was approximately 15 mm and the maximum exceeded 25 mm. After K-MC the largest residual motion score was 1.5 mm (inset in Fig. 6.10a). Residual translations along any axis did not exceed 0.7 mm and are below 0.5 mm for all but one subject. The maximum residual rotation along any axis was 1° for one subject and was much lower for all other subjects. Figure 6.10b equivalently shows the artefact score, which is generally reduced substantially for K-MC compared to I-MC.

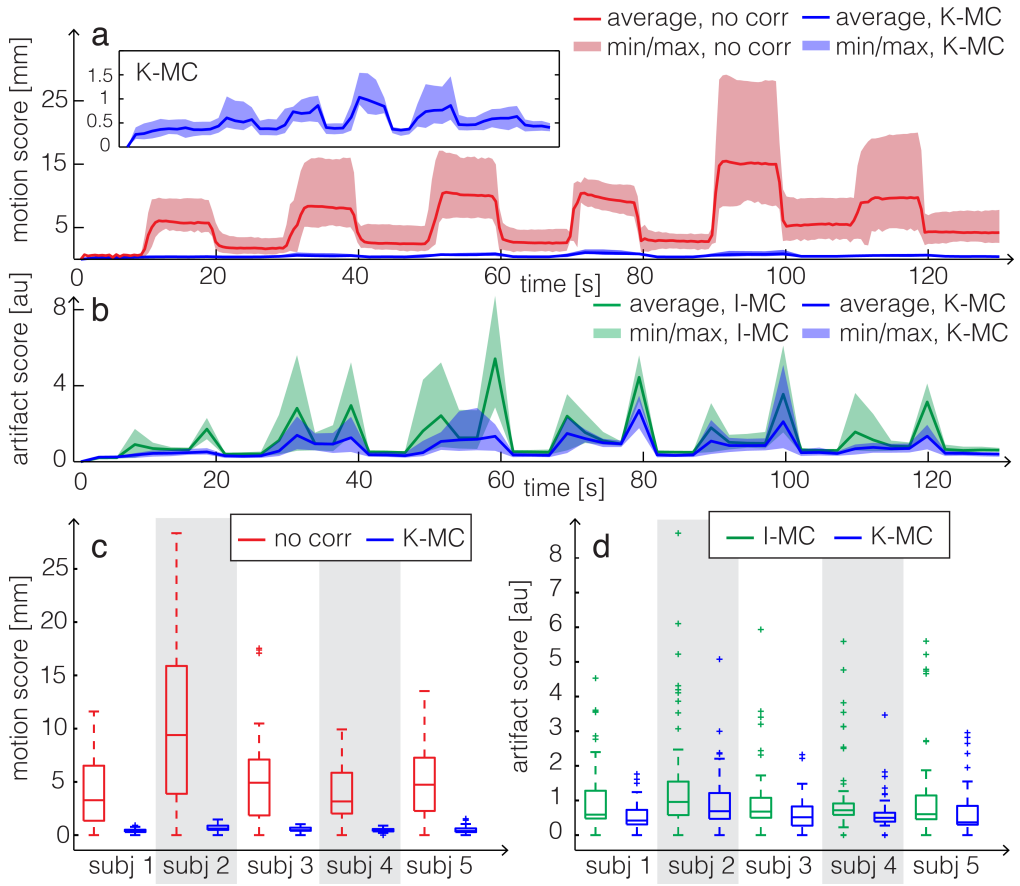


Figure 6.10: Cued motions (experiment 1): (a) motion score and (b) artefact score over time summarized for all subjects: The solid line shows the average over all subjects and the shaded area denotes the minimum/maximum value over all subjects. The inset shows the motion score after K-MC on an inflated scale. (c) Box whisker plot of the motion score before correction and after K-MC. This corresponds to an overall reduction in motion score of 90 %. (d) Box whisker plot of the artefact score, comparing K-MC to I-MC, showing a reduction of over 40 %.

However, for certain subject motions (e.g. at t = 100 s) there are substantial residual

artefacts. Averaged over all subjects and time points, the motion scores were reduced by 90.0 % and artefact scores were reduced by 40 % when using K-MC relative to I-MC. The box-whisker plots for the individual subjects (Fig. 6.10c-d) show large inter-subject differences between motion amplitudes and artefact level. In no cases were artefacts observed to increase following correction, either through visual inspection or based on artefact score.

6.3.2.2 fMRI with Natural Motion

In the natural motion experiment the subjects moved using all degrees of freedom (example subject in Fig. 6.11a), but z translations and x rotations had greater magnitude and were more frequent than the other motions. Presumably this reflects movement constraints by the coil and padding along the other directions and/or the most natural types of head motion. The translations were within a range of 10 mm, and rotations did not exceed a range of 10°, both larger than would normally be expected of naïve but compliant subjects. Using K-MC, the average motion scores were reduced by 90 % and 76 % and the average artefact scores were reduced by 41 % and 46 % for subjects 1 and 2 respectively, compared to I-MC (Fig. 6.11b).

Figure 6.12 shows the result of the functional evaluation of the visual-motor task fMRI experiment with natural motion. The activation maps for subject 1 (Fig. 6.12a) show visual and motor activation for both K-MC and I-MC. In the motor cortex the activation for K-MC localizes to the characteristic "omega" shape of the hand region of the precentral gyrus and also follows the folding of the central sulcus in the coronal view ("hook" shape). This is partially missing in the I-MC activation maps, indicating that the K-MC reconstruction recovered some of the activation, which had been lost due to motion (green arrows). Some of the visual activation in the I-MC data crosses white matter boundaries (yellow arrows), suggesting the activation is artefactual or misplaced. In the maps for K-MC (see zoomed subplots of the sagittal and transversal view in Fig. 6.12a) the activation follows the sulcal anatomy more accurately. The blue colour scheme for z-stats

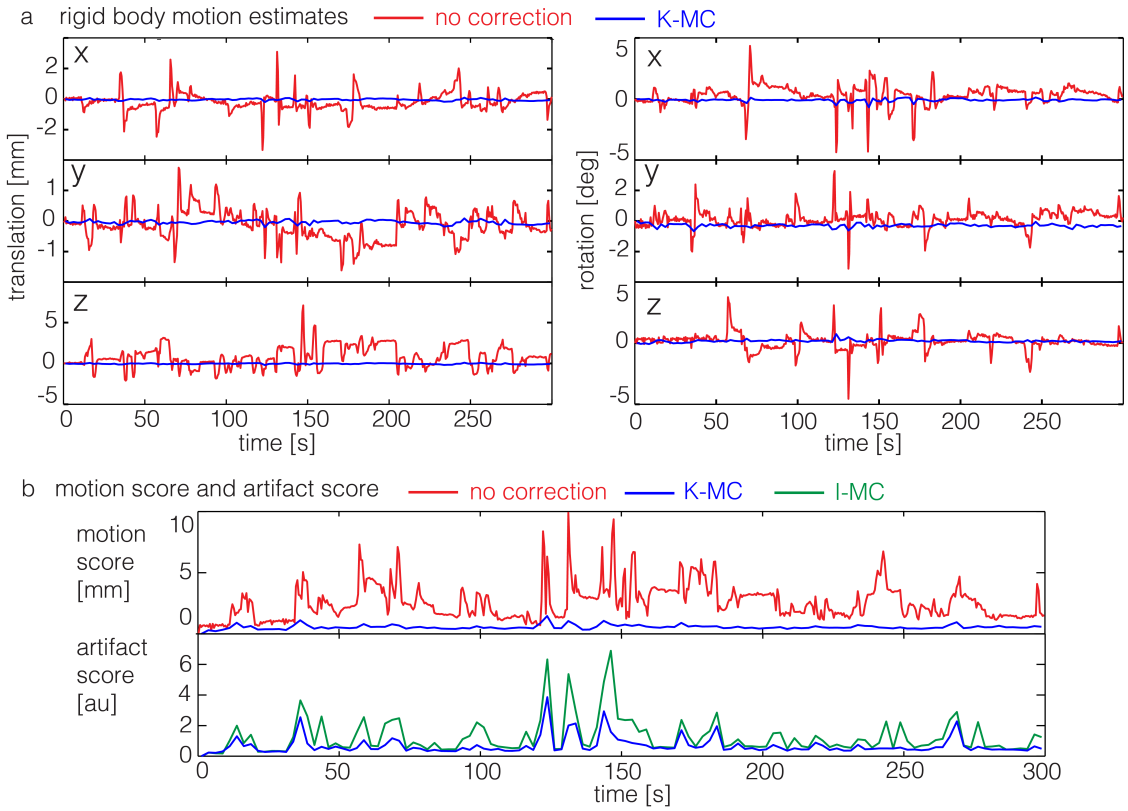


Figure 6.11: fMRI with natural motion (experiment 2): (a) Translation and rotation estimates based on the navigator time-series (red) for a representative subject for the fMRI experiment with deliberate but uncued motion. Residual motion after performing motion correction shown in blue.(b) Motion score summarising motion prior to correction (red) and residual motion after K-MC (blue) in top panel; artefact score on final images using K-MC (blue) or I-MC (green) in bottom panel.

represents artefactual negative correlations, likely driven by stimulus-correlated motion, which are reduced in K-MC compared to I-MC (visible in sagittal view in Fig. 6.12a).

For quantitative analysis of the z-stats we derived ROIs from the data with no explicit motion (as an approximation to a ground truth estimate for the correct activation). The number of voxels with activation ($z \geq 2.3$) in visual and motor cortex ROIs increased respectively by 16 % and 113 % for subject 1, and by 5 % and 41 % for subject 2. The average z-stats in the same ROIs increased by 13 % and 54 % for subject 1, and by 18 % and 40 % for subject 2.

Figure 6.13a plots time-series for using refined visual and motor masks that include only the 100 most active voxels in each mask (based on the low-motion functional reference scan). The time courses follow the stimulus design more clearly for K-MC time courses than for

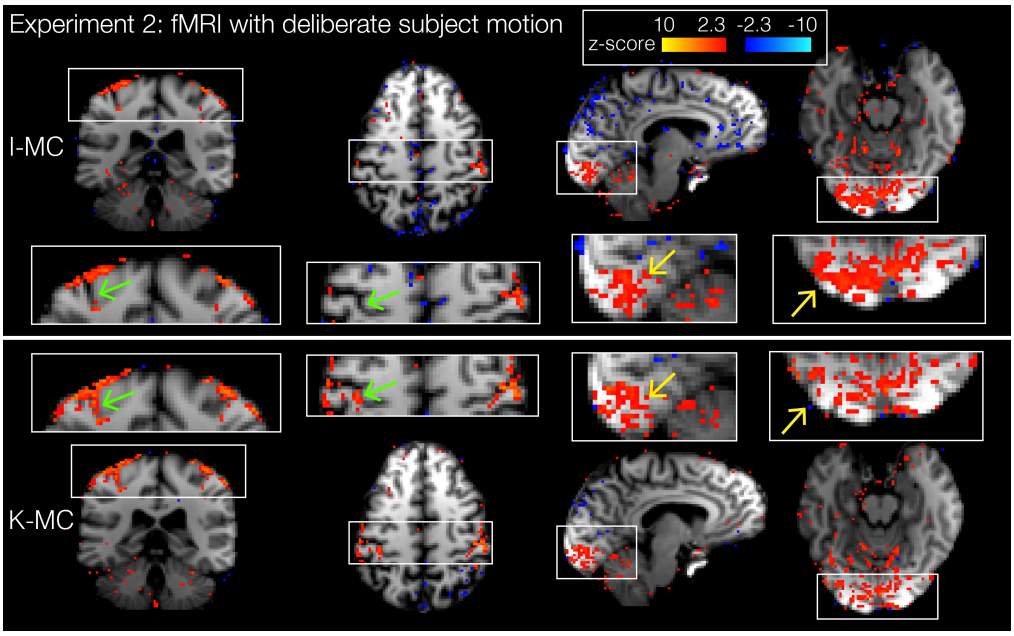


Figure 6.12: fMRI with natural motion (experiment 2): Activation maps for subject 1 using I-MC (top) and K-MC (bottom). Green arrows in zoomed images indicate examples of activation recovered using K-MC, the yellow arrows highlight reduction of misplaced activation when using K-MC.

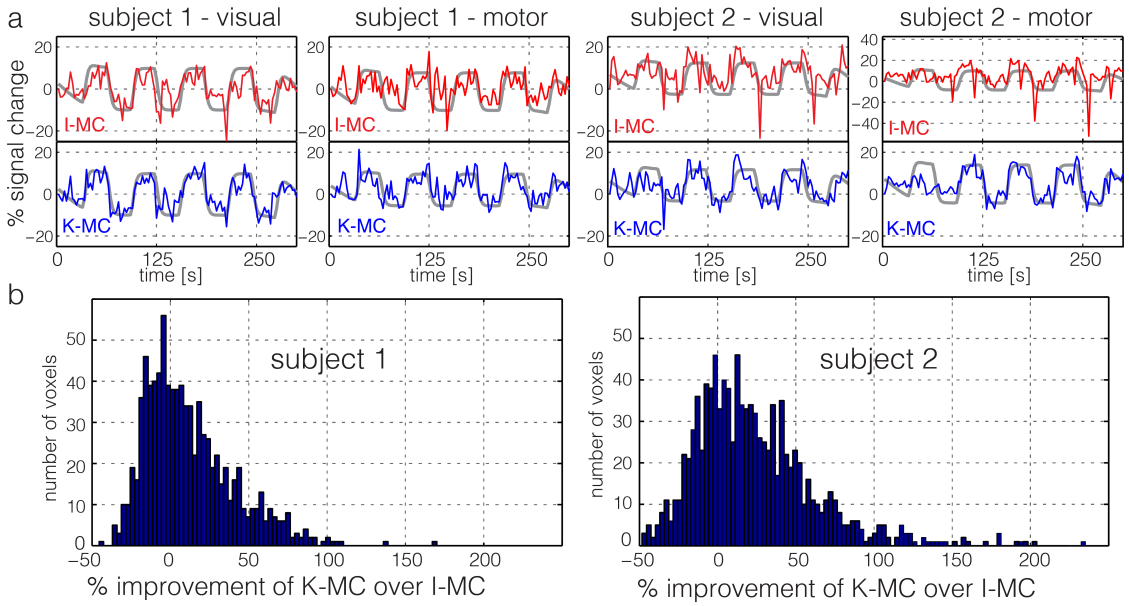


Figure 6.13: fMRI with natural motion (experiment 2): (a) Average time courses for motor and visual activation for both subjects using a high z-stat ROI based on the activation maps from the data with no explicit motion. Plots are displayed for K-MC (blue) and I-MC (red) with the FSL model underlaid in grey. (b) Histograms of percent improvement of K-MC over I-MC for active ($z \geq 2.3$ in both K-MC and I-MC) voxels in visual and motor cortex ROI.

I-MC. We evaluated percent improvement of the z-stats in the combined visual and motor ROI with the additional conditions of voxels included having $z \geq 2.3$ both for the I-MC

and K-MC data set (see histogram in Fig. 6.13b). For both subjects the histograms exhibit a strong skew, indicating that the z-stats for majority of voxels in the ROI are improved.

6.3.2.3 fMRI with no Explicit Motion

In the visual-motor task fMRI experiment with no explicit motion, subjects moved only minimally: the maximum motion score did not exceed 2 mm for either of the subjects, and the average motion score was 0.5 mm and 1 mm for subjects 1 and 2. The z-stat maps (Fig. 6.14) for subject 1 with I-MC show activation in the visual and motor cortices typical for the task used, but also a substantial amount of false positives (active voxels at boundaries of the head and isolated active voxels within structures that are not expected to be active in the task). Analysis of the motion estimates showed that both subjects moved (small amounts) in a stimulus correlated fashion, clearly visible in the z-translation motion estimate (Fig. 6.14a), which is likely due to the subject moving unintentionally when finger-tapping. The K-MC reconstruction substantially reduces motion-induced false positives, with the activation maps showing specific activation following the grey matter outlines in the visual and motor cortex (Fig. 6.14b).

6.3.2.4 Temporal SNR

Finally, the temporal SNR in the experiments with natural motion was averaged over the entire brain. Temporal SNR increased by 28 % (from 8.3 to 10.7) and 36 % (from 9.9 to 13.5) for subjects 1 and 2 respectively when using K-MC over I-MC. For the same scan with no deliberate motion the whole-brain averaged tSNR was 31.0 and 27.9 for subjects 1 and 2 respectively. In ROIs defined for the visual and motor cortex areas the tSNR was 40.1 and 39.8 respectively for subject 1 and 28.9 and 37.7 for subject 2.

6.4 Discussion

In this work we have demonstrated the motion correction capabilities of the TURBINE trajectory for fMRI, even in the presence of severe motion. The approach requires no

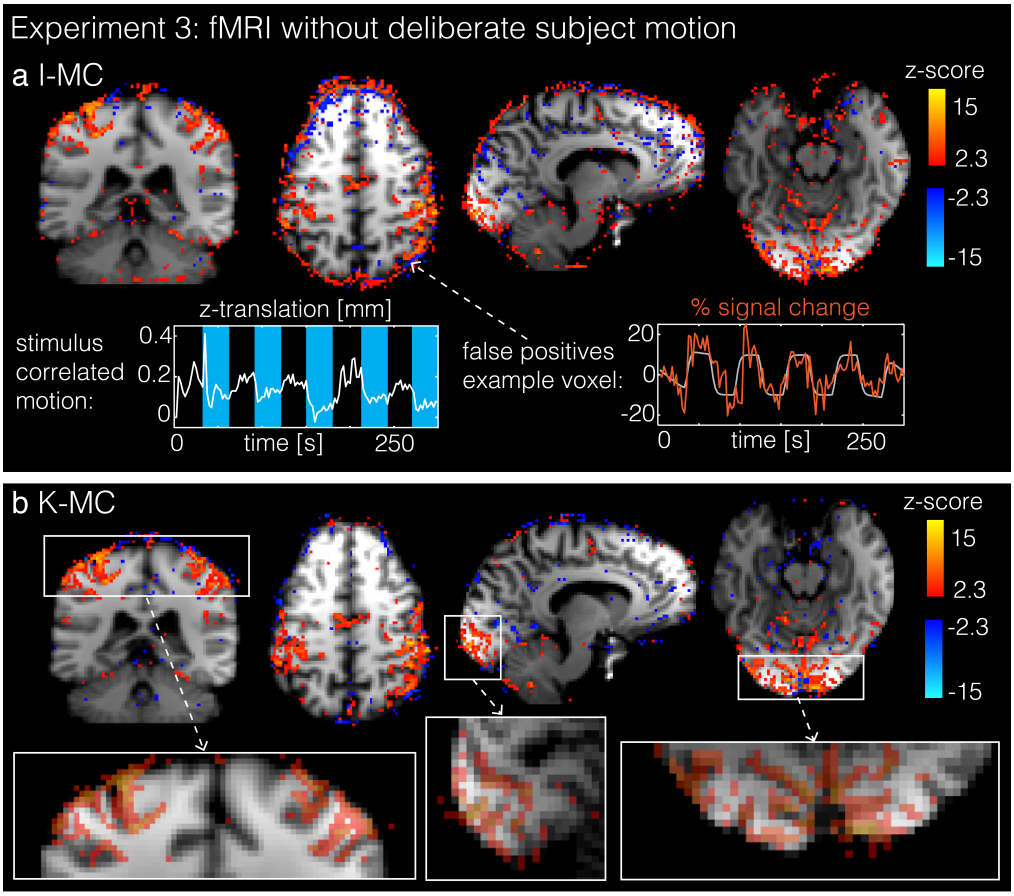


Figure 6.14: fMRI with no deliberate motion (experiment 3): Activation (z-stat) maps for an example subject with a small amount of (non-deliberate) motion. (a) I-MC z-stats maps containing a large amount of false positives, primarily at tissue boundaries. The false positives result from stimulus correlated motion as illustrated by insets (motion in the z-direction and time course for an example edge voxel). (b) K-MC z-stats maps, with reduced false positives. Zoomed activation maps below display positive z-stats maps overlaid (semi-transparent) on structural image to show correspondence of activation with grey matter architecture.

designated hardware and no separate navigator scans and is therefore highly compatible with fMRI, where we require short volume TR and re-acquisition is not possible. We used the intrinsically flexible temporal resolution of the k-space acquisition to derive motion estimates from the reconstruction of a high temporal resolution navigator time-series.

Using TURBINE with K-MC, most of the bulk motion was corrected, and artefacts were reduced substantially for all motion corrupted data sets. We were able to improve functional results both in data sets with very large subject motion and more subtle stimulus correlated motion, as reflected by recovered functional activity and reduced false positives. Residual motion artefacts (signal dropout and radial streaking) were observed

mainly for the most highly corrupted volumes (i.e. those with very high artefact score before K-MC). This could be due to the rigid body motion estimate being less accurate due to high levels of artefact in the navigator volumes, which are reconstructed at a relatively coarse 0.5 s temporal resolution. A disadvantage of the image-based motion detection approach is that the accuracy depends on the spatial resolution and image quality of the input time-series, leading to a trade-off between motion estimate fidelity and temporal resolution of the motion estimates. The 0.5 s windows used here for navigator reconstruction limits the temporal fidelity of the motion estimates, but in future work we will explore the use of sliding window filters (e.g. temporal hourglass [209]), which have been shown to improve the temporal resolution of motion detection. Residual errors might also come from variation in distortion and coil sensitivity when the head moves within the B_0 field and the coil, invalidating the calibrated GRAPPA or SPIRiT kernels. Further inconsistencies might be driven by changes in the MRI signal due to spin-history near slab boundaries, especially in the data containing large-scale motion, which K-MC cannot correct for. For very heavily corrupted data sets the tSNR remains substantially lower than for an uncorrupted data set even after correction, due to residual image artefacts.

Our approach cannot correct for intra-blade motion, but the impact of intra-blade motion could potentially be minimised by using a data rejection criterion, to discard blades acquired during very rapid motion transitions. Preliminary tests of blade rejection did not improve the quality of the reconstructed time series substantially. However, this was tested in data sets where the residual artefacts were largely driven by inconsistency between parts of the acquired 3D k-space rather than individual corrupted blades. Furthermore, data blade rejection results in temporally varying SNR, which could be problematic for fMRI. Another approach to reduce intra-blade motion corruption would be to apply the motion estimates on a line-by-line basis by interpolating between the detected motion estimates.

We observed that the zero order phase correction is very important in motion corrupted TURBINE data, especially when translations in the phase-encode direction occur. If uncorrected the phase discontinuities cause dropout artefacts. The zero order phase correction also allows removing respiratory induced B_0 fluctuations. In the study presented

here the zero order phase correction was applied together with the motion correction step (after the GRAPPA step) and the calibration data was not phase corrected. Most likely this is not the best approach and for future work we plan to apply the zero order phase correction to the reference data prior to the calibration of the GRAPPA and SPIRiT kernels. The removal of phase fluctuations induced by respirations might improve the parallel imaging reconstruction. Note, that while the zero order phase correction reduces physiological noise caused by respiration, it does not help remove any artefacts induced by cardiac pulsation as cardiac pulsation does not cause global phase shifts (instead it has more local impact). Therefore it would be beneficial to additionally perform RETROICOR on the final data sets.

The sampling SNR efficiency of non-Cartesian trajectories is intrinsically lower than uniform sampling [14], but the aliasing artifacts are more incoherent and therefore higher under-sampling factors are tolerated. Further increasing the robustness to under-sampling is important to allow TURBINE to compete with fast Cartesian 3D EPI methods, some of which are capable of achieving whole brain coverage at 2 mm isotropic in under a second [85]. The acquisition time of the navigators dictate the limit of motion correction in TURBINE. If a Cartesian sampling strategy could acquire the entire 3D k-space as fast as these navigators, the benefits of TURBINE will solely be due to correction of global phase inconsistencies. However, the fastest theoretical 3D navigator in TURBINE consists of two blades, because each pair of blades provides information about motion in all directions. How many blades are required in practice for an accurate k-space or image-space motion detection depends on the amount of noise in the data. With sequence refinements and integration of the motion detection into the image reconstruction pipeline, improved temporal resolution (compared to the 0.5 s navigators used in this work) of motion detection should be achievable.

The benefit of 3D in terms of SNR efficiency primarily relates to the ability to rapidly repeat volumetric excitations, whereas 2D methods, including simultaneous multi-slice (SMS), must cycle through slices before a given slice is excited again. The Achilles heel of 3D EPI is that it is segmented, making it sensitive to physiological noise and motion, while 2D

EPI or SMS are relatively insensitive to motion (although still susceptible to spin-history artefacts [210]). TURBINE, while segmented, has the potential to correct for motion and therefore combine high frame rate (conferring statistical benefits) with high SNR efficiency.

The absolute stability of the approach, in terms of tSNR, is what ultimately matters. It will be important to systematically evaluate whether 3D TURBINE with K-MC provides better fMRI results than a Cartesian 3D EPI or 2D (multiband) EPI acquisition with the best motion correction that can be done on all data sets. In future work, we hope to assess under what conditions (e.g. resolution, amount of motion) TURBINE confers the most benefit. See also chapter 7 for more details on the planned evaluation. We expect TURBINE to be most beneficial at high-resolution fMRI at ultra-high field, where 3D methods provide larger advantages over 2D and high resolution requires relatively long and therefore motion-sensitive acquisitions. For higher spatial resolution, more blades will be required to form an image with the same under-sampling factor. However, this could cause a reduction in the temporal resolution of the navigator time. In order to guarantee the same motion correction accuracy, it will be important to allow additional acceleration, for example by implementing the CAIPI shift. Potential sequence improvements are discussed in chapter 7.

In this study, we used compliant subjects to study the ability to correct for motion during the main phase of data acquisition. These subjects were specifically directed not to move during parallel imaging calibration acquisitions (first $\sim$20 s of scan), as motion corrupted calibration data can result in poor image reconstructions. If the subject cannot lie still for this amount of time, the protocol could be modified to be auto-calibrating by acquiring a fully sampled central region of k-space for every blade or altering the phase-encoding steps such that a fully sampled k-space is acquired after some number of blades. In that case any motion free segment of data could be selected to perform the parallel imaging calibrations. This would have further advantages, because non auto-calibrated parallel imaging in a motion-corrupted data set is problematic, as the coil sensitivities can change due to the motion.

In order to evaluate detection of activation in the presence of motion, we chose a relatively simple and robust task involving a visual stimulus with simultaneous finger-tapping. This task was simple enough to enable subjects to simultaneously conduct deliberate motions. However, there may be some downsides to this experimental design. Visual activation in the occipital lobe is located at a relatively fixed pivot point, making this area particularly insensitive to subject motion. Similarly, subjects' motor cortices were engaged in both finger-tapping task and cued motions, which could lead to poor task compliance and/or interactions between task and artefact. It will therefore be useful to consider more realistic settings, including higher-level tasks and more realistic motions.

Incoherent sampling schemes, such as TURBINE are well-suited for combination with reconstructions that build on concepts in data compression; for example we have used it in combination with the fixed-rank reconstruction scheme k-t-FASTER [168, 169]. As motion can violate the low-rank assumptions of this approach, we hope to improve the quality of reconstruction by combining it with the motion correction presented in this work.

6.5 Conclusion

We have demonstrated that the hybrid radial-Cartesian EPI is suitable for motion correction in fMRI, and has much potential for further development using more sophisticated motion estimation and correction schemes. Using the flexible spatio-temporal resolution of this sampling strategy we were able to correct data corrupted by severe subject motion and improve recovery of fMRI activation. Unlike many other methods for motion correction, this approach does not prolong scan time and no extra hardware for motion tracking is required.

7

Summary and Future Work

Contents

7.1 Summary 147
7.2 Future Work 148
7.2.1 TURBINE Sequence Improvements and Variants 149
7.2.2 Image Reconstruction Improvements Using Gradient Characterisation and B_0 Correction 150
7.2.3 Improved TURBINE Motion Correction 153
7.3 Final Remarks 155

In this chapter we briefly reiterate the main findings of this thesis and we present potential future directions building on this work.

7.1 Summary

This thesis focused on the investigation of functional MRI using 3D sequences. We initially analysed the behaviour of standard segmented 3D EPI in the context of physiological noise. Consistent with what has been reported in the literature, we found that 3D EPI is only beneficial compared to a 2D approach, when used in the thermal noise dominated regime or when performing physiological noise correction.

In order to address the sensitivity of 3D EPI to temporal fluctuations, we implemented the TURBINE sequence, an alternative 3D sequence using a hybrid radial Cartesian readout. We tested the sequence and the reconstruction pipeline on fMRI data and showed that, even without using the special properties of the TURBINE sampling scheme, we can achieve comparable results to standard acquisition methods.

We then made use of the self-navigating properties of TURBINE to correct for subject motion and respiration. We used the flexible temporal resolution of the readout to derive rigid body motion parameters from a spatially coarse, high temporal resolution navigator time-series for adjustment of sampled data directly in k-space. Furthermore, we exploited the self-navigated properties of the readout to remove global phase fluctuations driven by subject motion and respiration. The motion was corrected in k-space as part of the image reconstruction and was evaluated for experiments using both cued and natural motion. We demonstrated that the motion correction worked robustly and that we could achieve substantial artefact reduction as well as improvements in temporal SNR and fMRI activation. We believe TURBINE has potential as an alternative to Cartesian 3D EPI, especially for subject groups with significant head motion. Notably, this motion correction approach does not prolong the scan and no extra hardware is required.

The golden angle acquisition ordering scheme results in temporally variable aliasing, which can reduce tSNR and the ability to accurately detect motion. Furthermore, we have seen that motion can happen on very fast time scales and our motion estimates are currently based on relatively temporally coarse bins and thus unable to capture some of the faster motion transitions. However, we have not yet exploited the full potential of the TURBINE sampling scheme and a number of improvements are possible to increase its efficiency and accuracy.

7.2 Future Work

We have divided our suggestions for future work into improvement suggestions for the TURBINE sequence 7.2.1 and reconstruction 7.2.2 and ideas for improved motion and

physiological noise correction with TURBINE 7.2.3.

7.2.1 TURBINE Sequence Improvements and Variants

The most important improvement we have planned is to include CAIPI gradient blips into the TURBINE sequence to improve its under-sampling properties. The easiest modification is to alternate the lines, which are skipped in the EPI plane from blade to blade. In a second step additional gradient blips perpendicular to the EPI blade can be introduced to further distribute the samples in 3D k-space.

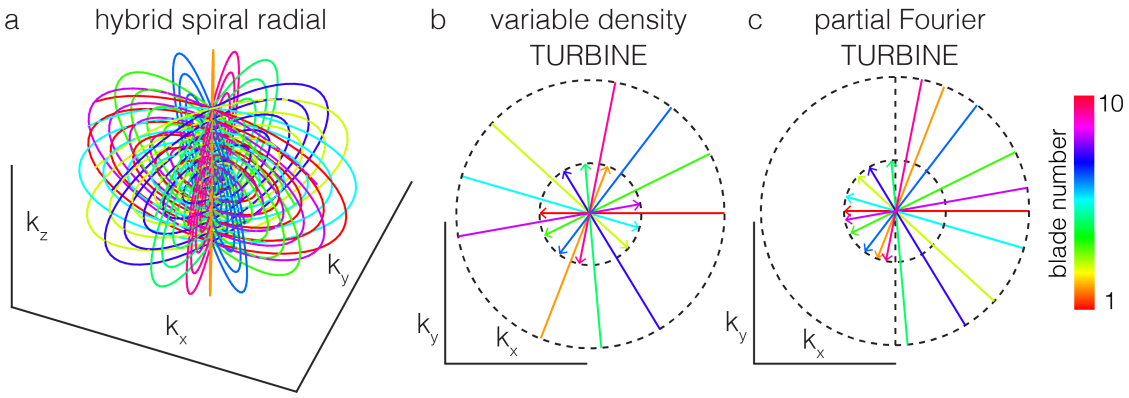


Figure 7.1: Illustration of example TURBINE sequence variants: (a) hybrid-spiral radial acquisition with golden angle update, (b) variable-density TURBINE trajectory (k_x-k_y plane) and (c) partial Fourier TURBINE acquisition (k_x-k_y plane).

It would be useful to explore auto-calibrating variants of the sequence, where a central portion of 3D k-space is always fully sampled. This would likely improve the parallel imaging reconstruction, especially in motion-corrupted data sets [211], as the parallel imaging reconstruction kernel could be re-calibrated for every volume. The TURBINE sequence could be made self-calibrating by sampling all the lines in a region around the k-space centre, resulting in a hybrid radial-EPIK (EPI with keyhole [212]) style acquisition.

Alternatively, the EPI blade could be replaced with a readout trajectory which more naturally oversamples the k-space centre, for example a variable-density spiral using spiral in-out interleaves for optimal TE for BOLD contrast [72]). This would result in a hybrid radial-spiral sequence where the sampled k-space volume is a ball, which could also be interpreted as a 3D extension of the SNAILS [213] approach (see Fig. 7.1a). This

approach would likely be more efficient than the variable density EPI approach suggested above. The down side would be that we would lose the benefits of having one Cartesian direction (for image reconstruction and artefact correction). We think that approaches like the hybrid radial-spiral would only work reliably with B_0 and trajectory correction incorporated into the reconstruction, as discussed below.

Especially at 7 T, there is need to shorten the very long TURBINE echo trains in high-resolution scans in order to achieve optimal BOLD contrast as well as to reduce distortion, off-resonance effects and T_2^* blurring. This could be achieved by using in-plane partial Fourier in addition to parallel imaging, an option which is already available in the current sequence implementation. This could be expanded to perform a radial partial Fourier reconstruction by sampling a smaller radius in one half of k-space (with all the blade starting between 0 and π, see Fig. 7.1c). Another option would be for alternate blades to sample alternate ends of k-space, resulting in a variable density radial distribution (as done in [214] for radial imaging, see Fig. 7.1b). While it is not clear if these approaches would considerably impact fMRI data quality and motion correction fidelity, they are nevertheless worthwhile avenues of exploration.

7.2.2 Image Reconstruction Improvements Using Gradient Characterisation and B_0 Correction

In chapter 5 we examined the image quality of TURBINE data at 3 T and 7 T; especially in the 7 T data, there were a substantial amount of residual artefact. The artefacts could be caused by deviations from the nominal trajectory or stem from B_0 inhomogeneity. We therefore plan to implement an improved reconstruction pipeline, which takes into account an estimate of the true k-space trajectory and corrects for static field inhomogeneity. We have already performed trajectory measurements and characterisation of the gradient system using a field camera. This work and how we are planning to use it for future work is introduced below.

7.2.2.1 Preliminary Work on Trajectory Correction

We are planning to perform trajectory correction by treating the gradient system as a linear time invariant (LTI) system and measuring its gradient impulse response (GIRF). The GIRF of the gradient system is defined as:

$$o(t) = \int_{-\infty}^{\infty} i(\tau) \cdot \mathrm{girf}(t-\tau)\mathrm{d}t \overset{\mathcal{F}}{\longleftrightarrow} O(\omega) = I(\omega) \cdot \mathrm{GIRF}(\omega), \tag{7.1}$$

where $i(\tau)/o(t)$ and $I(\omega)/O(\omega)$ correspond to the input and output of the system in the time and frequency domain respectively. To measure the GIRF, the output of the system needs to be measured for known inputs. We have shown in the past that using GIRFs to predict k-space trajectories for reconstruction results in improved image quality compared to simple gradient correction schemes such, as a gradient delay correction [42, 215]. The GIRF method provides a simple tool for artefact reduction for arbitrary k-trajectories, and depends only on a one-time calibration of the system.

For the GIRF measurement on the FMRIB 7 T scanner we used a dynamic field camera as part of a collaboration project with ETH Zurich (Skope Magnetic Resonance Technologies, Zurich, Switzerland). The field camera consists of 16 transmit-receive ^{1}H NMR probes distributed on a sphere of 10 cm radius. Spherical harmonic basis functions up to 1st order were fitted to the probe phase measurements, yielding the linear x/y/z gradient field terms and the 0^{th} order field, which corresponds to global phase fluctuations. We also used the field camera to monitor the TURBINE trajectory for a range of protocols.

A measured TURBINE trajectory is shown in Fig. 7.2. This example shows how many readout lines miss the nominal k-space centre by a substantial amount. The spread of readout lines was up to $\frac{1}{3} \cdot \Delta k$ in this acquisition, and highly variable depending on the blade angle. Fig. 7.2b shows that the three non phase-encoded navigator lines for Nyquist ghost correction are also shifted perpendicular to the readout direction. This invalidates the assumption that the non-phase encoded odd and even lines in the navigator lines are only separated by a shift in the direction of the readout. These discrepancies

to nominal trajectories might be responsible for some of the remaining image artefacts but the exact impact remains to be investigated.

Unfortunately, the gradient coil used in the above characterisation broke and had to be replaced shortly after the measurements had been performed. However, we will likely have another opportunities to borrow a dynamic field camera; if not we may use a phantom based method (for example the Duyn method [216]) to perform the GIRF measurement of the new gradient hardware.

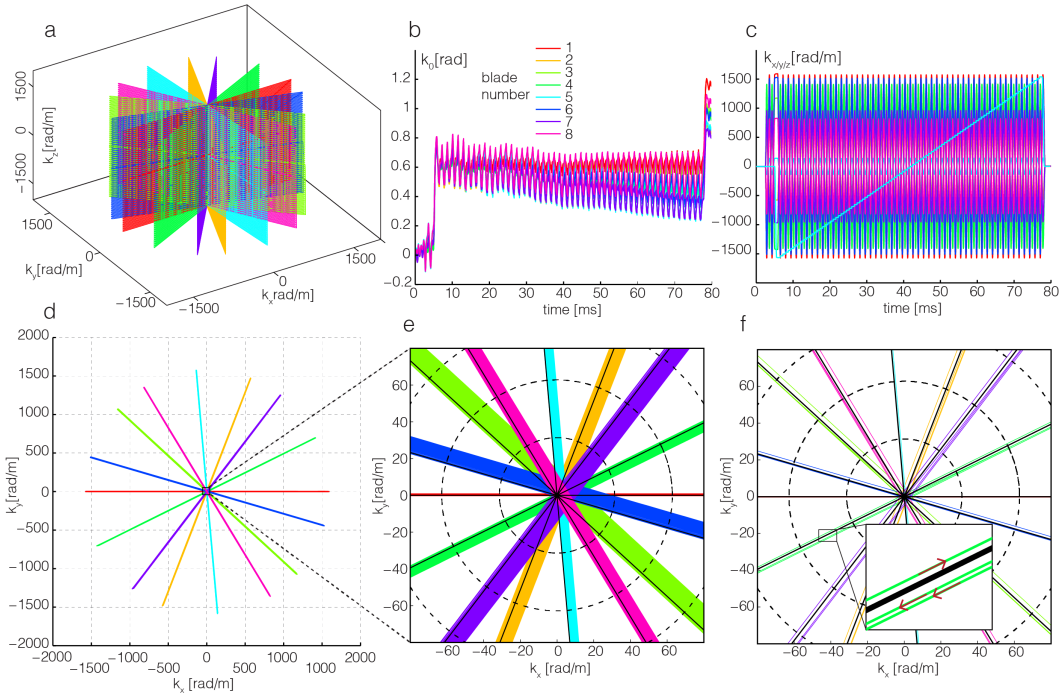


Figure 7.2: Monitored TURBINE trajectory: (a) 3D trajectory as measured with the dynamic field camera. (b) Measured zero-order phase term, (c) first order phase coefficients (k_x, k_y and k_z). (d) view of the radial plane (e) zoom of the centre, the dashed circles are spaced by Δk. (f) Non-phase encoded navigator lines for Nyquist ghost correction. Zoomed inset shows shift between forward and reverse lines perpendicular to the readout.

7.2.2.2 Other Reconstruction Improvements

One of the drawbacks of non-Cartesian imaging is the more complicated off-resonance distortion, which is not straightforward to correct. A number of solutions have been proposed for this problem, including conjugate phase reconstruction [217], or multi-frequency interpolation [218, 219]. We are planning to explore and incorporate B_0

correction into the reconstruction pipeline, to reduce image artefacts.

As mentioned previously, the TURBINE sampling scheme is very well suited for reconstruction schemes, which rely on incoherent sampling, such as compressed sensing, because under-sampling can only be employed in multiple directions. To use TURBINE to its full potential it would be useful to combine the motion correction approach with a k-t acceleration scheme, such as k-t-FASTER [168, 169].

7.2.3 Improved TURBINE Motion Correction

Our proposed motion correction pipeline was able to remove a substantial amount of intra-volume motion, however there is a large potential for improvement.

7.2.3.1 Refinement of the Navigator Approach

Variants of the existing pipeline might improve the motion estimation by incorporating a sliding window approach. A conventional sliding window is likely to result in similar estimates to the temporal interpolation step we already perform in our pipeline, with the additional cost of needing to reconstruct more navigator volumes. A sliding window using a spatio-temporal filter however, for example a temporal hourglass filter (also called tornado temporal filter), has been shown to increase the temporal resolution of motion estimates [209]. Such a filter is designed to be narrow in the oversampled central region of k-space, using only very few blades. It increases smoothly towards higher spatial frequencies including information from more blades. The higher spatial frequencies are therefore temporally averaged over a longer window, but the reduced under-sampling reduces radial streaking artefacts, which can reduce the accuracy of the registration step to determine the motion estimates.

7.2.3.2 Motion Correction Within the Image Reconstruction

It would be useful to incorporate the motion estimation directly into the reconstruction model. This means that we effectively need to estimate 6 additional parameters for each shot in the image acquisition within the image reconstruction. A few groups have explored approaches like this recently [206, 220] using either an alternating or joint minimisation approach to find the reconstructed image and the motion parameters. We are planning to collaborate with researchers from Kings College, London to test their recently developed motion correction approach [206] on TURBINE data.

7.2.3.3 Motion Camera

We temporarily have access to an optical motion camera on the 7 T scanner at FMRIB. Using this would be beneficial for a number of reasons: we could use the data from the motion camera to provide a ground truth motion estimate and evaluate our approach against it in order to choose the best TURBINE motion correction pipeline. This would also allow evaluation of our method against a state-of-the-art prospective correction.

Head motion is heavily constrained by the geometry of the head and its placement within the receive coil. Potentially, studying data from a motion camera would allow a model of typical patient motion to be built, or at least the determination of physiological limits (in maximum motion and maximum speed of motion) for possible motions. This could potentially help constrain motion estimates, improving the conditioning of the motion estimation problem and faster convergence of the motion estimation algorithms discussed above.

7.2.3.4 Static B_0 and Shimming in the Presence of Motion

The field distribution around air-tissue interfaces depends on the geometry of those interfaces and their orientation with respect to the main magnetic field. Subject movement can cause nonlinear changes in this field distribution and its gradients. This is likely a

source of error in our data. As mentioned above we are planning to incorporate static B_0 correction into the image reconstruction, and it would be useful to be able to update the B_0 estimate after motion has occurred. As a first step we want to test if a simple adjustment, such as rotating the B_0 map (acquired before the fMRI experiment) according to the motion determined with the TURBINE navigators can improve the final reconstructions.

Another challenge is the shim settings. Potentially, this could be approached in a similar way by using dynamic shimming based on a rotated field map (using the field map collected before the functional scan and updating it according to the navigators). However, this approach would require the ability to reconstruct the navigator images and estimate the motion in near-real time.

7.3 Final Remarks

Since its invention over 25 years ago, fMRI has made remarkable progress establishing itself as the primary method for mapping human brain function. The rapid advancement of the field has been driven by a unique interplay between engineering and neuroscience. The field currently faces a number of challenges. These include the need for higher spatial and temporal resolution, as well as ensuring reproducibility and robust single-subject results for clinical adaptation.

The work in this thesis aims to contribute to methods development in fMRI. The goal of the work was to investigate 3D EPI methods, with a focus on the implementation and testing of the TURBINE sequence for fMRI. As the core result we were able to show that TURBINE motion correction reduces image corruption and improves fMRI results in the presence of motion. We hope that, especially after implementing some of the improvements laid out above, TURBINE will contribute towards making fMRI acquisition more flexible and robust.

It is a fascinating time to be part of fMRI research and it will be interesting to see which methods will shape the fMRI experiment of the future.

References

[1] M. A. Bernstein, K. F. King, and X. J. Zhou. *Handbook of MRI Pulse Sequences*. London, UK: Elsevier Academic Press, 2004.

[2] R. W. Brown, Y. C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging. Physical Principles and Sequence Design*. Chichester, UK: John Wiley & Sons, 2014.

[3] D. G. Nishimura. *Principles of magnetic resonance imaging*. Stanford, CA: Stanford University Press, 1996.

[4] F. Bloch. "Nuclear Induction". *Physical Review* 70.7-8 (1946), pp. 460–474.

[5] E. M. Purcell, H. C. Torrey, and R. V. Pound. "Resonance Absorption by Nuclear Magnetic Moments in a Solid." *Physical Review* 69.1-2 (1946), pp. 37–38.

[6] P. C. Lauterbur. "Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance." *Nature* 242.5394 (1973), pp. 190–191.

[7] A. Kumar, D. Welti, and R. R. Ernst. "NMR Fourier zeugmatography. 1975." *Journal of Magnetic Resonance* 18 (1975), pp. 69–83.

[8] P. Mansfield. "Multi-planar image formation using NMR spin echoes." *Journal of Physics C: Solid State Physics* 10.3 (1977), pp. L55–L58.

[9] T. R. Brown, B. M. Kincaid, and K. Ugurbil. "NMR chemical shift imaging in three dimensions." *Proceedings of the National Academy of Sciences* 79.11 (1982), pp. 3523–3526.

[10] S. Ljunggren. "A simple graphical representation of fourier-based imaging methods." *Journal of Magnetic Resonance (1969)* 54.2 (1983), pp. 338–343.

[11] D. B. Twieg. "The k-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods." *Medical Physics* 10.5 (1983), p. 610.

[12] D. A. Feinberg, J. D. Hale, J. C. Watts, L. Kaufman, and A. Mark. "Halving MR imaging time by conjugation: demonstration at 3.5 kG." *Radiology* 161.2 (1986), pp. 527–531.

[13] W. A. Edelstein, J. M. Hutchison, G. Johnson, and T. Redpath. "Spin warp NMR imaging and applications to human whole-body imaging." *Physics in Medicine and Biology* 25.4 (1980), pp. 751–756.

[14] C. M. Tsai and D. G. Nishimura. "Reduced aliasing artifacts using variable-density k-space sampling trajectories." *Magnetic Resonance in Medicine* 43.3 (2000), pp. 452–458.

[15] Y. Zur, M. L. Wood, and L. J. Neuringer. "Spoiling of transverse magnetization in steady-state sequences." *Magnetic Resonance in Medicine* 21.2 (1991), pp. 251–263.

[16] W. Lin and H. K. Song. "Improved signal spoiling in fast radial gradient-echo imaging: Applied to accurate T1 mapping and flip angle correction". *Magnetic Resonance in Medicine* 62.5 (2009), pp. 1185–1194.

[17] R. R. Ernst and W. A. Anderson. "Application of Fourier Transform Spectroscopy to Magnetic Resonance." *Review of Scientific Instruments* 37.1 (1966), pp. 93–102.

[18] J. W. Cooley and J. W. Tukey. "An Algorithm for the Machine Calculation of Complex Fourier Series." *Mathematics of Computation* 19.90 (1965), pp. 297–301.

[19] J. I. Jackson, C. H. Meyer, and D. G. Nishimura. "Selection of a convolution function for Fourier inversion using gridding [computerised tomography application]." *IEEE Transactions on Medical Imaging* 10.3 (1991), pp. 473–478.

[20] J. D. O'Sullivan. "A fast sinc function gridding algorithm for Fourier inversion in computer tomography." *IEEE Transactions on Medical Imaging* 4.4 (1985), pp. 200–207.

[21] P. J. Beatty, D. G. Nishimura, and J. M. Pauly. "Rapid gridding reconstruction with a minimal oversampling ratio." *IEEE Transactions on Medical Imaging* 24.6 (2005), pp. 799–808.

[22] J. A. Fessler and B. P. Sutton. "Nonuniform fast fourier transforms using min-max interpolation." *IEEE Transactions on Signal Processing* 51.2 (2003), pp. 560–574.

[23] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. "The NMR phased array." *Magnetic Resonance in Medicine* 16.2 (1990), pp. 192–225.

[24] M. Blaimer, F. Breuer, M. Mueller, R. M. Heidemann, M. A. Griswold, and P. M. Jakob. "SMASH, SENSE, PILS, GRAPPA: How to Choose the Optimal Method." *Topics in Magnetic Resonance Imaging* 15.4 (2004), pp. 223–236.

[25] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. "SENSE: sensitivity encoding for fast MRI." *Magnetic Resonance in Medicine* 42.5 (1999), pp. 952–962.

[26] D. K. Sodickson and W. J. Manning. "Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays." *Magnetic Resonance in Medicine* 38.4 (1997), pp. 591–603.

[27] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. "Generalized autocalibrating partially parallel acquisitions (GRAPPA)." *Magnetic Resonance in Medicine* 47.6 (2002), pp. 1202–1210.

[28] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig. "ESPIRiT–an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA." *Magnetic Resonance in Medicine* 71.3 (2014), pp. 990–1001.

[29] W. S. Hoge, D. H. Brooks, B. Madore, and W. E. Kyriakos. "A tour of accelerated parallel MR imaging from a linear systems perspective." *Concepts in Magnetic Resonance Part A* 27A.1 (2005), pp. 17–37.

[30] N. Seiberlich, P. Ehses, J. Duerk, R. Gilkeson, and M. Griswold. "Improved radial GRAPPA calibration for real-time free-breathing cardiac imaging." *Magnetic Resonance in Medicine* 65.2 (2011), pp. 492–505.

[31] K. L. Wright, J. I. Hamilton, M. A. Griswold, V. Gulani, and N. Seiberlich. "Non-Cartesian parallel imaging reconstruction." *Journal of Magnetic Resonance Imaging* 40.5 (2014), pp. 1022–1040.

[32] M. Lustig and J. M. Pauly. "SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space." *Magnetic Resonance in Medicine* 64.2 (2010), pp. 457–471.

[33] A. Haase, J. Frahm, W. Hanicke, and D. Matthaei. "^{1}H NMR chemical shift selective (CHESS) imaging." *Physics in Medicine and Biology* 30.4 (1985), pp. 341–344.

[34] P. J. Hors. "A new method for water suppression in the proton NMR spectra of aqueous solutions." *Journal of Magnetic Resonance* 54.3 (1983), pp. 539–542.

[35] P. Jezzard and S. Clare. "Sources of distortion in functional MRI data." *Human Brain Mapping* 8.2-3 (1999), pp. 80–85.

[36] P. Jezzard and R. S. Balaban. "Correction for geometric distortion in echo planar images from B_0 field variations." *Magnetic Resonance in Medicine* 34.1 (1995), pp. 65–73.

[37] P. Jezzard. "Correction of geometric distortion in fMRI data." *NeuroImage* 62.2 (2012), pp. 648–651.

[38] C. Boesch, R. Gruetter, and E. Martin. "Temporal and spatial analysis of fields generated by eddy currents in superconducting magnets: Optimization of corrections and quantitative characterization of magnet/gradient systems." *Magnetic Resonance in Medicine* 20.2 (1991), pp. 268–284.

[39] Y. Wu, B. A. Chronik, C. Bowen, C. K. Mechefske, and B. K. Rutt. "Gradient-induced acoustic and magnetic field fluctuations in a 4T whole-body MR imager." *Magnetic Resonance in Medicine* 44.4 (2000), pp. 532–536.

[40] S. J. Vannesjo, M. Haeberlin, L. Kasper, M. Pavan, B. J. Wilm, C. Barmet, and K. P. Pruessmann. "Gradient system characterization by impulse response measurements with a dynamic field camera." *Magnetic Resonance in Medicine* 69.2 (2013), pp. 583–593.

[41] S. J. Vannesjo, B. E. Dietrich, M. Pavan, D. O. Brunner, B. J. Wilm, C. Barmet, and K. P. Pruessmann. "Field camera measurements of gradient and shim impulse responses using frequency sweeps." *Magnetic Resonance in Medicine* 72.2 (2014), pp. 570–583.

[42] S. J. Vannesjo, N. N. Graedel, L. Kasper, S. Gross, J. Busch, M. Haeberlin, C. Barmet, and K. P. Pruessmann. "Image reconstruction using a gradient impulse response model for trajectory prediction." *Magnetic Resonance in Medicine* 76.1 (2016), pp. 45–58.

[43] F. Schmitt, M. K. Stehling, and R. Turner. *Echo-Planar Imaging. Theory, Technique and Application.* Theory, Technique and Application. Heidelberg, Germany: Springer, 1998.

[44] P. A. Bandettini. "What's New in Neuroimaging Methods?" *Annals of the New York Academy of Sciences* 1156.1 (2009), pp. 260–293.

[45] D. G. Norris. "Principles of magnetic resonance assessment of brain function." *Journal of Magnetic Resonance Imaging* 23.6 (2006), pp. 794–807.

[46] H. Lu, X. Golay, J. J. Pekar, and P. C. M. Van Zijl. "Functional magnetic resonance imaging based on changes in vascular space occupancy." *Magnetic Resonance in Medicine* 50.2 (2003), pp. 263–274.

[47] J. A. Detre, J. S. Leigh, D. S. Williams, and A. P. Koretsky. "Perfusion imaging". *Magnetic Resonance in Medicine* 23.1 (1992), pp. 37–45.

[48] D. S. Williams, J. A. Detre, J. S. Leigh, and A. P. Koretsky. "Magnetic resonance imaging of perfusion using spin inversion of arterial water." *Proceedings of the National Academy of Sciences* 89.1 (1992), pp. 212–216.

[49] N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann. "Neurophysiological investigation of the basis of the fMRI signal." *Nature* 412.6843 (2001), pp. 150–157.

[50] L. Pauling and C. D. Coryell. "The Magnetic Properties and Structure of Hemoglobin, Oxyhemoglobin and Carbonmonoxyhemoglobin." *Proceedings of the National Academy of Sciences* 22.4 (1936), pp. 210–216.

[51] K. R. Thulborn, J. C. Waterton, P. M. Matthews, and G. K. Radda. "Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field." *Biochimica Et Biophysica Acta* 714.2 (1982), pp. 265–270.

[52] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. "Brain magnetic resonance imaging with contrast dependent on blood oxygenation." *Proceedings of the National Academy of Sciences* 87.24 (1990), pp. 9868–9872.

[53] S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S. G. Kim, H. Merkle, and K. Ugurbil. "Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging." *Proceedings of the National Academy of Sciences* 89.13 (1992), pp. 5951–5955.

[54] K. K. Kwong, J. W. Belliveau, D. A. Chesler, I. E. Goldberg, R. M. Weisskoff, B. P. Poncelet, D. N. Kennedy, B. E. Hoppel, M. S. Cohen, and R. Turner. "Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation." *Proceedings of the National Academy of Sciences* 89.12 (1992), pp. 5675–5679.

[55] D. A. Yablonskiy and E. M. Haacke. "Theory of NMR signal behavior in magnetically inhomogeneous tissues: The static dephasing regime". *Magnetic Resonance in Medicine* 32.6 (1994), pp. 749–763.

[56] E. M. Haacke, S. Lai, D. A. Yablonskiy, and W. Lin. "In vivo validation of the bold mechanism: A review of signal changes in gradient echo functional MRI in the presence of flow." *International Journal of Imaging Systems and Technology* 6.2-3 (1995), pp. 153–163.

[57] R. B. Buxton, K. Uludağ, D. J. Dubowitz, and T. T. Liu. "Modeling the hemodynamic response to brain activation." *NeuroImage* 23 (2004), S220–S233.

[58] R. B. Buxton, E. C. Wong, and L. R. Frank. "Dynamics of blood flow and oxygenation changes during brain activation: the balloon model." *Magnetic Resonance in Medicine* 39.6 (1998), pp. 855–864.

[59] K. Setsompop, D. A. Feinberg, and J. R. Polimeni. "Rapid brain MRI acquisition techniques at ultra-high fields." *NMR in Biomedicine* 29.9 (2016), pp. 1198–1221.

[60] K. Scheffler and P. Ehses. "High-resolution mapping of neuronal activation with balanced SSFP at 9.4 Tesla." *Magnetic Resonance in Medicine* 76.1 (2015).

[61] T. Schild, G. Aubert, J. Belorgey, C. Berriaud, A. Bourquard, P. Bredy, O. Dubois, G. Gilgrass, F. P. Juster, H. Lannou, F. Molinie, M. Nusbaum, F. Nunio, A. Payn, L. Quettier, L. Scola, A. Sinanna, V. Stepanov, and P. Vedrine. "Iseult/INUMAC 11.7-T MRI Assembly Status." *IEEE Transactions on Applied Superconductivity* 26.4 (2016), pp. 1–4.

[62] L. DelaBarre, R. L. Lagore, Y. Eryaman, G. Adriany, and J. T. Vaughan. "Porcine Imaging in a 10.5T Whole-Body Human MRI." *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 541.

[63] S. Ogawa, R. S. Menon, D. W. Tank, S. G. Kim, H. Merkle, J. M. Ellermann, and K. Ugurbil. "Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model." *Biophysical Journal* 64.3 (1993), pp. 803–812.

[64] J. Gati, R. Menon, K. Ugurbil, and B. Rutt. "Experimental determination ofthe BOLD field strength dependence in vessels and tissue." *Magnetic Resonance in Medicine* 38 (1997), pp. 296–302.

[65] R. Turner. "How much cortex can a vein drain? Downstream dilution of activation-related cerebral blood oxygenation changes." *NeuroImage* 16.4 (2002), pp. 1062–1067.

[66] R. Boyacioğlu, J. Schulz, N. C. Müller, P. J. Koopmans, M. Barth, and D. G. Norris. "Whole brain, high resolution multiband spin-echo EPI fMRI at 7 T: A comparison with gradient-echo EPI using a color-word Stroop task." *NeuroImage* 97 (2014), pp. 142 –150.

[67] G. Krüger and G. H. Glover. "Physiological noise in oxygenation-sensitive magnetic resonance imaging." *Magnetic Resonance in Medicine* 46.4 (2001), pp. 631–637.

[68] W. A. Edelstein, G. H. Glover, C. J. Hardy, and R. W. Redington. "The intrinsic signal-to-noise ratio in NMR imaging". *Magnetic Resonance in Medicine* 3.4 (1986), pp. 604–618.

[69] A. Shmuel, E. Yacoub, D. Chaimow, N. K. Logothetis, and K. Ugurbil. "Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla." *NeuroImage* 35.2 (2007), pp. 539–552.

[70] P. A. Bandettini, E. C. Wong, R. S. Hinks, R. S. Tikofsky, and J. S. Hyde. "Time course EPI of human brain function during task activation." *Magnetic Resonance in Medicine* 25.2 (1992), pp. 390–397.

[71] A. M. Blamire, S. Ogawa, K. Ugurbil, D. Rothman, G. McCarthy, J. M. Ellermann, F. Hyder, Z. Rattner, and R. G. Shulman. "Dynamic mapping of the human visual cortex by high-speed magnetic resonance imaging." *Proceedings of the National Academy of Sciences* 89.22 (1992), pp. 11069–11073.

[72] G. H. Glover and C. S. Law. "Spiral-in/out BOLD fMRI for increased SNR and reduced susceptibility artifacts." *Magnetic Resonance in Medicine* 46.3 (2001), pp. 515–522.

[73] G. H. Glover and S. Lai. "Self-navigated spiral fMRI: Interleaved versus single-shot." *Magnetic Resonance in Medicine* 39.3 (1998), pp. 361–368.

[74] D. A. Feinberg, S. Moeller, S. M. Smith, E. Auerbach, S. Ramanna, M. Gunther, M. F. Glasser, K. L. Miller, K. Ugurbil, and E. Yacoub. "Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging." *PloS ONE* 5.12 (2010), e15710.

[75] S. M. Smith, K. L. Miller, S. Moeller, J. Xu, E. J. Auerbach, M. W. Woolrich, C. F. Beckmann, M. Jenkinson, J. Andersson, M. F. Glasser, D. C. Van Essen, D. A. Feinberg, E. S. Yacoub, and K. Ugurbil. "Temporally-independent functional modes of spontaneous brain activity." *Proceedings of the National Academy of Sciences of the United States of America* 109.8 (2012), pp. 3131–3136.

[76] H.-L. Lee, B. Zahneisen, T. Hugger, P. LeVan, and J. Hennig. "Tracking dynamic resting-state networks at higher frequencies using MR-encephalography." *NeuroImage* 65 (2013), pp. 216–222.

[77] E. Yacoub, N. Harel, and K. Ugurbil. "High-field fMRI unveils orientation columns in humans." *Proceedings of the National Academy of Sciences of the United States of America* 105.30 (2008), pp. 10607–10612.

[78] J. R. Polimeni, B. Fischl, D. N. Greve, and L. L. Wald. "Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1." *NeuroImage* 52.4 (2010), pp. 1334–1346.

[79] A. Devor, S. Sakadžić, V. J. Srinivasan, M. A. Yaseen, K. Nizar, P. A. Saisan, P. Tian, A. M. Dale, S. A. Vinogradov, M. A. Franceschini, and D. A. Boas. "Frontiers in optical imaging of cerebral blood flow and metabolism." *Journal of Cerebral Blood Flow and Metabolism* 32.7 (2012), pp. 1259–1276.

[80] D. J. Larkman, J. V. Hajnal, A. H. Herlihy, G. A. Coutts, I. R. Young, and G. Ehnholm. "Use of multicoil arrays for separation of signal from multiple slices simultaneously excited." *Journal of Magnetic Resonance Imaging* 13.2 (2001), pp. 313–317.

[81] B. A. Poser, P. J. Koopmans, T. Witzel, L. L. Wald, and M. Barth. "Three dimensional echo-planar imaging at 7 Tesla." *NeuroImage* 51.1 (2010), pp. 261–266.

[82] F. A. Breuer, S. A. R. Kannengiesser, M. Blaimer, N. Seiberlich, P. M. Jakob, and M. A. Griswold. "General formulation for quantitative G-factor calculation in GRAPPA reconstructions." *Magnetic Resonance in Medicine* 62.3 (2009), pp. 739–746.

[83] F. A. Breuer, M. Blaimer, R. M. Heidemann, M. F. Mueller, M. A. Griswold, and P. M. Jakob. "Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging." *Magnetic Resonance in Medicine* 53.3 (2005), pp. 684–691.

[84] B. A. Poser, V. G. Kemper, D. Ivanov, and K. Uludag. "CAIPIRINHA-accelerated 3D EPI for high temporal and/or spatial resolution EPI acquisitions." *Proceedings of the ESMRMB Congress*. Toulouse, France, 2013, p. 287.

[85] M. Narsude, D. Gallichan, W. van der Zwaag, R. Gruetter, and J. P. Marques. "Three-dimensional echo planar imaging with controlled aliasing: A sequence for high temporal resolution functional MRI." *Magnetic Resonance in Medicine* 75.6 (2016), pp. 2350–2361.

[86] A. C. Silva, E. L. Barbier, I. J. Lowe, and A. P. Koretsky. "Radial Echo-Planar Imaging". *Journal of Magnetic Resonance* 135.1 (1998), pp. 242–247.

[87] M. Krämer, T. H. Jochimsen, and J. R. Reichenbach. "Functional magnetic resonance imaging using PROPELLER-EPI." *Magnetic Resonance in Medicine* 68.1 (2012), pp. 140–151.

[88] W. van der Zwaag, S. Francis, and R. Bowtell. "Improved echo volumar imaging (EVI) for functional MRI." *Magnetic Resonance in Medicine* 56.6 (2006), pp. 1320–1327.

[89] S. Posse, E. Ackley, R. Mutihac, J. Rick, M. Shane, C. Murray-Krezan, M. Zaitsev, and O. Speck. "Enhancement of temporal resolution and BOLD sensitivity in real-time fMRI using multi-slab echo-volumar imaging." *NeuroImage* 61.1 (2012), pp. 115–130.

[90] G. R. Lee, M. A. Griswold, and J. A. Tkach. "Rapid 3D radial multi-echo functional magnetic resonance imaging." *NeuroImage* 52.4 (2010), pp. 1428–1443.

[91] Y. Hu and G. H. Glover. "Three-dimensional spiral technique for high-resolution functional MRI." *Magnetic Resonance in Medicine* 58.5 (2007), pp. 947–951.

[92] W. Deng, B. Zahneisen, and V. A. Stenger. "Rotated stack-of-spirals partial acquisition for rapid volumetric parallel MRI." *Magnetic Resonance in Medicine* 76.1 (2016), pp. 127–135.

[93] B. Zahneisen, T. Grotz, K. J. Lee, S. Ohlendorf, M. Reisert, M. Zaitsev, and J. Hennig. "Three-dimensional MR-encephalography: Fast volumetric brain imaging using rosette trajectories." *Magnetic Resonance in Medicine* 65.5 (2011), pp. 1260–1268.

[94] B. Zahneisen, T. Hugger, K. J. Lee, P. LeVan, M. Reisert, H.-L. Lee, J. Assländer, M. Zaitsev, and J. Hennig. "Single shot concentric shells trajectories for ultra fast fMRI." *Magnetic Resonance in Medicine* 68.2 (2012), pp. 484–494.

[95] J. A. McNab, D. Gallichan, and K. L. Miller. "3D steady-state diffusion-weighted imaging with trajectory using radially batched internal navigator echoes (TURBINE)." *Magnetic Resonance in Medicine* 63.1 (2010), pp. 235–242.

[96] D. C. V. Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, K. Ugurbil, and the WU-Minn HCP Consortium. "The WU-Minn Human Connectome Project: An overview." *NeuroImage* 80 (2013), pp. 62–79.

[97] T. Stoecker. "Big Data: The Rhineland Study." *24th Annual Meeting of the ISMRM*. Talk 09/05/2016. Singapore, 2016.

[98] R. Stirnberg, D. Brenner, T. Stöcker, and N. J. Shah. "Rapid fat suppression for three-dimensional echo planar imaging with minimized specific absorption rate." *Magnetic Resonance in Medicine* (2015). In press.

[99] P. Kragel, R. M. Carter, and S. Huettel. "What Makes a Pattern? Matching Decoding Methods to Data in Multivariate Pattern Analysis." *Frontiers in Neuroscience* 6 (2012), p. 162.

[100] M. McKeown, S. Makeig, G. Brown, T. Jung, S. Kindermann, A. Bell, and T. Sejnowski. "Analysis of fMRI data by blind separation into independent spatial components." *Human Brain Mapping* 6.3 (1998), pp. 160–188.

[101] M. S. Dagli, J. E. Ingeholm, and J. V. Haxby. "Localization of Cardiac-Induced Signal Change in fMRI." *NeuroImage* 9.4 (1999), pp. 407–415.

[102] C. Windischberger, H. Langenberger, T. Sycha, E. M. Tschernko, G. Fuchsjäger-Mayerl, L. Schmetterer, and E. Moser. "On the origin of respiratory artifacts in BOLD-EPI of the human brain." *Magnetic Resonance Imaging* 20.8 (2002), pp. 575–582.

[103] D. Raj, A. W. Anderson, and J. C. Gore. "Respiratory effects in human functional magnetic resonance imaging due to bulk susceptibility changes." *Physics in Medicine and Biology* 46.12 (2001), pp. 3331–3340.

[104] R. G. Wise, K. Ide, M. J. Poulin, and I. Tracey. "Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal." *NeuroImage* 21.4 (2004), pp. 1652–1664.

[105] R. M. Birn, J. B. Diamond, M. A. Smith, and P. A. Bandettini. "Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI." *NeuroImage* 31.4 (2006), pp. 1536–1548.

[106] R. M. Birn. "The role of physiological noise in resting-state functional connectivity." *NeuroImage* 62.2 (2012), pp. 864–870.

[107] B. Biswal, E. A. Deyoe, and J. S. Hyde. "Reduction of physiological fluctuations in fMRI using digital filters." *Magnetic Resonance in Medicine* 35.1 (1996), pp. 107–113.

[108] G. H. Glover, T. Q. Li, and D. Ress. "Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR." *Magnetic Resonance in Medicine* 44.1 (2000), pp. 162–167.

[109] X. Hu, T. H. Le, T. Parrish, and P. Erhard. "Retrospective estimation and correction of physiological fluctuation in functional MRI." *Magnetic Resonance in Medicine* 34.2 (1995), pp. 201–212.

[110] K. Shmueli, P. van Gelderen, J. A. de Zwart, S. G. Horovitz, M. Fukunaga, J. M. Jansma, and J. H. Duyn. "Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal." *NeuroImage* 38.2 (2007), pp. 306–320.

[111] C. Hutton, O. Josephs, J. Stadler, E. Featherstone, A. Reid, O. Speck, J. Bernarding, and N. Weiskopf. "The impact of physiological noise correction on fMRI at 7T." *NeuroImage* 57.1 (2011), pp. 101–112.

[112] A. Lutti, D. L. Thomas, C. Hutton, and N. Weiskopf. "High-resolution functional MRI at 3 T: 3D/2D echo-planar imaging with optimized physiological noise correction." *Magnetic Resonance in Medicine* 69.6 (2013), pp. 1657–1664.

[113] J. Jorge, P. Figueiredo, W. van der Zwaag, and J. P. Marques. "Signal fluctuations in fMRI data acquired with 2D-EPI and 3D-EPI at 7 Tesla." *Magnetic Resonance Imaging* 31.2 (2013), pp. 212–220.

[114] R. H. N. Tijssen, M. Jenkinson, J. C. W. Brooks, P. Jezzard, and K. L. Miller. "Optimizing RetroICor and RetroKCor corrections for multi-shot 3D FMRI acquisitions." *NeuroImage* 84 (2014), pp. 394–405.

[115] S. Särkkä, A. Solin, A. Nummenmaa, A. Vehtari, T. Auranen, S. Vanni, and H. Lin. "Dynamic retrospective filtering of physiological noise in BOLD fMRI: DRIFTER." *NeuroImage* 60.2 (2012), pp. 1517–1527.

[116] C. G. Thomas, R. A. Harshman, and R. S. Menon. "Noise reduction in BOLD-based fMRI using component analysis." *NeuroImage* 17.3 (2002), pp. 1521–1537.

[117] X. Hu and S.-G. Kim. "Reduction of signal fluctuation in functional MRI using navigator echoes." *Magnetic Resonance in Medicine* 31.5 (1994), pp. 495–503.

[118] E. Durand, P.-F. van de Moortele, M. Pachot-Clouard, and D. Le Bihan. "Artifact Due to B0 Fluctuations in fMRI: Correction Using the k-Space Central Line." *Magnetic Resonance in Medicine* 46.1 (2001), pp. 198–201.

[119] J. Pfeuffer, P.-F. van de Moortele, K. Ugurbil, X. Hu, and G. H. Glover. "Correction of physiologically induced global off-resonance effects in dynamic echo-planar and spiral functional imaging." *Magnetic Resonance in Medicine* 47.2 (2002), pp. 344–353.

[120] R. H. N. Tijssen, T. W. Okell, and K. L. Miller. "Real-time cardiac synchronization with fixed volume frame rate for reducing physiological instabilities in 3D FMRI." *NeuroImage* 57.4 (2011), pp. 1364–1375.

[121] P. Kundu, S. J. Inati, J. W. Evans, W.-M. Luh, and P. A. Bandettini. "Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI." *NeuroImage* 60.3 (2012), pp. 1759–1770.

[122] K. J. Friston, S. Williams, R. Howard, R. S. Frackowiak, and R. Turner. "Movement-related effects in fMRI time-series." *Magnetic Resonance in Medicine* 35.3 (1996), pp. 346–355.

[123] J. V. Hajnal, R. Myers, A. Oatridge, J. E. Schwieso, I. R. Young, and G. M. Bydder. "Artifacts due to stimulus correlated motion in functional imaging of the brain." *Magnetic Resonance in Medicine* 31.3 (1994), pp. 283–291.

[124] M. Soellinger, A. K. Rutz, S. Kozerke, and P. Boesiger. "3D cine displacement-encoded MRI of pulsatile brain motion." *Magnetic Resonance in Medicine* 61.1 (2009), pp. 153–162.

[125] G. Salimi-Khorshidi, G. Douaud, C. F. Beckmann, M. F. Glasser, L. Griffanti, and S. M. Smith. "Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers." *NeuroImage* 90 (2014), pp. 449 –468.

[126] L. Griffanti, G. Salimi-Khorshidi, C. F. Beckmann, E. J. Auerbach, G. Douaud, C. E. Sexton, E. Zsoldos, K. P. Ebmeier, N. Filippini, C. E. Mackay, S. Moeller, J. Xu, E. Yacoub, G. Baselli, K. Ugurbil, K. L. Miller, and S. M. Smith. "ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging." *NeuroImage* 95 (2014), pp. 232 –247.

[127] E. Seto, G. Sela, W. E. McIlroy, S. E. Black, W. R. Staines, M. J. Bronskill, A. R. McIntosh, and S. J. Graham. "Quantifying head motion associated with motor tasks used in fMRI." *NeuroImage* 14.2 (2001), pp. 284–297.

[128] Y. Duerst, B. J. Wilm, B. E. Dietrich, S. J. Vannesjo, C. Barmet, T. Schmid, D. O. Brunner, and K. P. Pruessmann. "Real-time feedback for spatiotemporal field stabilization in MR systems." *Magnetic Resonance in Medicine* 73.2 (2014), pp. 884–893.

[129] Q. Tian, E. Gong, C. W. Leuze, J. Pauly, and J. McNab. "Effects of Motion on Coupling of Coil Elements and Parallel Imaging Reconstruction at 3T and 7T." *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, Canada, 2015, p. 2419.

[130] G. H. Glover and J. M. Pauly. "Projection Reconstruction Techniques for Reduction of Motion Effects in MRI." *Magnetic Resonance in Medicine* 28.2 (1992), pp. 275–289.

[131] F. Godenschweger, U. Kägebein, D. Stucht, U. Yarach, A. Sciarra, R. Yakupov, F. Lüsebrink, P. Schulze, and O. Speck. "Motion correction in MRI of the brain." *Physics in Medicine and Biology* 61.5 (2016), R32–56.

[132] M. Zaitsev, J. Maclaren, and M. Herbst. "Motion artifacts in MRI: A complex problem with many partial solutions." *Journal of Magnetic Resonance Imaging* 42.4 (2015), pp. 887–901.

[133] M. Jenkinson, P. Bannister, M. Brady, and S. Smith. "Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images." *NeuroImage* 17.2 (2002), pp. 825–841.

[134] J. D. Power, K. A. Barnes, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen. "Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion." *NeuroImage* 59.3 (2012), pp. 2142–2154.

[135] M. D. Tisdall, A. T. Hess, M. Reuter, E. M. Meintjes, B. Fischl, and A. J. W. van der Kouwe. "Volumetric navigators for prospective motion correction and selective reacquisition in neuroanatomical MRI." *Magnetic Resonance in Medicine* 68.2 (2011), pp. 389–399.

[136] J. Maclaren, M. Herbst, O. Speck, and M. Zaitsev. "Prospective motion correction in brain imaging: a review." *Magnetic Resonance in Medicine* 69.3 (2013), pp. 621–636.

[137] O. Speck, J. Hennig, and M. Zaitsev. "Prospective real-time slice-by-slice motion correction for fMRI in freely moving subjects." *Magnetic Resonance Materials in Physics, Biology and Medicine* 19.2 (2006), pp. 55–61.

[138] C. C. Lee, C. R. Jack, R. C. Grimm, P. J. Rossman, J. P. Felmlee, R. L. Ehman, and S. J. Riederer. "Real-time adaptive motion correction in functional MRI." *Magnetic Resonance in Medicine* 36.3 (1996), pp. 436–444.

[139] E. B. Welch, A. Manduca, R. C. Grimm, H. A. Ward, and C. R. Jack Jr. "Spherical navigator echoes for full 3D rigid body motion measurement in MRI." *Magnetic Resonance in Medicine* 47.1 (2002), pp. 32–41.

[140] N. F. Ramsey, J. S. Van den Brink, A. M. C. van Muiswinkel, P. J. M. Folkers, C. T. W. Moonen, J. M. Jansma, and R. S. Kahn. "Phase Navigator Correction in 3D fMRI Improves Detection of Brain Activation: Quantitative Assessment with a Graded Motor Activation Procedure." *NeuroImage* 8.3 (1998), pp. 240–248.

[141] D. Atkinson, D. Hill, and P. Stoyle. "Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion." *IEEE Transactions on Medical imaging* 16.6 (1997), pp. 903–910.

[142] W. van der Zwaag, S. Francis, K. Head, A. Peters, P. Gowland, P. Morris, and R. Bowtell. "fMRI at 1.5, 3 and 7 T: characterising BOLD signal changes." *NeuroImage* 47.4 (2009), pp. 1425–1434.

[143] W. van der Zwaag, J. P. Marques, T. Kober, G. Glover, R. Gruetter, and G. Krueger. "Temporal SNR characteristics in segmented 3D-EPI at 7T." *Magnetic Resonance in Medicine* 67.2 (2012), pp. 344–352.

[144] A. Kristoffersen and P.-E. Goa. "Cardiac-induced physiological noise in 3D gradient echo brain imaging: effect of k-space sampling scheme." *Journal of Magnetic Resonance* 212.1 (2011), pp. 74–85.

[145] U. Goerke, H. E. Möller, D. G. Norris, and C. Schwarzbauer. "A comparison of signal instability in 2D and 3D EPI resting-state fMRI." *NMR in Biomedicine* 18.8 (2005), pp. 534–542.

[146] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. J. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews. "Advances in functional and structural MR image analysis and implementation as FSL." *NeuroImage* 23 Suppl 1 (2004), S208–19.

[147] S. M. Smith. "Fast robust automated brain extraction." *Human Brain Mapping* 17.3 (2002), pp. 143–155.

[148] J. C. W. Brooks, C. F. Beckmann, K. L. Miller, R. G. Wise, C. A. Porro, I. Tracey, and M. Jenkinson. "Physiological noise modelling for spinal functional magnetic resonance imaging studies." *NeuroImage* 39.2 (2008), pp. 680–692.

[149] N. Todd, S. Moeller, E. J. Auerbach, E. Yacoub, G. Flandin, and N. Weiskopf. "Evaluation of 2D multiband EPI imaging for high-resolution, whole-brain, task-based fMRI studies at 3T: Sensitivity and slice leakage artifacts." *NeuroImage* 124.Pt A (2016), pp. 32–42.

[150] C. Triantafyllou, R. D. Hoge, and L. L. Wald. "Effect of spatial smoothing on physiological noise in high-resolution fMRI." *NeuroImage* 32.2 (2006), pp. 551–557.

[151] F. A. Breuer, M. Blaimer, M. F. Mueller, N. Seiberlich, R. M. Heidemann, M. A. Griswold, and P. M. Jakob. "Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA)." *Magnetic Resonance in Medicine* 55.3 (2006), pp. 549–556.

[152] M. Narsude, W. Zwaag, T. Kober, R. Gruetter, and J. P. Marques. "Improved temporal resolution for functional studies with reduced number of segments with three-dimensional echo planar imaging." *Magnetic Resonance in Medicine* 72.3 (2014), pp. 786–792.

[153] R. Stirnberg, D. Brenner, and T. Stöcker. "Variable Flip Angles and Echo Train Lengths in Segmented 3D-EPI at 3 and 7 Tesla." *Proceedings of the 22nd Annual Meeting of the ISMRM.* Milan, Italy, 2014, p. 609.

[154] J. R. Polimeni, H. Bhat, T. Witzel, T. Benner, T. Feiweier, S. J. Inati, V. Renvall, K. Heberlein, and L. L. Wald. "Reducing sensitivity losses due to respiration and motion in accelerated echo planar imaging by reordering the autocalibration data acquisition." *Magnetic Resonance in Medicine* 75.2 (2015), pp. 665–679.

[155] R. Stirnberg and T. Stöcker. "Conventional 2D-EPI or Segmented 3D-EPI? A Temporal SNR Study at 3 and 7 Tesla." *Proceedings of the 22nd Annual Meeting of the ISMRM.* Milan, Italy, p. 868.

[156] R. Stirnberg, W. Huijbers, B. A. Poser, and T. Stöcker. "Ultra-fast gradient echo EPI with controlled aliasing at 3T: simultaneous multi-slice vs. 3D-EPI." *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 941.

[157] O. Reynaud, M. Narsude, J. Jorge, J. P. Marques, R. Gruetter, and W. van der Zwaag. "Comparison of physiological signals characterization efficiency of SMS-EPI and 3D-EPI-CAIPI at 7T." *Proceedings of the 22nd Annual Meeting of the OHBM*. 2016, p. 2074.

[158] J. Ladstein, H. R. Evensmoen, A. K. Håberg, A. Kristoffersen, and P. E. Goa. "Effect of Task-Correlated Physiological Fluctuations and Motion in 2D and 3D Echo-Planar Imaging in a Higher Cognitive Level fMRI Paradigm." *Frontiers in Neuroscience* 10 (2016), p. 225.

[159] C. Triantafyllou, R. D. Hoge, G. Krueger, C. J. Wiggins, A. Potthast, G. C. Wiggins, and L. L. Wald. "Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters." *NeuroImage* 26.1 (2005), pp. 243–250.

[160] C. Triantafyllou, J. R. Polimeni, and L. L. Wald. "Physiological noise and signal-to-noise ratio in fMRI with multi-channel array coils." *NeuroImage* 55.2 (2011), pp. 597–606.

[161] M. A. Lindquist, C.-H. Zhang, G. Glover, and L. Shepp. "Rapid three-dimensional functional magnetic resonance imaging of the initial negative BOLD response." *Journal of Magnetic Resonance* 191.1 (2008), pp. 100–111.

[162] S. V. Jonathan, P. Vakil, Y. I. Jeong, R. G. Menon, S. A. Ansari, and T. J. Carroll. "RAZER: a pulse sequence for whole-brain bolus tracking at high frame rates." *Magnetic Resonance in Medicine* 71.6 (2014), pp. 2127–2138.

[163] M. Magnusson and O. D. Leinhard. "A 3D-Plus-Time Radial-Cartesian Hybrid Sampling of K-Space With High Temporal Resolution and Maintained Image Quality for MRI and FMRI." *Proceedings of the 19th Annual Meeting of the ISMRM*. Montreal, Canada, 2011, p. 2851.

[164] N. N. Graedel, M. Chiew, and K. L. Miller. "Motion correction for functional MRI with hybrid radial-Cartesian 3D EPI." *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 940.

[165] N. N. Graedel, M. Chiew, J. A. McNab, and K. L. Miller. "FMRI using a 3D radial-Cartesian trajectory: spatio-temporal tunability and artifact correction." *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, Canada, 2015, p. 360.

[166] P. Ehses, G. Shajan, and K. Scheffler. "3D Radial GRE-EPI with up to 8-fold acceleration for functional imaging at 9.4T". *Proceedings of the 22nd Annual Meeting of the ISMRM*. Milan, Italy, 2014, p. 2992.

[167] S. Winkelmann, T. Schaeffter, T. Koehler, H. Eggers, and O. Dössel. "An Optimal Radial Profile Order Based on the Golden Ratio for Time-Resolved MRI." *IEEE Transactions on Medical Imaging* 26.1 (2007), pp. 68–76.

[168] M. Chiew, N. N. Graedel, J. A. McNab, S. M. Smith, and K. L. Miller. "Accelerating functional MRI using fixed-rank approximations and radial-cartesian sampling." *Magnetic Resonance in Medicine* (2016). In press.

[169] M. Chiew, S. M. Smith, P. J. Koopmans, N. N. Graedel, T. Blumensath, and K. L. Miller. "k-t FASTER: Acceleration of functional MRI data acquisition using low rank constraints." *Magnetic Resonance in Medicine* 74.2 (2015), pp. 353–364.

[170] M. Lustig, D. Donoho, and J. M. Pauly. "Sparse MRI: The application of compressed sensing for rapid MR imaging." *Magnetic Resonance in Medicine* 58.6 (2007), pp. 1182–1195.

[171] L. Feng, R. Grimm, K. T. Block, H. Chandarana, S. Kim, J. Xu, L. Axel, D. K. Sodickson, and R. Otazo. "Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI." *Magnetic Resonance in Medicine* 72.3 (2014), pp. 707–717.

[172] M. Chiew, N. N. Graedel, and K. L. Miller. "Promoting incoherence of radial x-f point spread functions using randomly perturbed golden angles." *Proceedings of the 24th Annual Meeting of the ISMRM*. SIngapore, 2016, p. 1441.

[173] D. M. Spielman and J. M. Pauly. "Magnetic resonance fluoroscopy using spirals with variable sampling densities." *Magnetic Resonance in Medicine* 34.3 (1995), pp. 388–394.

[174] D. Ivanov, M. Barth, K. Uludag, and B. A. Poser. "Robust ACS acquisition for 3D echo planar imaging." *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, Canada, 2015, p. 2059.

[175] J. H. Duyn. "Steady state effects in fast gradient echo magnetic resonance imaging." *Magnetic Resonance in Medicine* 37.4 (1997), pp. 559–568.

[176] P. Vakil, S. A. Ansari, M. C. Hurley, H. Bhat, H. H. Batjer, B. R. Bendok, C. S. Eddleman, and T. J. Carroll. "Magnetization spoiling in radial FLASH contrast-enhanced MR digital subtraction angiography." *Journal of Magnetic Resonance Imaging* 36.1 (2012), pp. 249–258.

[177] G. N. Ramachandran and A. V. Lakshminarayanan. "Three-dimensional Reconstruction from Radiographs and Electron Micrographs: Application of Convolutions instead of Fourier Transforms." *Proceedings of the National Academy of Sciences* 68.9 (1971), pp. 2236–2240.

[178] F. Aurenhammer. "Voronoi diagrams - a survey of a fundamental geometric data structure." *ACM Computing Surveys (CSUR)* 23.3 (1991), pp. 345–405.

[179] H. Bruder, H. Fischer, H. E. Reinfelder, and F. Schmitt. "Image reconstruction for echo planar imaging with nonequidistant k-space sampling." *Magnetic Resonance in Medicine* 23.2 (1992), pp. 311–323.

[180] K. Block and M. Uecker. "Simple method for adaptive gradient-delay compensation in radial MRI." *Proceedings of the 19th Annual Meeting of the ISMRM*. Montreal, Canada, 2011, p. 2816.

[181] A. Moussavi, M. Untenberger, M. Uecker, and J. Frahm. "Correction of Gradient-Induced Phase Errors in Radial MRI." *Proceedings of the 21st Annual Meeting of the ISMRM*. Salt Lake City, USA, 2013, p. 3776.

[182] N.-k. Chen and A. M. Wyrwicz. "Removal of EPI Nyquist ghost artifacts with two-dimensional phase correction." *Magnetic Resonance in Medicine* 51.6 (2004), pp. 1247–1253.

[183] A. Deshmane, M. Blaimer, F. Breuer, P. Jakob, J. Duerk, N. Seiberlich, and M. Griswold. "Self-calibrated trajectory estimation and signal correction method for robust radial imaging using GRAPPA operator gridding." *Magnetic Resonance in Medicine* 75.2 (2016), pp. 883–896.

[184] C. Barmet, N. De Zanche, and K. P. Pruessmann. "Spatiotemporal magnetic field monitoring for MR." *Magnetic Resonance in Medicine* 60.1 (2008), pp. 187–197.

[185] C. Barmet, N. De Zanche, B. J. Wilm, and K. P. Pruessmann. "A transmit/receive system for magnetic field monitoring of in vivo MRI." *Magnetic Resonance in Medicine* 62.1 (2009), pp. 269–276.

[186] D. Gallichan, C. A. Cocosco, A. Dewdney, G. Schultz, A. Welz, J. Hennig, and M. Zaitsev. "Simultaneously driven linear and nonlinear spatial encoding fields in MRI." *Magnetic Resonance in Medicine* 65.3 (2011), pp. 702–714.

[187] D. C. Noll. "Rapid MR image acquisition in the presence of background gradients." *ISBI* (2002), pp. 725–728.

[188] D. C. Noll, J. A. Fessler, and B. P. Sutton. "Conjugate phase MRI reconstruction with spatially variant sample density correction." *IEEE Transactions on Medical Imaging* 24.3 (2005), pp. 325–336.

[189] Y. Ding, H. Xue, R. Ahmad, C. Chang, S. T. Ting, and O. P. Simonetti. "Paradoxical effect of the signal-to-noise ratio of GRAPPA calibration lines: A quantitative study." *Magnetic Resonance in Medicine* 74.1 (2014), pp. 231–239.

[190] P. Qu, K. Zhong, B. Zhang, J. Wang, and G. X. Shen. "Convergence behavior of iterative SENSE reconstruction with non-Cartesian trajectories." *Magnetic Resonance in Medicine* 54.4 (2005), pp. 1040–1045.

[191] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. "Advances in sensitivity encoding with arbitrary k-space trajectories." *Magnetic Resonance in Medicine* 46.4 (2001), pp. 638–651.

[192] M. Buehrer, K. P. Pruessmann, P. Boesiger, and S. Kozerke. "Array compression for MRI with large coil arrays." *Magnetic Resonance in Medicine* 57.6 (2007), pp. 1131–1139.

[193] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, and M. Lustig. "Fast l_1-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime." *IEEE Transactions on Medical Imaging* 31.6 (2012), pp. 1250–1262.

[194] M. Uecker, F. Ong, J. I. Tamir, D. Bahri, P. Virtue, J. Y. Cheng, T. Zhang, and M. Lustig. "Berkeley Advanced Reconstruction Toolbox". *Proceedings of the 23rd Annual Meeting of the ISMRM*. Toronto, Canada, 2015, p. 2486.

[195] C. F. Beckmann and S. M. Smith. "Probabilistic independent component analysis for functional magnetic resonance imaging." *IEEE Transactions on Medical Imaging* 23.2 (2004), pp. 137–152.

[196] J. C. W. Brooks, O. K. Faull, K. T. S. Pattinson, and M. Jenkinson. "Physiological noise in brainstem FMRI." *Frontiers in Human Neuroscience* 7 (2013), p. 623.

[197] T. A. Yousry, U. D. Schmid, H. Alkadhi, D. Schmidt, A. Peraud, A. Buettner, and P. Winkler. "Localization of the motor hand area to a knob on the precentral gyrus. A new landmark." *Brain* 120.1 (1997), pp. 141–157.

[198] S. Wundrak, J. Paul, J. Ulrici, E. Hell, M.-A. Geibel, P. Bernhardt, W. Rottbauer, and V. Rasche. "Golden ratio sparse MRI using tiny golden angles." *Magnetic Resonance in Medicine* 75.6 (2016), pp. 2372–2378.

[199] K. Setsompop, B. A. Gagoski, J. R. Polimeni, T. Witzel, V. J. Wedeen, and L. L. Wald. "Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty." *Magnetic Resonance in Medicine* 67.5 (2012), pp. 1210–1224.

[200] S. Moeller, E. Yacoub, C. A. Olman, E. Auerbach, J. Strupp, N. Harel, and K. Ugurbil. "Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI." *Magnetic Resonance in Medicine* 63.5 (2010), pp. 1144–1153.

[201] N. Todd, O. Josephs, M. F. Callaghan, A. Lutti, and N. Weiskopf. "Prospective motion correction of 3D echo-planar imaging data for functional MRI using optical tracking." *NeuroImage* 113 (2015), pp. 1–12.

[202] T. Schäffter, V. Rasche, and I. C. Carlsen. "Motion compensated projection reconstruction." *Magnetic Resonance in Medicine* 41.5 (1999), pp. 954–963.

[203] G. Cruz, D. Atkinson, C. Buerger, T. Schaeffter, and C. Prieto. "Accelerated motion corrected three-dimensional abdominal MRI using total variation regularized SENSE reconstruction." *Magnetic Resonance in Medicine* 75.4 (2016), pp. 1484–1498.

[204] G. Vaillant, C. Prieto, C. Kolbitsch, G. P. Penney, and T. Schaeffter. "Retrospective Rigid Motion Correction in k-Space for Segmented Radial MRI." *IEEE Transactions on Medical Imaging* 33.1 (2014), pp. 1–10.

[205] M. Costagli, R. A. Waggoner, K. Ueno, K. Tanaka, and K. Cheng. "Correction of 3D rigid body motion in fMRI time series by independent estimation of rotational and translational effects in k-space". *NeuroImage* 45.3 (2009), pp. 749 –757.

[206] L. Cordero-Grande, R. P. A. G. Teixeira, E. J. Hughes, J. Hutter, A. N. Price, and J. Hajnal. "Sensitivity encoding for aligned multishot magnetic resonance reconstruction." *IEEE Transactions on Computational Imaging* 2.3 (2016), pp. 266–280.

[207] P. G. Batchelor, D. Atkinson, P. Irarrazaval, D. L. G. Hill, J. Hajnal, and D. Larkman. "Matrix description of general motion correction applied to multishot images." *Magnetic Resonance in Medicine* 54.5 (2005), pp. 1273–1280.

[208] K. Shoemake. "Animating rotation with quaternion curves." *SIGGRAPH* 19.3 (1985), pp. 245–254.

[209] M. Magnusson, O. D. Leinhard, and P. Brynolfsson. "Improved temporal resolution in radial k-space sampling using an hourglass filter." *Proceedings of the 17th Annual Meeting of the ISMRM*. Honolulu, Hawaii, 2009, p. 2836.

[210] M. E. Kelly and K. Miller. "An assessment of motion artefacts in multi band EPI for high spatial and temporal resolution resting state fMRI." *Proceedings of the 21st Annual Meeting of the ISMRM*. Salt Lake City, USA, 2013, p. 3275.

[211] Z. Faraji-Dana, F. Tam, J. Chen, and S. Graham. "Interactions between head motion and coil sensitivity in accelerated fMRI". *Journal of Neuroscience Methods* 270 (2016), pp. 46 –60.

[212] M. Zaitsev, K. Zilles, and N. J. Shah. "Shared k-space echo planar imaging with keyhole." *Magnetic Resonance in Medicine* 45.1 (2001), pp. 109–117.

[213] C. Liu, D. Bammer R. Kim, and M. Moseley. "Self-navigated interleaved spiral (SNAILS): application to high-resolution diffusion tensor imaging." *Magnetic Resonance in Medicine* 52.6 (2004), pp. 1388–1396.

[214] A. V. Barger, W. F. Block, Y. Toropov, T. M. Grist, and C. A. Mistretta. "Time-Resolved Contrast-Enhanced Imaging WithIsotropic Resolution and Broad Coverage Using an Undersampled 3D Projection Trajectory." *Magnetic Resonance in Medicine* 48.2 (2002), pp. 297–305.

[215] N. N. Graedel, S. J. Vannesjo, L. Kasper, S. Gross, B. E. Dietrich, C. Barmet, and K. P. Pruessmann. "Image Reconstruction Using the Gradient Impulse Response for Trajectory Prediction." *Proceedings of the 21st Annual Meeting of the ISMRM*. Salt Lake City, USA, 2013, p. 552.

[216] J. H. Duyn, Y. Yang, J. A. Frank, and J. W. van der Veen. "Simple correction method for k-space trajectory deviations in MRI." *Journal of Magnetic Resonance (San Diego, Calif. : 1997)* 132.1 (1998), pp. 150–153.

[217] W. Chen, C. T. Sica, and C. H. Meyer. "Fast conjugate phase image reconstruction based on a Chebyshev approximation to correct for B0 field inhomogeneity and concomitant gradients". *Magnetic Resonance in Medicine* 60.5 (2008), pp. 1104–1111.

[218] L.-C. Man, J. M. Pauly, and A. Macovski. "Multifrequency interpolation for fast off-resonance correction". *Magnetic Resonance in Medicine* 37.5 (1997), pp. 785–792.

[219] B. Sutton, D. Noll, and J. Fessler. "Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities." *IEEE Transactions on Medical Imaging* 22.2 (2003), pp. 178–188.

[220] M. H. S. Cauley and L. Wald. "TArgeted Motion Estimation and Reduction (TAMER): Data Consistency Based Motion Mitigation using a Reduced Model Joint Optimization." *Proceedings of the 24th Annual Meeting of the ISMRM*. Singapore, 2016, p. 1849.