

Novel Methods for Brainstem FMRI



H.N. (Rob) Tijssen

Physics Group, FMRIB Centre,
Nuffield Department of Clinical Neurosciences
and University College

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Supervisors: Dr. Karla L. Miller and Prof. Peter Jezzard

Trinity Term 2011

Novel Methods for Brainstem FMRI

Rob H.N. Tijssen

University College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Trinity Term 2011

Abstract

The brainstem plays a crucial role in a great number of vital functions such as respiratory regulation, visual reflexes, and the perception of pain. The small size and close proximity of the nuclei requires high-resolution functional magnetic resonance imaging (fMRI). However, brainstem fMRI using conventional gradient-echo echo-planar imaging (GRE-EPI) techniques is challenging due to the increased signal dropout and geometric distortions in the brainstem.

The primary aim of the work presented in this thesis was to investigate alternative methods for brainstem fMRI in order to overcome some of the challenges associated with single-shot GRE-EPI techniques. Towards this goal 3D segmented sequences were explored, which have the advantage that the size of the geometric distortions is not proportional to the resolution at which is scanned. In particular, two sequences were investigated: balanced steady-state free-precession (bSSFP) and spoiled gradient echo (SPGR).

First, a set of experiments was conducted, in which each experiment aimed to isolate a limited range of sequence properties in order to characterize and assess the potential of the candidate sequences. It was found that bSSFP has better noise characteristics compared to GRE-EPI when applied with a 2D acquisition, but when 3D readouts were used the signal instabilities increased dramatically.

Based on these findings, experiments that investigate the influence of multi-shot acquisitions on signal instabilities caused by physiological noise were performed. The signal instabilities were found to mainly originate from regions of CSF and blood and were highly correlated to the cardiac cycle. Several correction methods were explored and one method was identified to be implemented *in vivo*.

A novel method that allows real-time cardiac synchronization of the k -space acquisition was developed. The developed methods used a custom parallel imaging reconstruction to allow for acquisition with a fixed volume frame rate, which is desirable for fMRI purposes. The method was found to reduce the signal instabilities in 3D SPGR and bSSFP significantly.

A comprehensive assessment of two currently available retrospective correction techniques was conducted and their practicalities were compared. Recommendations are made to improve the robustness of the investigated correction methods. A novel optimization method was implemented, which was developed to determine the optimal regressor set for retrospective corrections. The method can be applied to image based as well as k -space based methods.

Acknowledgements

Over the last four years many people have contributed to the realisation of this thesis. Without their academic or personal support writing this thesis would have been impossible. I'm extremely grateful to all the people around me, who have offered their kind support during my time as a D.Phil. student. Although I would have liked nothing more than to use this section to thank everyone in person, it is only possible to mention a few.

First and foremost I would like to thank my supervisor Dr. Karla Miller. When I first came to FMRIB back in 2006 to do a short internship I was already deeply impressed by Karla's passion for science and her mentoring skills. Karla has played a key role in my decision to pursue a career in science by offering me the opportunity to do a D.Phil. under her supervision. Karla, I'm thankful for your all-encompassing supervision. For not only advising me on my projects, but especially for your dedication to teach me how to think like a scientist. You always managed to find the right balance between providing guidance where needed, but at the same time giving me enough space to develop my own ideas. Your enthusiasm, diligence, and creativity is something that has truly inspired me, and something I hope to be able to match in my future career.

A big token of my appreciation also goes out to my second supervisor, Prof. Peter Jezzard. Thank you Peter for your trust in me and for giving me the opportunity to come to Oxford. Your experience and unsurpassed knowledge about the field is truly admirable, and your advice has been invaluable for the preparation of this thesis.

Apart from my supervisors I had the pleasure to get to know many other talented scientists; professionally as well as personally. One person in particular, who has been there for me "since the beginning" is Saad Jbabdi. Saad is one of the warmest personalities you could ever meet. Thank you Saad for not only being a great friend, but also for your infectious (ever so geeky) enthusiasm about science. I'll never forget how you bought Einstein's Relativity for me to read when I was recovering from surgery, because you were worried I'd waste my time watching television all day long!

Thanks to all my office mates with whom I shared the downstairs open plan area. Especially Tom Okell, my "friendly neighbour" and good friend. I have benefitted tremendously from our four-year lasting collaboration. I don't think it's possible to even estimate how many hours Tom has saved me by patiently helping me out whenever I got stuck with a problem or by generously sharing

his latest discoveries on the intricacies of pulse programming, Matlab quirks, or handy shell scripts (big thanks to Adrian for the scripts as well!). I thoroughly enjoyed the many discussions we had during coffee breaks on exciting new stuff we had seen at the ISMRM or MR physics in general. Thanks to Rob Frost and James Meakin for the many late night and weekend sessions at the lab. It was good to know that, no matter what time a day, at least one of you would be there! And to the ever good-natured Way Charng Chen for never arriving nor leaving the lab without a big smile on his face. Thanks to Dan Bulte for resisting the imminent move to the upstairs offices and for organising a great trip around Hawaii in 2009. Thanks to Wilfred Lam for “oxygenating” the office with all his home grown plants, and to the analysis students Adrian Groves, Reza Salimi, and Aleksandar Petrovic for helping out us physicists out with silly maths and stats problems.

Amongst the post-docs and principle investigators in FMRIB there are also many people that I owe my gratitude to. Thanks to Stuart Clare, Jesper Andersson, Mark Woolrich, Eugene Duff, Jill O’Reilly, Charlie Stagg, Chris Rodgers, and Matt Robson for generously sharing their expertise with the rest of the lab in the form of lectures, clinics, and informal chats. Thanks to Mark Jenkinson, Jonathan Brooks, Manus Donahue, Nicola Filippini, Prof. Steve Smith, and the Trondheim group, for their collaboration on various projects. Thank you Marilyn, Sue, Dave, Duncan, Iain, and Matt for keeping FMRIB and Jalapeno up and running. Without you guys no theses would have been produced this year!

It’s been fantastic to experience how scientific and social life blends at FMRIB. Often in the form of discussing spin physics or Bayesian statistics over a pint and cheesy chips at the White Hart. This to me has emphasised the fact that being a researcher is a life-style rather than a job and has made my D.Phil. time a very stimulating time of my life. Many people have contributed to the great atmosphere in FMRIB. Thanks to Brad Macintosh, Stefan Piechnik, Daniel Gallichan, Jennifer McNab, Molly Bright, Jingyi Xie, Verena Heise, Velicia Bachtiar, Christel Gudberg, Adam Thomas, Cassandra Sampaio, Lucy Matthews, Claire Allman, Gwen Douaud, Katie Warnaby, Ricarda Menke, Catherine Thiel, Inga Kröger, Lisa Bruckert, Andri Tziortzi, Stacey Davie, Thomas Vaessen, Dian de Vries, and Karolina Wartolowska for all the drinks, movies, and dinners we shared. To Natalie Voets for the recipes, wine suggestions, and amazing buttery biscuit base. And to Michael Kelly and Jamie Near and Saad, for their fantastic musical contributions at many of these social events, including my goodbye party.

I would also like to thank all my friends outside the lab, friends at the OU Kickboxing club, the Dutch Society, and Univ for making my time in Oxford an unforgettable experience. People who have certainly aided to this as well are my housemates, Adrian, Andy, and Lee. The house pranks at the Mansion will certainly be missed!

Finally a huge thank you to Mom, Dad, Esther, and Amber for all their support throughout my D.Phil. and all the years leading up to it. Thanks for always being there for me. Especially during the rollercoaster that was 2009. This really made me realise how fortunate I am to have the family that I have.

Contents

Contents	v
List of Figures	ix
List of Tables	xi
List of Publications	xii
1 Introduction	1
1.1 Motivation	1
1.2 Thesis Outline	3
2 Principles of MR Imaging	6
2.1 The MR Signal	6
2.1.1 Magnetisation	6
2.1.2 Relaxation	10
2.1.3 The Bloch Equation	12
2.2 Image Formation	13
2.2.1 Gradient Fields	13
2.2.2 Slice Selection	14
2.2.3 k -space	15
2.2.4 Readout Trajectories	17
2.2.5 3D Trajectories	19
2.3 Gradient Echo Sequences	20
2.3.1 Longitudinal Steady-State	20
2.3.2 Transverse Steady-State	22
2.4 Steady-State Free Precession	24
2.4.1 Balanced SSFP	24
2.4.2 Balanced-SSFP Signal Equations	26
2.5 Image Reconstruction	28
2.5.1 Standard Image Reconstruction for EPI	29
2.5.2 Partial Fourier Reconstruction	30
2.5.3 Accelerated Parallel Imaging Techniques	31

3	Functional MRI methods	35
3.1	The BOLD Contrast	35
3.1.1	The fMRI Experiment and Data Analysis	36
3.2	Imaging Sequences for fMRI	38
3.2.1	2D Single-Shot GRE-EPI	38
3.2.2	3D fMRI	39
3.2.3	Balanced SSFP fMRI	40
3.3	Brainstem Anatomy and Physiology	41
3.4	Practical Considerations Specific to Brainstem fMRI	43
3.4.1	Resolution	43
3.4.2	Distortion	44
3.4.3	Physiological Noise	45
3.5	Retrospective Correction of Physiological Noise	45
4	Characterization of 3D Sequences for Brainstem fMRI: Preliminary Investigations	47
4.1	Introduction	47
4.2	The Effect of Spin Echo Contrast on Physiological Noise	49
4.2.1	The Temporal Noise Model	50
4.2.2	Methods	51
4.2.3	Results	54
4.2.4	Discussion	54
4.3	Physiological Noise in 2D balanced SSFP and SPGR	58
4.3.1	Time-Series Power Spectrum Analysis	59
4.3.2	Comparing 2D bSSFP and SPGR to GRE-EPI	60
4.4	3D balanced SSFP; fMRI of the Hippocampus	66
4.4.1	3D Stack-of-segmented EPI	67
4.4.2	Methods	68
4.4.3	Results	69
4.5	Temporal Stability in the Brainstem using 3D Readouts	74
4.5.1	Methods	74
4.5.2	Results and Discussion	75
4.6	Discussion	75
4.7	Conclusion	78
5	Physiological Noise: Characterization and Correction Methods	80
5.1	Introduction	80
5.1.1	Background	81
5.2	Regression of Physiological Processes with 2D data	82
5.2.1	Methods	82
5.2.2	Results	84
5.2.3	Discussion	85
5.3	Characterization of the Effects of Flow	88
5.3.1	Quantification of the First-order Gradient Moments in a bSSFP Acquisition with a 3D Stack-of-Segmented EPI Readout	89

5.3.2	The Effects of Flow-induced Phase perturbations on the Steady-state signal <i>in vivo</i>	92
5.4	Spoiling CSF in Balanced SSFP	98
5.4.1	Introduction	98
5.4.2	Methods	99
5.4.3	Results and Discussion	100
5.5	Readout Synchronization Simulations	104
5.5.1	Theory	105
5.5.2	Methods	106
5.5.3	Results and Discussion	108
5.6	Conclusions	110
6	Cardiac synchronization	111
6.1	Introduction	111
6.2	Theory	111
6.3	Description of the Method	113
6.3.1	Synchronization Method	114
6.4	Sequence Implementation	117
6.5	Sequence Optimization	118
6.5.1	Optimization of the Reordering Strategy	119
6.5.2	Optimization of the Image Reconstruction	121
6.5.3	Trigger Delay	123
6.6	Assessment of Cardiac Synchronization in 3D SPGR and bSSFP	125
6.6.1	Results	127
6.7	Discussion	130
6.7.1	Characterization of Signal Instabilities	130
6.7.2	Current Implementation of the Correction Method	133
6.8	Conclusion	138
7	Retrospective Correction Strategies	139
7.1	Introduction	139
7.1.1	Chapter Overview	141
7.2	Data Acquisition Methods	142
7.2.1	Data acquisition	142
7.2.2	Physiological Monitoring	145
7.2.3	Definition of the Regressors	146
7.2.4	Performing the Regressions; RETROICOR vs. RETROKCOR	147
7.2.5	ROI Analysis	150
7.3	Method Development and Characterization	150
7.3.1	Initial Characterization of RETROICOR and RETROKCOR	151
7.3.2	Model Optimization; Implementation of the Method	157
7.4	Results	163
7.4.1	Model Selection Results on 3DSHORT-data	163
7.4.2	Model Selection Results on 3DLONG-data	166
7.4.3	Optimized Sequence Comparison	170

7.5	General Discussion	175
7.5.1	Optimization of RETROKCOR	175
7.5.2	Optimization of the Regressor Set	176
7.5.3	The Potential for 3D Brainstem FMRI using bSSFP	177
7.6	Conclusion	178
8	Summary and Future Work	181
8.1	Thesis Summary	181
8.2	Future Work	183
	References	187
A	Definition of the Confound Regressors	200

List of Figures

1.1	Physiological Noise Sensitivity	4
2.1	Polarisation and Precession	8
2.2	Excitation	10
2.3	Slice Selective Excitation	15
2.4	2D Readout Trajectories	19
2.5	3D Segmented EPI	20
2.6	Balanced SSFP	25
2.7	Balanced SSFP Signal Profiles	28
2.8	Partial Fourier	31
3.1	BOLD physiology	37
3.2	bSSFP Signal Profile	42
3.3	Brainstem Anatomy	43
4.1	Temporal Noise in GRE-EPI and SE-EPI	55
4.2	Signal vs. Flip Angle in GRE-EPI and SE-EPI	56
4.3	Comparison of Results with Literature	57
4.4	Power Spectrum Example	62
4.5	Sequence Comparison	63
4.6	Effect of Segmentation	66
4.7	3D Stack-of-Segmented EPI	67
4.8	Raw Image Quality; GRE-EPI vs. 3D bSSFP	71
4.9	Distortion; GRE-EPI vs. 3D bSSFP	72
4.10	Activation Maps	73
4.11	tSNR in 2D and 3D bSSFP	76
5.1	Regression Results of 2D SPGR and SSFP	84
5.2	M1 Moments in 3D EPI	93
5.3	RF Flow Simulation	97
5.4	Steady-state Approach with Bipolar Spoiling	101
5.5	Steady-state Profiles with Bipolar Spoiling	103
5.6	Synchronization Simulations Theory	107
5.7	Simulated Synchronization Results	109
6.1	Synchronization Theory	113

6.2	Synchronization Method	116
6.3	Reorder Optimization Results	119
6.4	GRAPPA Optimization Results	123
6.5	Synchronization Delay Results	124
6.6	ROI Analysis Results	128
6.7	SPGR tSNR Maps	131
6.8	bSSFP tSNR Maps	132
6.9	Eddy Current Assessment	136
7.1	Simulation Observations	153
7.2	RETROKCOR Optimization	155
7.3	<i>in vivo</i> Comparison on 2D data	158
7.4	Correlation Matrix	162
7.5	Single EV regression on 3DSHORT-Data	164
7.6	BIC Selection on 3DSHORT-Data	167
7.7	Single EV regression on 3DLONG-Data	168
7.8	tSNR in optimized sequences	171
7.9	Synchronization Efficacy	174
7.10	Non-linear Phase Uncertainty	176

List of Tables

2.1	Relaxation Times	12
5.1	Flow Simulations	96
7.1	Sequence Parameters 3DLONG-Data	144
7.2	Sequence Parameters 3DSHORT-Data	145
7.3	Full Regressor List	148
7.4	BIC Selected Regressors	180

List of Publications Arising from this Thesis

Journal Articles

1. **R. H. Tijssen**, T. W. Okell, K. L. Miller, "A real-time cardiac synchronized readout with fixed frame rate for reducing physiological instabilities in 3D FMRI", *NeuroImage*, 57(4):1364–75, 2011
2. K. L. Miller, **R. H. Tijssen**, N. Stikov, T. W. Okell, "Steady-State MRI: Methods for neuroimaging", *Imaging in Medicine*, 3(1):93–105, 2011

Conference Proceedings

1. **R. H. Tijssen**, T. W. Okell, K. L. Miller, "A real-time cardiac synchronization method for reducing flow-induced instabilities in SSFP FMRI of the brainstem", in *Proceedings 19th International Society of Magnetic Resonance in Medicine*, Montreal, 2011
2. **R. H. Tijssen**, K. L. Miller, "The effect of cardiac synchronization on the temporal characteristics of 3D SSFP and 3D SPGR", in *Proceedings 19th International Society of Magnetic Resonance in Medicine*, Montreal, 2011
3. T. W. Okell, P. Schmitt, X. Bi, M. A. Chappell, **R. H. Tijssen**, K. L. Miller, P. Jezzard, "4D Vessel-Encoded Arterial Spin Labeling Angiography", in *Proceedings 19th International Society of Magnetic Resonance in Medicine*, Montreal, 2011
4. **R. H. Tijssen**, K. L. Miller, "Cardiac synchronization with fixed frame rate for reducing physiological instabilities in 3D Brainstem FMRI", in *Proceedings 16th Annual Meeting, British Chapter of the International Society of Magnetic Resonance in Medicine*, Nottingham, 2010
5. **R. H. Tijssen**, K. L. Miller, "Reducing physiological noise in 3D brainstem FMRI: Cardiac synchronization with fixed frame rate", in *Proceedings 16th Scientific Meeting, Organization Human Brain Mapping*, Barcelona, 2010
6. **R. H. Tijssen**, S. M. Smith, P. Jezzard, R. Frost, M. Jenkinson, K. L. Miller, "Characterization and Correction of Physiological Instabilities in 3D FMRI",

- in *Proceedings 18th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Stockholm, 2010
7. M. H. Chappell, H. Lehn, P. E. Goa, A. Kristoffersen, **Tijssen RH**, A. K. Häberg, K. L. Miller, "fMRI of the medial temporal lobe using balanced steady-state free precession", in *Proceedings 18th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Stockholm, 2010
 8. W. .C. Chen, **R. H. Tijssen**, C. Rodgers, J. Near, K. L. Miller, "Balanced SSFP profile asymmetries reflect frequency distribution asymmetries: Evidence from chemical shift imaging (CSI)", in *Proceedings 18th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Stockholm, 2010
 9. P. E. Goa, A. Kristoffersen, M. H. Chappell, **R. H. Tijssen**, A. K. Häberg, K. L. Miller, "Passband bSSFP: Functional contrast compared to GRE-EPI and SE-EPI at 3T", in *Proceedings 18th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Stockholm, 2010
 10. **R. H. Tijssen**, N. van Strien, M. H. Chappell, H. Lehn, N. Fillipini, A. K. Häberg, K. L. Miller, "High resolution, low distortion, fMRI of the Hippocampus using balanced SSFP at 3T", in *Proceedings 15th Scientific Meeting, Organization of Human Brain Mapping*, San Francisco, 2009
 11. M. J. Vaessen, P. A. Hofman, J. F. Jansen, **R. H. Tijssen**, A. Aldenkamp, W. H. Backes, "Reproducibility of small world metrics from diffusion tensor tractography at 3 Tesla", in *Proceedings 15th Scientific Meeting, Organization of Human Brain Mapping*, San Francisco, 2009
 12. **R. H. Tijssen**, M. Jenkinson, P. Jezzard, K. L. Miller, "Characterising physiological noise in the brainstem: passband SSFP vs GRE-EPI", in *Proceedings 17th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Honolulu, 2009
 13. M. J. Donahue, S. K. Piechnik, **R. H. Tijssen**, D. Gallichan, K. L. Miller, P. Jezzard, "A theoretical and experimental investigation of vascular-space-occupancy (VASO) blood nulling times: influence of hematocrit and oxygenation on null times and CBV quantification", in *Proceedings 17th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Honolulu, 2009
 14. M. Jenkinson, **R. H. Tijssen**, J. Brooks, K. L. Miller, "Voxel-wise regressor selection for physiological noise correction with RETROICOR", in *Proceedings 17th Scientific Meeting, International Society of Magnetic Resonance in Medicine*, Honolulu, 2009

Chapter 1

Introduction

Functional imaging of the brainstem allows the investigation of various important neuronal processes, such as breathing control, cardiovascular control, and the reception and desensitisation of pain. However, compared to the cerebrum, the brainstem still remains a challenging area of the brain to image.

This D.Phil. thesis describes the development of image acquisition methods to improve the quality of functional magnetic resonance imaging (fMRI) of the brainstem. In this chapter the motivation for this work is described, and an outline of the remainder of the thesis given.

1.1 Motivation

The small size of the brainstem belies its important role in the control of sympathetic nervous system. The brainstem regulates a great number of vital functions such as respiratory regulation and arousal reflexes [1] and plays an important role in the perception and desensitization of pain. Chronic pain is a disabling disease with a high prevalence (19%) and forms considerable direct health care costs, and indirect costs such as reduced productivity. It is estimated that chronic pain costs the European economy over €34 billion each year [2]. A few studies have re-

ported preliminary fMRI evidence of the role of the brainstem in the perception of pain [3, 4, 5, 6]. However, a large portion of the nuclei in the brainstem are relatively small and densely spaced, which makes accurate identification of the active nuclei based on a functional activation map extremely difficult [7]. Moreover, the inferior location of the brainstem causes fMRI data acquisition to be highly vulnerable to image artifacts (e.g., geometric distortions) and physiological noise (i.e., artifactual temporal signal fluctuations) [8]. New methods that are able to provide high resolution brainstem fMRI with high image fidelity are therefore desirable and may provide new insights in the role of key brainstem structures in the wide variety of functions the brainstem is involved in.

Although the main workhorse of fMRI is Gradient-Echo Echo-Planar Imaging (GRE-EPI), there has been renewed interest in 3D acquisition strategies; mainly to overcome some of the limitations in cortical fMRI at higher magnetic field strengths (e.g., 7 Tesla) [9]. The same properties that make these types of sequences suitable for high field fMRI are potentially advantageous for high resolution brainstem fMRI.

Balanced steady-state free-precession (bSSFP) fMRI has been successfully applied in high resolution cortical fMRI [10, 11]. The advantage of bSSFP fMRI is that functional contrast is achieved at short echo time (TE), thereby decoupling functional contrast from sources of image artifacts such as signal dropout and image distortion (as short readouts are inherent to this sequence). One potential drawback is the reduced functional sensitivity of bSSFP compared to conventional (long TE) GRE-EPI. At 3 Tesla the relative signal changes in response to a visual stimulus have been reported to be $\sim 2\%$ for bSSFP compared to $\sim 4\%$ for GRE-EPI [12, 13].

However, the contrast-to-noise ratio (CNR), which ultimately determines the

statistical power of the fMRI experiment, is not only determined by the functional sensitivity, but also by the temporal stability of the signal. The CNR is determined by the product of the relative signal change, $\Delta S/\bar{S}$, and the signal-to-noise ratio (tSNR), \bar{S}/σ : $\text{CNR} = \Delta S/\sigma = \Delta S/\bar{S} \times \bar{S}/\sigma$, where ΔS is the signal difference between rest and activity, \bar{S} is the mean signal over time, and σ is the temporal noise. In order to overcome the reduced functional sensitivity, and thus achieve comparable CNR, the tSNR in bSSFP needs to be roughly twice that of GRE-EPI.

Due to comparable tSNR in cortical areas, the CNR of bSSFP is generally lower compared to GRE-EPI for most visual studies (although sufficient for reliably detecting activation) [12]. A recent study by Miller *et al.* [13], however, showed that the sensitivity to physiological noise in bSSFP is TE dependent. Unlike GRE, which shows a reduction in CNR when the TE is reduced, bSSFP demonstrates a fairly constant CNR over the entire range of tested TEs (Fig. 1.1(a)), even though the contrast for both sequences decreases for short TE (Fig. 1.1(b)). The dissimilarity in CNR between GRE and bSSFP can be explained by the reduced noise sensitivity for bSSFP at short TE (λ is a measure of the noise sensitivity) as shown in Fig. 1.1(c). We therefore hypothesize that bSSFP fMRI will be advantageous over GRE-EPI in areas where the physiological noise is expected to be larger than in the cortex. The combination of reduced physiological noise sensitivity and minimal signal dropout is expected to result in a favourable tSNR for bSSFP. bSSFP may thus provide a comparable or even higher CNR compared to GRE-EPI in the brainstem, while offering low distortion images.

1.2 Thesis Outline

This thesis describes developments of 3D acquisition strategies and reconstruction methods to improve the quality of brainstem fMRI. Chapter 2 and 3 provide




Figure 1.1: Plots illustrating the differences in physiological noise sensitivity between GRE and bSSFP. Panel (a) shows the contrast-to-noise ratio given as the z -statistic for GRE and bSSFP over a range of TRs (the TE at which the data were acquired is approximately half the TR value). Panel (b) shows the functional contrast (relative signal change in %) and Panel (c) the calculated noise sensitivity measure λ . Note that low λ depicts low sensitivity to physiological noise. (figure courtesy of Dr. Karla L. Miller)

an overview of the background relevant to the work presented in this thesis, including brainstem anatomy and physiology, the physical principles underlying MRI, image acquisition strategies, image reconstruction methods, and steady-state FMRI sequences.

In Chapter 4 a characterization of two related 3D FMRI methods (passband bSSFP and RF-spoiled SPGR) is performed in order to assess their potential for brainstem FMRI. Due to the significant departure from traditional GRE-EPI in terms of functional contrast and k -space acquisition, each experiment aims to isolate a limited range of relevant sequence properties.

Chapter 5 presents a series of experiments that were aimed to identify the source of physiological signal fluctuations that are specific to 3D multi-shot readouts. Various potential correction methods are explored and validated with simulations. The limitations of each of the methods are discussed and a real-time synchronization method is identified for development *in vivo*.

A method to prospectively correct for physiological fluctuations in 3D acquisition methods is presented in Chapter 6. This method uses a real-time approach to order the k -space acquisition with respect to the cardiac cycle. It is shown

that this method significantly reduces cardiac related noise in 3D bSSFP and RF-spoiled SPGR.

Chapter 7 focusses on retrospective correction methods to further improve the temporal stability in 2D and 3D imaging sequences. Image based and k -space based correction techniques are optimized and compared, and a selection procedure based on Bayesian modeling comparison is used to identify the optimal model to correct the data. This chapter concludes with a comparison of the temporal stability between the developed 3D sequences and the “gold-standard” GRE-EPI sequence.

Finally, a summary of this work is given in Chapter 8. The merits of the developed techniques are discussed and prospects for future research are given.

Chapter 2

Principles of MR Imaging

This chapter covers the basic physics of magnetic resonance (MR) and imaging principles at a level suitable for understanding the ideas and methods presented in this thesis. Although the quantum mechanical description is used briefly to explain the concept of magnetic spin, the majority will be explained using the classical description. The classical view provides an intuitive description of the MR phenomenon, which is often more useful in the context of understanding MR imaging concepts than the quantum mechanical description [14]. For a more comprehensive background the reader is kindly referred to Refs. [15, 16, 17, 18, 19], which were valuable sources of information used to write this chapter.

2.1 The MR Signal

2.1.1 Magnetisation

The source of the MR signal arises from the interaction of particles that have a magnetic moment and two types of magnetic field: the static field B_0 , which causes polarisation and precession, and the radio frequency (RF) field B_1 , responsible for the excitation.

Polarisation

Charged particles, including hydrogen, possess a net spin angular momentum. This property gives rise to a magnetic dipole moment $\boldsymbol{\mu}$, given by:

$$\boldsymbol{\mu} = \gamma \hbar \mathbf{I} \quad (2.1)$$

where γ is the gyromagnetic ratio, \hbar is the Planck's constant divided by 2π and \mathbf{I} is the nuclear spin. The magnetic moment, $\boldsymbol{\mu}$, is a vector quantity with components parallel and orthogonal to the main static magnetic field B_0 . For the remainder of the chapter we will use classical descriptions, which are adequate to describe the MR phenomena relevant to MR imaging [14].

Fig. 2.1 shows the effect of the magnetic field on the distribution of spin orientations. Spins in a sample are in constant motion and each spin alters the magnetic field experienced by its neighbours. When no external field is applied (Panel (a)), the magnetic interactions cause the spins to point randomly in all directions. In the presence of an external magnetic field the spins have a preference to align towards the direction of the magnetic field. The energy levels associated with the alignment, however, are small compared to the thermal energies and thus the tendency to align with the field is small. The distribution of the spins therefore demonstrates only a small skew towards the direction of the external magnetic field as shown by Panel (b), which is also implied by the Boltzman distribution.

The small excess in spins parallel to the magnetic field creates a net magnetisation moment along B_0 . The net magnetisation $\mathbf{M} = \sum \boldsymbol{\mu}$, represents the macroscopic MR effect and can be depicted as a vector in 3D space (depicted by the large vertical arrow in Panel (b)). The net magnetisation vector \mathbf{M} is directed along the z -axis, usually referred to as the longitudinal axis. The plane perpen-




Figure 2.1: (a) In the absence of a magnetic field the directions of the spins are randomly distributed. (b) The distribution of the spins in a magnetic field is slightly skewed to the axis of the applied field, which creates a net magnetisation \mathbf{M} (large vertical arrow). The spins also undergo precession about the z -axis due to the torque force applied by B_0 (depicted by the white arrow). (Remake of Figs. 1 and 3 in [14], courtesy of Dr. K.L. Miller)

dicular to the longitudinal axis is called the transverse plane.

Precession

The second effect that the static field has on the nuclear spins is precession. Precession is caused by a torque force that is applied to $\boldsymbol{\mu}$ in the presence of B_0 (Fig. 2.1(b)). This torque force causes a change in direction of the magnetic moment that is given by:

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

The torque forces will thus cause a rotation about the axis parallel to the magnetic field at a rate that is given by the resonance frequency. For the external field B_0 , the resonance frequency ω_0 is given by:

$$\omega_0 = \gamma B_0 \quad (2.3)$$

The resonance frequency, also known as the Larmor frequency, is directly proportional to the magnetic field, and related by γ , which for ^1H is 42.58 MHz/T [18].

Excitation

At thermal equilibrium, \mathbf{M} is pointed along the longitudinal axis and the distribution of the spins in the transverse plane is even. In order to detect signal a net transverse magnetisation needs to be created. This is achieved by applying a second magnetic field, orthogonal to the main field, and rotating at the Larmor frequency. This causes the magnetisation to also precess around this additional field towards the transverse plane. Upon excitation the magnetisation vector \mathbf{M} will follow a helical path towards the transverse plane. Fig. 2.2(a) shows the helical path of \mathbf{M} for a flip-angle α of 45° . Fig. 2.2(b) shows the same excitation in a frame of reference that is rotating about the z -axis at the Larmor frequency for a simpler visualisation (the x and y are replaced by x' and y' in the rotating frame of reference).

Signal Detection

After excitation the net magnetisation will cause a fluctuating magnetic field due to its precession. The fluctuating magnetic field will produce a magnetic flux in the receiver coils that are placed orthogonal to the main magnetic field. In textbook examples usually two receiver coils are drawn: one along the x -axis, picking up the real signal; and one along the y -axis, picking up the imaginary signal [18]. In reality, however, the number of receiver coils can vary, and the real




Figure 2.2: The path of the net magnetisation vector, \mathbf{M} as it is tipped through 45° by an RF pulse in (a) the laboratory frame of reference and (b) the rotating frame of reference

and imaginary parts of the signal are formed when the signal is demodulated from the Larmor frequency (see Ref. [15] for more details).

2.1.2 Relaxation

As mentioned above in thermal equilibrium the net magnetisation \mathbf{M} is aligned with the magnetic field \mathbf{B} . After excitation, the magnetisation will return to its equilibrium position, in a process called relaxation. Relaxation consists of two processes, longitudinal and transverse relaxation, which we will describe separately. During relaxation the transverse component will decay away while the longitudinal component recovers to its equilibrium magnitude, referred to as M_0 .

Longitudinal Relaxation: T_1

T_1 relaxation describes the re-alignment of the proton spins towards the direction of the magnetic field, after RF excitation. T_1 relaxation is caused by the magnetic field fluctuations at the Larmor frequency that protons experience produced by the magnetic moments of other nuclei. These can be either neighbouring hydro-

gen nuclei or other molecules (e.g., protein).

The field fluctuations act as very short RF pulses causing re-orientations of the magnetic moment. Due to the slight preference of the protons to align to the magnetic field the distribution will slowly return to its equilibrium state re-aligning the net magnetisation along the longitudinal axis. The rate of return is described by the time constant T_1 .

Transverse Relaxation: T_2

T_2 relaxation describes the “non-recoverable” loss in transverse magnetisation. The same field fluctuations responsible for T_1 -relaxation also induce transverse relaxation. However, in addition to the field fluctuations at the Larmor frequency, slowly varying field fluctuations also introduce a loss in phase coherence between the nuclei and thus contribute to the loss in transverse magnetisation. The time-constant that describes the rate of decay of the transverse magnetisation is T_2 . Because transverse relaxation is affected by both low frequency fluctuations and fluctuations at the resonance frequency, T_2 is typically much shorter than T_1 .

Enhanced Transverse Relaxation: T_2^*

Additional to field fluctuations caused by thermal motion of surrounding spins, spatially varying static field offsets (i.e., field inhomogeneities) also induce an amount of dephasing to the spins. The additional decay in transverse magnetisation associated with this effect is characterized by T_2' . Unlike T_2 -effects this effect is reversible. By applying a 180° RF pulse the phase of the magnetisation will be inverted. After this inversion the spins will continue to precess and will rephase producing an echo of the original signal. Such RF pulses are therefore called refocusing pulses. Pulse sequences that include a refocusing pulse are generally referred to as Spin-Echo (SE) sequences [20]. The combination of T_2 and

T_2' , finally, is described by a time-constant called T_2^* , which is defined as:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \quad (2.4)$$

Relaxation plays an important role in the contrast observed on MR images, as relaxation times are highly tissue dependent. The different brain tissues, like grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF), each have unique relaxation times. The table below lists the T_1 , T_2 , and T_2^* -relaxation times of each tissue type at 3 Tesla.

<i>Tissue</i>	$T_1(ms)$	$T_2(ms)$	$T_2^*(ms)$
GM	1330 [21]	110 [21]	45 [21]
WM	830 [21]	80 [21]	45 [21]
CSF	4300 [22]	503 [23]	-

Table 2.1: MR relaxation times at 3.0T

2.1.3 The Bloch Equation

The Bloch equation [24] is a concise way of relating the time evolution of the magnetisation to the external fields and relaxation by means of a single differential equation:

$$\frac{d\mathbf{M}}{dt} = \mathbf{M} \times \gamma \mathbf{B} - \frac{M_x \mathbf{x} + M_y \mathbf{y}}{T_2} + \frac{(M_0 - M_z) \mathbf{z}}{T_1} \quad (2.5)$$

where \mathbf{M} is the magnetisation, \mathbf{B} the magnetic field, and T_1 and T_2 the longitudinal and transverse relaxation times, respectively. \mathbf{x} , \mathbf{y} , and \mathbf{z} are unit vectors in x , y , and z . The Bloch equation was first described in 1946, but still remains a powerful tool in pulse sequence development.

2.2 Image Formation

In the previous sections, the origin of the MR signal was discussed. This section describes the basic methods that allow us to determine the location of the protons in order to construct images of the measured object. In addition, specific acquisition strategies commonly used in neuro-imaging and crucial for the methods developed in this thesis are outlined.

2.2.1 Gradient Fields

All the spatial localisation methods rely on the ability to manipulate the z -component of \mathbf{B} , the main magnetic field. Most scanners contain three gradient coils that are able to generate linear gradient fields along the x , y , and z directions called G_x , G_y , and G_z , respectively. Note, that only the z -component of \mathbf{B} is altered. The gradient coils do not generate magnetic fields along x and y . For example, G_x describes the slope of the z -component of the main magnetic field along the direction x , given by dB_z/dx . In an MRI experiment the gradient fields are varied over time so one generally refers to gradient “pulses”. The gradient strength is reported in mT/m and is small compared to the main magnetic field. The maximum gradient strength for a human 3.0 Tesla scanner is typically ~ 40 mT/m. The local magnetic field in the presence of all three gradient fields is:

$$\mathbf{B}(\mathbf{r}, t) = \mathbf{B}_0 + G_x(t)x + G_y(t)y + G_z(t)z \quad (2.6)$$

where $\mathbf{r}=(x,y,z)$ is the vector representing the spatial location and B_0 the main magnetic field. The effect of changing the magnetic field locally is that the resonance frequency is now spatially dependent. The change in resonance frequency in the presence of a gradient field is:

$$\Delta\omega(\mathbf{r}, t) = \gamma\mathbf{G}(t) \cdot \mathbf{r} \quad (2.7)$$

which plays a key role in the spatial localisation techniques described below.

2.2.2 Slice Selection

As discussed above the net magnetisation can be tipped towards the transverse plane by applying an RF pulse provided its frequency corresponds to the resonance frequency of the spins. In the presence of B_0 , without linear gradient fields, all the spins precess at the same frequency (assuming a perfectly homogeneous field) and will all experience the same amount of excitation when an RF pulse is applied at the Larmor frequency ω_0 . Selective excitation can be achieved by applying an RF pulse while simultaneously applying a linear gradient field \mathbf{G} . For example, by turning on G_z we introduce a variation in the precession frequency along the z -axis. By altering the frequency of the RF pulse the excitation can be confined to only those spins that are at a particular location along z . In order to excite a slice with a thickness Δz and thus spins at a range of precession frequencies, an RF pulse is applied with a certain frequency bandwidth rather than a single frequency as shown in Fig. 2.3.

RF pulses are commonly sinc-shaped with a carrier frequency that corresponds with the resonance frequency of the spins at the centre of the slice of interest. The frequency bandwidth of the pulse, and thus excitation profile, can be approximated by the Fourier transform of the RF envelope (assuming small tip angles). A perfect rectangular slice profile, however, requires an infinitely long sinc function in the time-domain. Due to the finite length of RF pulses the achieved slice profile in practice will always contain a certain amount of imperfections in the form of

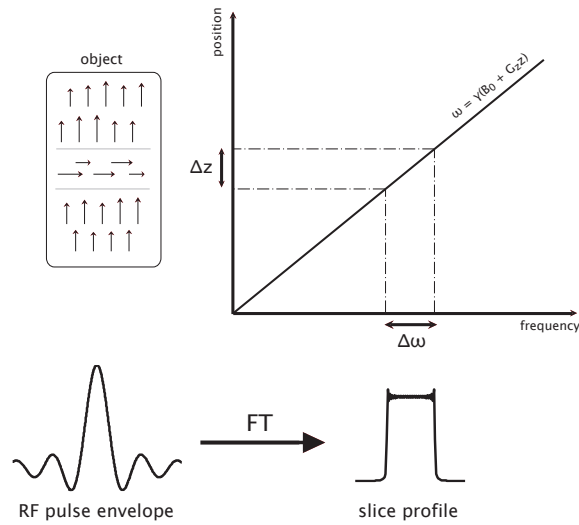


Figure 2.3: Illustration depicting slice excitation. The RF pulse envelope determines the RF bandwidth in the frequency domain. By applying a linear gradient along z only the protons at a particular location are excited, leaving the protons that are precessing outside the frequency bandwidth of the RF pulse unaffected.

ripples in the flat top region or even additional side lobes outside (Fig. 2.3). A measure of the selectivity of the pulse (i.e., the quality of the slice profile) is the dimensionless time-bandwidth product (TBW) of the RF pulse [16]. The TBW is simply $T\Delta f$, where T is duration of the pulse and Δf the bandwidth, which can be estimated by the inverse of the full-width-half-maximum (FWHM) of the central lobe of the sinc function. For slice selective pulses a TBW between 2 and 8 is typically used.

2.2.3 k -space

After selective excitation the field gradients are used for spatial localisation within the excited slice. As discussed earlier the field gradients allow us to change the local precession frequency. The local frequency offset at a given position r is given by Eq. 2.7. Under the influence of gradient fields the MR signal will thus consist of a range of frequency components that is dependent on the positions of the measured protons.

It is useful to express the MR signal in terms of the phase of the spins at a particular position in the presence of time-varying field gradients. The accumulated phase at position \mathbf{r} is the dot product of the time integral of the gradient with position:

$$\varphi(\mathbf{r}, t) = \gamma \int_0^t \mathbf{G}(\tau) \cdot \mathbf{r} d\tau \quad (2.8)$$

Assuming a receive coil with uniform sensitivity across the slice, the measured MR signal is the sum of all the magnetisation vectors that were excited. Neglecting relaxation terms the MR signal can be expressed as:

$$S(t) = \int_{\mathbf{r}} \rho(\mathbf{r}) e^{-i\varphi(\mathbf{r}, t)} d\mathbf{r} \quad (2.9)$$

$$S(t) = \int_{\mathbf{r}} \rho(\mathbf{r}) e^{-i\gamma \int_0^t \mathbf{G}(\tau) \cdot \mathbf{r} d\tau} d\mathbf{r} \quad (2.10)$$

where $\rho(\mathbf{r})$ is the proton density across the excited slice, and $\varphi(\mathbf{r}, t)$ the phase of the spins at position \mathbf{r} and time t as derived in Eq. 2.8. We can further simplify this equation by defining:

$$\mathbf{k} = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(\tau) d\tau \quad (2.11)$$

where \mathbf{k} is the time integral of the gradient waveform \mathbf{G} . We can write the MR signal as a function of \mathbf{k} :

$$S(\mathbf{k}) = \int_{\mathbf{r}} \rho(\mathbf{r}) e^{-i2\pi\mathbf{k} \cdot \mathbf{r}} d\mathbf{r} \quad (2.12)$$

Note, that this equation represents a Fourier transform. Therefore, \mathbf{k} represents the coordinates in the spatial frequency domain, referred to as k -space. k -space is the native domain in which the MR signal is acquired. The image of the object ($\rho(\mathbf{r})$) is obtained by an inverse 2D or 3D Fourier transform of $S(\mathbf{k})$.

2.2.4 Readout Trajectories

Data collection in MRI is done by excitation of the spins by means of an RF pulse after which a set of gradient waveforms is used to sample k -space. The order in which k -space is sampled is referred to as the k -space trajectory. The trajectory that is followed is determined by the shape of the gradient pulses (since the k -value at a given time t is determined by the time integral of the gradient waveform). Note that k -space does not have to be entirely sampled following a single RF pulse, but can be acquired step wise over a series of RF pulses, where only a portion of k -space is acquired during each repetition time (TR). Although many trajectories exist, the specific trajectories that are relevant to the remainder of this thesis will be outlined in this subsection.

2D Trajectories

In 2DFT, only a single line is acquired after each excitation (Fig. 2.4(a)). After excitation, but before signal acquisition, a brief gradient is applied along one axis (typically the y -axis is designated for this). This gradient has a variable area from one TR to the next in such a way that it moves the readout to a different location along k_y in each TR. At the end of this gradient, the magnetisation has been prepared with a particular phase profile corresponding to the desired k_y . This gradient is referred to as the 'phase-encode' gradient. At the same time, a negative gradient is applied along the other in-plane axis (the x -axis) to move to the edge of the desired k -space coverage. During signal acquisition the frequency encoding

gradient G_x is turned on which moves the sampling along k_x collecting a single line of data. Because only a single line in k -space is acquired per excitation, each line readout need not be particularly fast (e.g., a few milliseconds, corresponding to a low bandwidth). Although this imaging technique is thus able to provide images with high signal-to-noise ratio (SNR), the large number of excitations that are required makes this technique fairly slow.

Echo-Planar-Imaging (EPI) was developed in 1977 by Sir Peter Mansfield and is still one of the fastest readout methods to date [25]. EPI differs from 2DFT imaging in the sense that multiple phase encode lines are acquired after a single RF excitation. In EPI the readout gradients are quickly reversed multiple times to generate a series of gradient echoes. Interspersed phase encode gradients are applied to spatially encode each gradient echo such that different k -space lines are sampled. This significantly increases acquisition speed.

When the EPI-readout is collected over multiple readouts, the readout is called a “multi-shot” or “segmented” EPI readout. An example of a segmented-EPI readout is shown by Fig. 2.4(b). With the current gradient hardware, however, it is also possible to sample the entire k -space from a single RF excitation as shown in Fig. 2.4(c). This is typically referred to as a single-shot EPI. Single-shot EPI allows a 2D image to be collected in only a few tens of milliseconds. A 3D object is typically imaged by exciting one slice at a time until the entire volume is collected. For applications in the brain the volume acquisition time (T_{vol}) is usually in the order of a few seconds. Because each slice is only excited once per volume acquisition, for typical volume acquisition rates, only the longitudinal magnetisation will form a steady-state, and the signal dynamics of this type of sequence are described by the equations given in Section 2.3.1 below.




Figure 2.4: Illustration of the relationship between the pulse sequence timing diagrams (top row) and the corresponding k -space trajectories. The examples shown are: 2DFT (a); segmented, or multi-shot EPI (b); and single-shot EPI (c). The color coding depicts lines acquired in separate readouts (i.e., shots). Dashed lines are the pre- and dephasers (only shown for the first shot in each sequence)

2.2.5 3D Trajectories

The 2D readouts presented above can be extended to a 3D readout in which a 3D k -space is collected, rather than multiple slices of 2D k -space matrices. In 3D imaging a thick slab covering the entire volume of interest is excited using Fourier encoding in all three dimensions for spatial localisation. A straightforward method to achieve this is by adding a second phase encode direction along the third dimension. Figure 2.5 depicts the extension of 2D segmented EPI to 3D 'stack-of-segmented EPI' readout, which is used extensively in this thesis. In its most straightforward form, the readout order takes place in a linear fashion. Starting with the most negative k_z -location and ending with the most positive k_z . Our conventional implementation acquires all segments for a given k_z -plane before moving to the next.

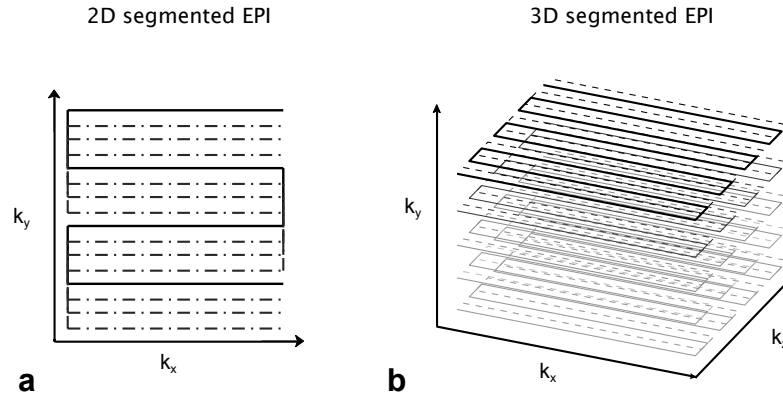


Figure 2.5: 2D (a) and 3D (b) segmented EPI trajectories. In 2D segmented EPI, a single slice is excited after which the k_x and k_y dimensions are collected using a segmented EPI trajectory. In 3D segmented EPI, the entire 3D volume is excited and the data are sampled in 3D k -space via a number of segmented EPI read-outs at multiple k_z locations.

2.3 Gradient Echo Sequences

Gradient echo (GRE) is a family of pulse sequences that are characterized by the gradient-refocused echoes they produce for signal formation. Unlike the SE sequences mentioned in the previous section, GRE sequences do not use refocusing RF pulses to form a spin-echo. First, an externally applied gradient in the magnetic field is used to dephase the spins, after which they are rephased using a gradient with opposite polarity to produce a gradient-echo [16]. GRE sequences are primarily used for fast scanning [26], and often operate under steady-state conditions.

2.3.1 Longitudinal Steady-State

As discussed above, in a typical MR imaging experiment the longitudinal magnetisation is tipped towards the transverse plane by an RF pulse, after which the MR signal is detected. Meanwhile the longitudinal magnetisation returns to its equilibrium state, M_0 , as an exponential recovery with time constant T_1 . However, allowing full recovery of the longitudinal magnetisation would require long TRs

leading to impracticably long scan times (e.g., a recovery to 95% of M_0 after a 90° excitation, requires $\sim 3 \times T_1$). In order to speed up the acquisition most imaging sequences employ a shorter TR. However, this means that the longitudinal magnetisation only partially recovers in between two RF pulses. In a train of RF pulses, such as in an imaging experiment, the longitudinal magnetisation will develop a steady-state; an equilibrium state where M_z is less than M_0 , but constant from one TR to the next. The steady-state can be defined as the condition when the relaxation processes are exactly cancelled by the RF pulse. Using the Bloch equation [24], we can calculate the steady-state signal by taking the magnetisation at any point in time, propagate it through one TR (accounting for RF pulses and relaxation effects), and then setting it equal to the magnetisation we started with.

In the simplest case, where T_2 is much shorter than TR, the transverse magnetisation has completely decayed away prior to the next RF pulse, and only the longitudinal steady-state magnetisation needs to be considered (assuming an infinitely short RF pulse). The effect of an RF pulse with a flip angle α at time t_n is given by:

$$M_z(t_n^+) = M_z(t_n^-) \cos(\alpha) \quad (2.13)$$

where $M_z(t_n^-)$ is the longitudinal magnetisation just prior and $M_z(t_n^+)$ the magnetisation immediately after the n^{th} RF pulse. During the TR interval, T_1 relaxation occurs, and thus the magnetisation prior to the next RF pulse is:

$$M_z(t_{n+1}^-) = M_z(t_n^+) e^{-T_R/T_1} + M_0(1 - e^{-T_R/T_1}) \quad (2.14)$$

Together these equations describe the time evolution of the magnetisation during a single TR. Substituting equation 2.13 into 2.14 and solving for $M_z^- = M_z(t_n^-) = M_z(t_{n+1}^-)$ then gives:

$$M_z^- = \frac{(1 - e^{-TR/T_1})}{1 - e^{-TR/T_1} \cos(\alpha)} M_0 \quad (2.15)$$

Note that this is the steady-state magnetisation prior to each RF pulse. The steady-state signal at an echo time TE is given by $M_{xy}(T_E) = M_z^- \sin(\alpha) e^{-T_E/T_2}$.

2.3.2 Transverse Steady-State

When the TR is of the order of T_2 the transverse magnetisation is not allowed enough time to fully decay to zero and will persist from one TR to the next. If both the longitudinal and the transverse relaxation are interrupted by RF pulses (i.e., if $TR \leq T_2 < T_1$) the pulse sequence tends to be referred to as a steady-state sequence. Steady-state sequences are often classified based on how the transverse magnetisation is dealt with at the end of the TR. When the transverse magnetisation develops into a steady-state, the sequence is classified as a *steady-state free-precession* (SSFP) sequence [16]. In this type of sequence the transverse magnetisation is effectively re-used in subsequent TRs leading to high signal levels (as outlined below). If, however, the transverse magnetisation is manipulated such that it does not contribute signal in subsequent TRs, the sequence is said to be *spoiled*. Spoiling aims to prevent residual magnetisation from creating echoes in subsequent acquisitions in order to increase the signal stability in short-TR sequences [27] or to create T_1 contrast [28]. Spoiled sequences, however, demonstrate smaller signal levels compared to SSFP sequences as the signal is purely determined by the magnitude of the longitudinal magnetisation just prior to the

excitation, which is small for short-TR sequences as shown by Eq. 2.15.

Spoiling

Two spoiling strategies are: 1) gradient spoiling and 2) RF spoiling. In gradient spoiling an additional gradient is applied that dephases the spins, which causes the transverse signal to decay (much like the T_2' -effect discussed above). The transverse magnetisation, however, is not destroyed and may be refocused in the subsequent TR producing a spin-echo-like signal. In order to prevent echoes from forming the gradient needs to be varied from one TR to the next, so that a different amount of dephasing is induced in each TR. However, even with varying gradients the efficacy of the spoiling will be spatially dependent, and ineffective at the centre of the gradients [27, 29].

A more effective strategy to spoil the transverse magnetisation is RF spoiling [27, 30, 31]. In RF spoiling the RF phase is varied from one TR to the next according to a predefined schedule. This introduces a phase offset between freshly excited magnetisation and the magnetisation that persists from previous excitations. If an appropriate phase cycling scheme is chosen, the net effect will be that the magnetisation from preceding TRs will phase cancel, while the freshly excited magnetisation (the magnetisation excited by the most recent RF pulse) will dominate [29]. A phase cycling scheme often used in RF spoiling is the quadratic phase cycling scheme proposed by Zur *et al.* [30]: $\varphi_n = n(n + 1)\Delta\varphi$, where φ_n is the phase of the n^{th} RF pulse and $\Delta\varphi$ the starting value that determines the size of the increments. A widely accepted value for $\Delta\varphi$ is 117° [30].

2.4 Steady-State Free Precession

The previous section discussed the case where the transverse magnetisation was zero at the end of each TR interval (either due to a long TR or due to spoiling). In this section the characteristics of steady-state free-precession (SSFP) sequences are discussed. Unlike the spoiled sequences, steady-state free precession acquisitions re-use the persisting transverse magnetisation to increase the signal and thus offer a particularly high SNR per unit of time [32]. SSFP sequences have a fixed TR with fixed dephasing gradients and a carefully chosen RF phase. Unlike spoiled sequences where the steady-state is driven by how much relaxation occurs within the TR, in SSFP sequences precession effects play an important role in the steady-state formation. In order to achieve a steady-state the RF pulse needs to cancel all sources of motion that the magnetisation experiences [33, 34]. This can be visualised as magnetisation vectors that follow an identical path each TR, ending up in the exact same position prior to each RF pulse. The exact path and magnitude of this vector is determined by the T_1 , T_2 , flip angle, and phase accrual during the TR, where phase accrual (either due to off-resonance precession or induced by gradients) is the dominant factor.

There are two types of SSFP sequences: balanced SSFP and non-balanced SSFP. For the purposes of this thesis, however, non-balanced SSFP sequences will not be covered in this chapter.

2.4.1 Balanced SSFP

Balanced-SSFP (bSSFP) is a special type of SSFP sequence in which the net time integral of gradient waveforms is exactly zero along every axis (i.e., the gradients are perfectly balanced), and the only source of phase accrual is off-resonance precession. As all imaging gradients are refocused, the magnetisation is a single

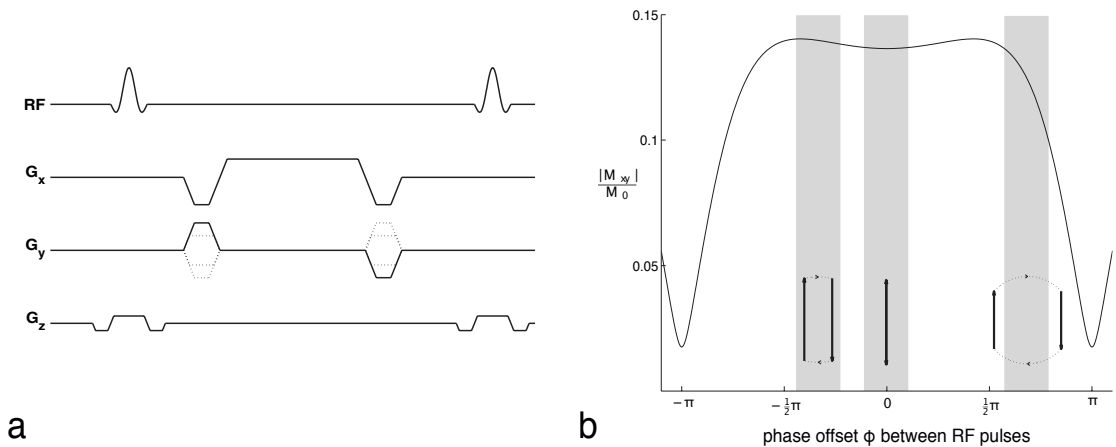


Figure 2.6: (a) the bSSFP timing diagram. All gradients have zero net gradient area, resulting in complete refocusing of the magnetisation at the end of each TR. Off-resonance effects, however, can introduce net phase accrual and thus a modulation of the steady-state. (b) shows the bSSFP signal as a function of phase offset in between RF-pulses. Projections of the magnetisation trajectory at three different offsets ($1/3\pi$, 0 , and $2/3\pi$) for a sequence with RF excitation with alternating phase ($-\alpha$, α , $-\alpha$) are shown at the bottom. Solid arrows depict RF flips, and the dashed lines describe precession between RF pulses (Figures inspired by Scheffler *et al.* 2003).

vector (i.e., is not dephased) at the end of the TR (Fig 2.6(a)).

Balanced SSFP normally consists of a train of alternating excitation pulses ($+\alpha$, $-\alpha$, $+\alpha$, $-\alpha$, ...). For on resonance spins that accrue no phase during the TR ($\varphi=0$), the magnetisation mostly rotates around the y-axis. However, any frequency offset, ν , with respect to the RF transmission frequency will cause the spins to accrue a phase of $\varphi = 2\pi\nu\text{TR}$ within two subsequent RF pulses. Phase accrual causes the angle between the RF pulse and the magnetisation to be different, so the RF pulse will tip the magnetisation to a very different part of the 3D magnetisation space. The longitudinal and transverse components of the steady-state can vary strongly depending on the precession angle. The bSSFP signal is therefore strongly dependent on local resonance frequency. Figure 2.6(b) shows the amplitude of the bSSFP signal as a function of off-resonance induced dephasing φ within the TR along with transverse projections of the magnetisation trajectories of spins at different resonance frequencies offsets (grey areas). Note that this is a 2π periodic

profile [35].

At certain phase offsets the signal profile exhibits a sharp decrease in signal, which are responsible for the well-known banding artifacts in SSFP. The location of these bands can be controlled by incrementing the RF phase by a fixed amount each TR (where the standard case described so far uses a π phase increment). This effectively shifts the signal profile relative to the absolute frequency of the spins. For example, a spin that accumulates a phase of 0.3π each TR in an experiment with no RF phase increment will end up in the same steady-state as a spin that precesses 0.7π each TR in an experiment with an RF phase increment of 0.4π . In both cases the net phase accrual is 0.3π relative to the RF excitation axis.

2.4.2 Balanced-SSFP Signal Equations

We can calculate the steady-state signal by equating the magnetisation in two corresponding time points in adjacent TR intervals. Because the transverse magnetisation persists at the end of each TR, we have to consider the magnetisation vector in all three dimensions. To allow a concise description of the equations in three dimensions a matrix formalism is used [36]. All interactions, such as excitation and relaxation are represented as matrix rotations and multiplications on the 3D magnetisation vector $\mathbf{M} = [M_x \ M_y \ M_z]^T$. Excitation and precession can be described by rotations about x and z , respectively, and are represented here by the rotation matrices $\mathbf{R}_x(\alpha)$ and $\mathbf{R}_z(\phi)$. The magnetisation just after the n^{th} RF pulse is given by:

$$\mathbf{M}_n^+ = \mathbf{R}_x(\alpha)\mathbf{M}_n^- \tag{2.16}$$

where \mathbf{M}^- is the magnetisation prior to the RF pulse and $\mathbf{R}_x(\alpha)$ represents a rotation about the x-axis by an angle α . During the TR interval precession and relaxation occurs, so the magnetisation just prior to the following RF pulse is:

$$\mathbf{M}_{n+1}^- = \mathbf{R}_z(\varphi)\mathbf{E}(T_R)\mathbf{M}_n^+ + (\mathbf{I} - \mathbf{E}(T_R))\mathbf{M}_0 \quad (2.17)$$

$\mathbf{R}_z(\phi)$ is the rotation due to phase accrual during the TR interval, and \mathbf{I} is the identity matrix. Longitudinal and transverse relaxation are represented by the diagonal matrix \mathbf{E} :

$$\mathbf{E} = \begin{bmatrix} e^{-t/T_2} & 0 & 0 \\ 0 & e^{-t/T_2} & 0 \\ 0 & 0 & e^{-t/T_1} \end{bmatrix}$$

By setting $\mathbf{M}^- = \mathbf{M}_n^- = \mathbf{M}_{n+1}^-$ we calculate \mathbf{M}^- , the steady-state magnetisation immediately prior to the RF. Substitution of Eq. 2.16 into 2.17 gives

$$\begin{aligned} \mathbf{M}^- &= \mathbf{R}_z(\varphi)\mathbf{E}(T_R)\mathbf{R}_x(\alpha)\mathbf{M}^- + (\mathbf{I} - \mathbf{E}(T_R))\mathbf{M}_0. \\ &= [\mathbf{I} - \mathbf{R}_z(\varphi)\mathbf{E}(T_R)\mathbf{R}_x(\alpha)]^{-1}(\mathbf{I} - \mathbf{E}(T_R))\mathbf{M}_0. \end{aligned} \quad (2.18)$$

Finally, using Eq 2.16, the magnetisation immediately after the RF pulse is given by:

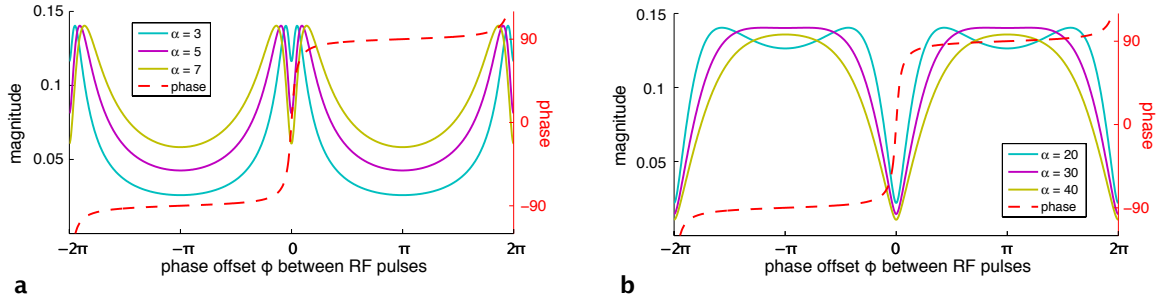


Figure 2.7: The steady-state signal magnitude (solid lines) and phase (dashed line) as a function of off-resonance frequency for a range of low (a) and high (b) flip angles. The portion of the profile, which exhibits a rapid phase change is referred to as the transition band, whereas the more constant portion of the profile is called the passband. Note that for these simulations, the RF pulses are applied with positive α only (i.e., not alternating between $+\alpha$ and $-\alpha$), which effectively shift the signal profiles by π compared to figure 2.6(b).

$$\mathbf{M}^+ = \mathbf{R}_x(\alpha)[\mathbf{I} - \mathbf{R}_z(\varphi)\mathbf{E}(T_R)\mathbf{R}_x(\alpha)]^{-1}(\mathbf{I} - \mathbf{E}(T_R))\mathbf{M}_0. \quad (2.19)$$

This expression shows that the signal amplitude has a strong dependency on off-resonance precession and flip angle. Figure 2.7 shows the steady-state signal profile calculated for a range of off-resonance frequencies and flip angles.

2.5 Image Reconstruction

Once the k -space is collected the spatial frequency data are converted into an image by means of image reconstruction. In theory if all the sampling is done on a perfect rectilinear grid a simple inverse Fourier transform would suffice. In practice, however, a few extra steps are often needed; either to correct for slight imperfections or to allow for subsampling of k -space.

In this section a few reconstruction techniques that are specific to EPI readouts are outlined and two more general reconstruction techniques that allow for sub-

sampling of k -space are discussed, which were relevant to the methods developed in this thesis.

2.5.1 Standard Image Reconstruction for EPI

The most pronounced artifact in EPI, when not corrected for, is ghosting. Eddy currents or unbalanced time delays in the electronic components can cause the acquired data to be offset relative to the true centre of k -space [37]. In a 2DFT acquisition, where all the lines are traversed in the same direction, this will not pose any problems as all the lines will have the same offset. In EPI, however, where the odd and even lines are traversed in opposite direction this will result in a modulation across the phase encode direction in k -space creating the infamous Nyquist ghosts [38].

Several methods have been proposed to reduce inconsistencies between the odd and even lines in k -space in order to mitigate the ghosting artifacts. The method typically used on clinical MRI scanners corrects for phase offsets and displacements between the odd and even lines by acquiring a few additional lines with no phase encoding (i.e., $k_y=0$.) [39]. On the Siemens systems used for this work, three additional lines are acquired: two forward lines and one backward line. The shift between the odd and even k -space lines is estimated by determining the phase difference between the forward and backward reference lines in the (1D) image domain (note that a shift in k -space corresponds to a linear phase ramp in the image domain). The image domain method proves to be more robust and allows correction of sub-voxel shifts. This type of correction is therefore referred to as phase correction.

2.5.2 Partial Fourier Reconstruction

Partial Fourier reconstruction refers to the reconstruction of data where k -space is only partially acquired as shown by Fig. 2.8. These methods rely on the property of the Fourier transform that real functions have conjugate symmetry in spatial frequency space. If the imaged object is entirely real the uncollected data can be synthesized by reflecting conjugate data across the origin as illustrated in Fig. 2.8(a). In practice, however, changes in local resonance frequency due to B_0 inhomogeneities can cause the object to have a non-constant phase and thus not be entirely real as depicted in Panel (b). Partial Fourier acquisitions that want to exploit the conjugate phase symmetry therefore always need some form of phase correction in order to synthesize the missing data.

The most pragmatic approach is to separately reconstruct the symmetrically acquired central portion of k -space (Fig. 2.8(c); red box) to phase correct the image initially reconstructed from the partial k -space data. The phase corrected image is then transformed back into k -space where the missing data are synthesized using conjugate symmetry. However, the phase correction in image space will introduce errors in the k -space data that are obtained from the phase corrected data. Because the multiplication with the phase correction function in image space is effectively a convolution of its Fourier transform in k -space, errors will arise near the boundary of the acquired data and the zero data.

Various iterative approaches have been suggested to reduce this effect. One of the methods is the POCS (for projection onto convex sets) algorithm. In POCS the partial k -space data are transformed into image space. This image is then forced to conform to the phase correction function obtained from the symmetrically acquired region at the centre of k -space. The phase constrained image is then transformed back into k -space to obtain the new estimate of the k -space data. At this point the lines that were actually acquired are replaced with the original




Figure 2.8: Illustration of a Partial Fourier acquisition. (a) Most PF reconstructions rely on the assumption of conjugate symmetry across k -space. (b) In practice, however, the phase variations cause the imaged object not to be entirely real. (c) Therefore, the acquisition requires more than half of Fourier space in order to perform phase correction using the central symmetrically region acquired at the centre of k -space (denoted by the red box). This example shows a 5/8 partial Fourier acquisition (the two back arrows depict the PF acquisition range)

k -space data and the process is repeated again. The process converges when the changes from one iteration to the next in the image domain are below the noise floor, which is usually after four or five iterations [40].

A method that is often applied due to its simplicity is to simply fill the missing data with zeroes. This, however, causes considerable blurring in image space. The way to understand this is by considering the data set as a fully acquired k -space data multiplied with a filter; in this case a step-function. The effect in image space is a convolution with the inverse Fourier transform of this filter (i.e., a sinc function with a linear phase ramp), which results in blurring.

2.5.3 Accelerated Parallel Imaging Techniques

During the early development of MRI, RF receive coils fell into one of two categories. Surface coils had a small receptive field, such that signal and noise only came from a small part of the anatomy, resulting in high SNR but limited coverage. Volume coils received signal and noise from a much larger volume, increasing coverage at the expense of SNR. In the early 2000s, MRI systems were fitted

with the hardware to receive from multiple coils simultaneously, and ‘array-coils’ consisting of multiple surface coils were developed that could achieve both high SNR and broad coverage [41]. It was soon realised that these coils had potential advantages beyond SNR, and that they could be used to aid reconstruction in a technique that has come to be known as ‘parallel imaging’.

Multi-channel receive coils can be used to accelerate the image acquisition by subsampling k -space (for example by skipping every other line) [42]. There exist various methods that utilize the different sensitivity profiles of the coils to allow for this undersampling that work either in k -space or in the image domain. Normally an acquisition where every other line is missing would result in an aliased image in the image domain. SENSE [43], the image domain variant, uses the fact that as long as the acquisition is subsampled with equidistant spacing the signal will alias in well defined locations. The aliased signal in the acquired images $S(x, y)$ can therefore be written as the sum of the signal from known locations in the unaliased image (i.e., (x, y) and $(x, y + FOV/R)$) weighted by the coil sensitivity profiles at these locations:

$$S_n(x, y) = C_n(x, y)\rho(x, y) + \sum_r^{R-1} C_n(x, y)\rho(x, y + r\frac{FOV}{R}) \quad (2.20)$$

where S_n is the measured signal of the n^{th} coil, C_n is the sensitivity profile per coil, R is the acceleration factor, and ρ is the unaliased image. As long as the coil profiles are different enough a reconstruction can be made.

Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [44], on the other hand, aims to fill in the missing data lines in k -space before the data are transformed into image space. The assumption that is made is that each point in k -space can be synthesized by a linear combination of the neighbouring

data points in the other coil images as long as their coil profiles are sufficiently different. In the image domain, the coils modulate the signal (i.e., the image is multiplied by each coil profile). The effect in k -space is a convolution with its Fourier transform. When acquiring a line of data, each coil therefore has effectively sampled a slightly different part of k -space, implying that more than one k -space line could be reconstructed from the data if one knew the relationship between the different coils.

GRAPPA takes advantage of the linearity of the coil effects (i.e., the convolution in k -space) and aims to find the linear coefficients (weights) between the coils necessary to effectively deconvolve the coil profiles and reconstruct multiple lines from different coil signals. The different weights between the target point (the missing point in k -space) and the source points (the acquired neighbouring data) are estimated by a reference scan, called the auto-calibration scan (ACS). This can either be a fully sampled central portion of the k -space or a separately acquired scan. Note, that the relationship between the source points and the target point is independent of the location in k -space, and therefore the weights that are estimated from the centre of k -space can be applied to the missing data points throughout the rest of k -space.

The GRAPPA kernel represents where the source points are in relation to the target point. The number of source points can vary, and the optimal value depends on the size of the ACS [45]. Typical values range between 6–20 source points. The GRAPPA kernel is estimated as the set of weights that is consistent with the ACS data in a least-squares sense by minimizing the error in reconstruction compared to measured values:

$$\mathbf{W} = (\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T \mathbf{T}_{CAL} \quad (2.21)$$

Here \mathbf{S} is a matrix with N_{ACS} rows, corresponding to the number of ACS points, and $n_c \times n_s$ columns (where n_c is the number of coils and n_s the number of source points). \mathbf{T}_{CAL} is a N_{ACS} by n_c matrix holding the target points. \mathbf{W} finally is a $n_c \times n_s$ by n_c matrix holding the estimated weights. Once the weights are estimated from the ACS data the kernel is applied to the subsampled k -space to fill in the missing data points:

$$\mathbf{T} = \mathbf{S}\mathbf{W} \quad (2.22)$$

where \mathbf{S} is a 1 by $n_c \times n_s$ matrix and \mathbf{T} a 1 by n_c matrix holding the estimated data points for each coil. GRAPPA thus reconstructs n_c fully sampled k -space matrices, corresponding to the estimated image as “seen” by a given coil (since the kernel was calculated specifically to estimate the data from one coil based on the signal in all the other coils). These are then transformed into image space and combined using, for example, a sum-of-squares or adaptive combine method [46].

Chapter 3

Functional MRI methods

This chapter presents the source of the BOLD contrast and various acquisition strategies to measure functional activation in the brain relevant to this thesis. The limitations inherent to each of the techniques are discussed as well as the complications inherent to brainstem FMRI.

3.1 The BOLD Contrast

Blood oxygenation level dependent (BOLD) imaging is a well established technique for the non-invasive investigation of brain function. In contrast to methods like electroencephalography (EEG) and magnetoencephalography (MEG), which measure direct electrical recordings, it relies on the haemodynamic response *associated* with neuronal activity. Due to the relatively slow vascular response the BOLD signal changes are blurred and delayed in time relative to the underlying neuronal activity.

BOLD FMRI is based on the fact that oxygenated haemoglobin (oxyHb) and deoxygenated haemoglobin (deoxyHb) have different magnetic properties. OxyHb is diamagnetic, whereas deoxyHb is paramagnetic. Fully deoxygenated blood therefore has a magnetic susceptibility, which is about 20% greater than fully

oxygenated blood [18]. As neurons become activated, their metabolic demands increase. The energy needed for this increase is provided by the vascular system. When a particular region in the brain becomes activated, oxygen demand in tissue increases, which is sourced from the blood by converting oxyHb into deoxygenated deoxyHb as depicted in Fig. 3.1(a). To compensate for the loss in oxygenated haemoglobin the local cerebral blood flow (CBF) and cerebral blood volume (CBV) is increased, resulting in a local *increase* of the oxyHb/deoxyHb ratio during brain activation (Fig. 3.1(b)). This increase is disproportionately large compared to the oxygen demand.

A higher oxyHb/deoxyHb ratio causes the local magnetic field to be less perturbed, which results in slower dephasing of the nearby spins and thus slower decay of the MR signal. The speed at which the signal decays is quantified by the relaxation rates R_2 and R_2^* (where R_2 and R_2^* are in relation to the transverse relaxation time constants as $1/T_2$ and $1/T_2^*$, respectively). So, using MR sequences sensitive to either T_2 or T_2^* one is able to detect differences in magnetic susceptibility and thus demonstrate differences in blood oxygenation [47]. The optimal TE at which to detect the R_2^* and R_2 changes is roughly equal to T_2^* and T_2 , respectively. For GRE BOLD at 3 Tesla this value is ~ 50 ms [48]. However, due to various limitations such as signal dropout (discussed below) GRE BOLD acquisitions typically use a TE of approximately 30 ms.

3.1.1 The fMRI Experiment and Data Analysis

In fMRI, brain activation is inferred from temporal changes in the MR signal that correlate to the performed task. In a typical fMRI experiment a series of images are collected while the subject is presented with a series of sensory stimuli or asked to perform a cognitive task interspersed with periods of rest. Due to the nature of the haemodynamic response the required temporal resolution is in

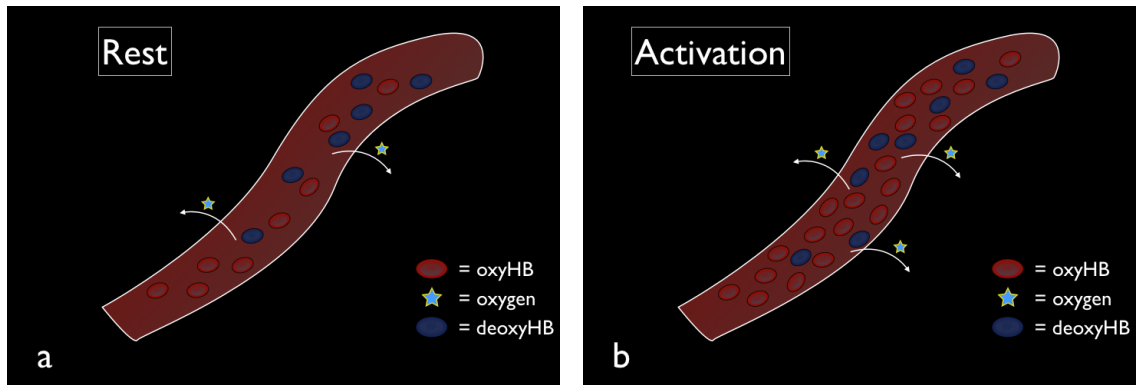


Figure 3.1: Diagram illustrating the haemodynamic response. In rest (**a**) a baseline amount of oxygen is delivered to the neurons converting a fraction of the oxyHb into deoxyHb. As a response to the increased energy demands during activation (**b**) the blood flow and blood volume increase, thereby overcompensating the energy demands. As a result the oxyHb/deoxyHb ratio is higher during activation.

the order of seconds and the required spatial resolution is in the order of a few millimetres [48].

Activation detection is done post-hoc by correlating the signal time course to the stimulus paradigm using a mathematical model called the General Linear Model (GLM). In linear modelling a model derived from the stimulus paradigm is used to fit to the data. In a simple on-off paradigm the stimulus is represented by a time-course vector holding zeros and ones, which represent rest and activation, respectively. To model the delay and blurring of the BOLD response with respect to the neuronal activity, the stimulus paradigm is convolved by the haemodynamic response function (HRF). The convolved time-course ($\underline{X}[t]$) is then fitted to the voxel time-series data ($\underline{Y}[t]$) by estimating the parameter estimate (PE), and the error of the model fit (ϵ).

In its simplest form the model then reads:

$$\underline{Y}[t] = PE \cdot \underline{X}[t] + \epsilon \quad (3.1)$$

To test whether the fit is significant, the PEs are divided by the error of the esti-

mate (as reflected in the residuals of the model fit), which results in a t -value. The obtained t -values are then converted to z - and p -values using standard statistical transformations and maps showing the z -statistics are generated that show the regions of activation.

3.2 Imaging Sequences for FMRI

In this section various acquisition strategies for FMRI are presented and their advantages and disadvantages are discussed.

3.2.1 2D Single-Shot GRE-EPI

Over the last two decades 2D single-shot EPI (Fig. 2.4(b)) has been the dominant technique for acquiring FMRI data; mostly in the form of a T_2^* -weighted experiment. We will refer to this technique as GRE-EPI throughout this thesis. The high efficiency of single-shot EPI allows the acquisition of whole brain images at a temporal and spatial resolution that is sufficient for most BOLD FMRI experiments. Moreover, with the current gradient hardware, the time at which the centre of k -space is acquired in single-shot EPI corresponds well with the optimal TE needed to maximise the BOLD contrast. For these reasons GRE-EPI has been a very successful acquisition technique for BOLD FMRI.

However, single-shot EPI acquisitions are particularly prone to distortion artifacts in regions that lie close to susceptibility boundaries (e.g., air/tissue interfaces) [49, 50, 51]. Spatial distortions arise when intrinsic field gradients (i.e., field inhomogeneities) cause local spins to precess at a slightly different rate than the expected precession rate based on the known applied gradients fields. Distortion artifacts are generally negligible in the readout direction due to the high bandwidth in this direction. In the phase encode direction, however, they can become

quite severe as the bandwidth in this direction is much lower. The severity of distortion artifacts (in the phase encode direction) is proportional to the total readout length, and thus increases with matrix size (e.g., when the resolution is increased or bandwidth is decreased).

Another limitation in GRE-EPI is signal dropout. The long echo time makes the acquisition sensitive to intra-voxel dephasing caused by other (non-functional) local field gradients. Signal dropout can be mitigated in several ways, for example: by reducing the TE, which allows less time for the spins to dephase; by reducing the voxel size, which effectively narrows the range of frequencies within a given voxel; or by using spin-echo EPI (SE-EPI), in which the static dephasing of the spins is refocused by an additional 180° RF pulse at TE/2. Unfortunately, however, all these approaches reduce the BOLD sensitivity of the acquisition as well.

3.2.2 3D FMRI

Recently 3D FMRI acquisitions have seen renewed interest [52, 53, 54] to overcome some of the limitations in 2D GRE-EPI. In contrast to 2D acquisitions, which use slice selective excitation to acquire the data in a slice-by-slice fashion, in 3D acquisitions a thick imaging slab, covering the entire field of view is excited. This has the advantage that the resolution in 3D acquisitions is not limited by imperfections in the excitation profile as in 2D.

In 3D imaging the time between excitations (TR) is often sufficiently short that the magnetisation does not fully recover between subsequent RF pulses (i.e., $TR \leq T_2 < T_1$). The data are thus generally acquired under steady-state conditions. Two steady-state sequences that are relevant to the methods developed in this thesis are spoiled gradient-echo (SPGR), which has GRE contrast, and balanced steady-state free-precession (bSSFP), which is discussed in the next subsection.

Due to the short TR, the magnetisation available for MR signal formation in these types of sequences is much lower (often only 5-15% of the magnetisation at equilibrium) than in 2D GRE-EPI. However, the fact that the signal is sampled from the entire volume with every excitation causes the intrinsic signal-to-noise ratio (SNR) to scale with $\sqrt{N_{k_z}}$, where N_{k_z} is the total number of phase encode steps in z . Further, steady-state sequences can employ highly efficient readouts (acquiring signal from a given spatial location for 60-90% of the TR, while 2D acquisitions in typical FMRI acquire signal from a given slice for 25–30 ms per TR=2–3 s) [29]. As a result, the image SNR in 3D acquired images is often comparable to, or higher than, the SNR in 2D EPI images.

3.2.3 Balanced SSFP FMRI

Over the past few years several methods have been proposed to detect functional activation using balanced SSFP. Balanced SSFP FMRI can be classified into two categories: transition-band bSSFP (tb-bSSFP) and passband bSSFP (pb-bSSFP). Earlier work used the transition band to achieve functional contrast [49, 55]. In tb-bSSFP the frequency sensitivity of the bSSFP signal profile is used to measure the deoxyHb frequency shift directly. To achieve this, the method matches the cortical region of interest to lie in the part of the bSSFP signal profile where the signal exhibits the largest signal transition (see Fig. 3.2). The proposed methods either exploited the transition in the magnitude profile (using a high flip-angle acquisition [55]) or the phase profile (using low flip-angle acquisitions [49]).

In this thesis, however, we will focus on the passband method, which utilizes the (flat) passband portion of the signal profile instead [56]. The functional contrast in passband SSFP does not rely on the sensitivity to off-resonance, but is determined by various, more complex, sources. The dominant source of contrast depends on the TE and TR that are used for the acquisition [13]. Studies have

shown that diffusion effects dominate the contrast at short TR [57], whereas at long TR T_2^* effects tend to dominate [12]. In order to assure that the data are acquired in the passband region of the profile it is important that a sufficiently good shim is achieved. Special care should also be taken to avoid temporal drifts in the resonance frequency caused by gradient heating or respiration. Methods to compensate for these effects include real-time correction for B_0 fluctuations [58].

The main advantage of bSSFP over GRE-EPI is that functional contrast is decoupled from image artifacts such as signal dropout and image distortion. The ability to achieve functional contrast at short TE (e.g., 3–6 ms) will reduce the signal dropout dramatically, whereas short repetition times allow highly segmented readouts, which mitigates distortion artifacts.

Because the remainder of the thesis will only focus on passband bSSFP FMRI we will refer to passband bSSFP simply as bSSFP in the following chapters for brevity.

3.3 Brainstem Anatomy and Physiology

The brainstem is located anterior to the cerebellum and connects the diencephalon to the spinal cord. Like the spinal cord the brainstem consists of a grey matter core surrounded by white matter fibre tracts. However, the brainstem contains additional grey matter nuclei embedded in the white matter [59]. The brainstem consists of three regions: the midbrain, the pons, and the medulla oblongata (see Fig. 3.3).

The midbrain is the most superior part of the brainstem. It includes the cerebral aqueduct, through which the CSF flows from the third to the fourth ventricle. The cerebral aqueduct is surrounded by periaqueductal grey matter, which is involved in the pain desensitization pathway. The brainstem also includes a number

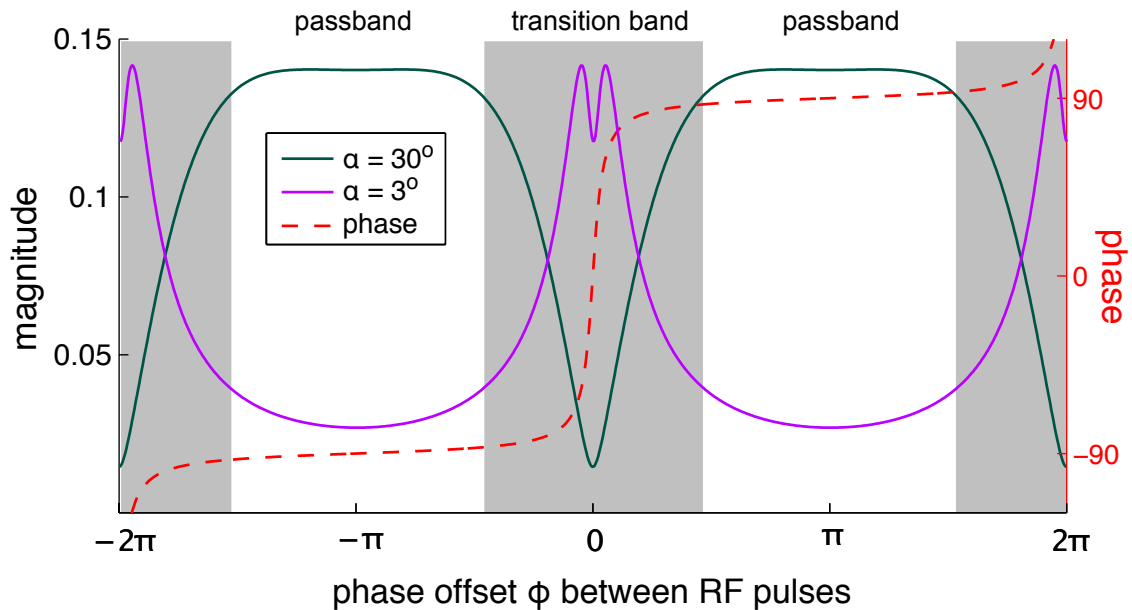


Figure 3.2: The steady-state signal magnitude (solid lines) and phase (dashed line) as a function of off-resonance frequency for a range of low for a flip angle of 30° (green) and 30° (pink). High flip-angle tb-bSSFP utilizes the magnitude variation to obtain functional contrast, whereas the long flip-angle acquisition takes advantage of the transition in phase. Passband bSSFP fMRI utilizes the flat portion of the profile to obtain a more BOLD-like contrast.

of nuclei such as the oculomotor nerve nucleus, the red nucleus and the superior and inferior colliculi. Their functions vary from acting as relay centres in motor pathways to visual reflex centres that control head and eye movement.

The pons is located inferior to the midbrain and encloses the fourth ventricle. The pons is the bulging structure of the brainstem and forms a highly connected hub, containing a broad range of nuclei that represent synapses connecting spinal, cortical and cerebellar neurons as well as a number of cranial nerves.

The medulla oblongata is the most inferior part of the brainstem. From superior to inferior the fourth ventricle narrows to form the central canal, which continues into the spinal cord. The ventral surface of the medulla consists of the pyramids, formed by pyramidal tracts descending from the motor cortex. At the medulla-spine transition these fibers cross over to the contra-lateral side causing each hemisphere to control the voluntary movements of muscles on the opposite side of the body. The medulla contains several important visceral motor cen-

tres such as the cardiovascular centre, the respiratory centres and several reflex centres, which regulate vomiting, swallowing, and coughing.

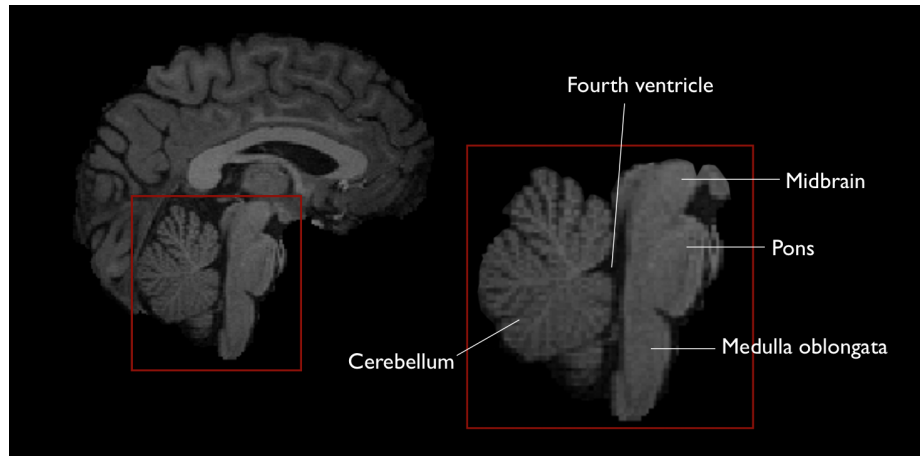


Figure 3.3: Mid-sagittal section of the brain illustrating the brainstem structures. From superior to inferior the three structures are: the Midbrain (i.e., Mesencephalon), the Pons, and the Medulla oblongata.

3.4 Practical Considerations Specific to Brainstem FMRI

Although the number of studies that include brainstem FMRI is increasing [60, 61, 62, 3, 63], the reliability and accuracy is still limited in comparison to the cortical regions of the brain (e.g., occipital lobe). Accurate functional localisation of the various nuclei is difficult as brainstem FMRI is impaired by a number of methodological and anatomical complications [7, 64].

3.4.1 Resolution

The brainstem is a relatively small structure that holds many different functional structures. The small size of these nuclei and the fact that they are closely spaced make functional localisation difficult. The lateral lemniscus and the superior olivary complexes, for example, only measure $2 \times 2 \times 5 \text{ mm}^3$ and $2 \times 2 \times 2 \text{ mm}^3$, respectively. At a typical imaging resolution of $3 \times 3 \times 3 \text{ mm}^3$, partial voluming ef-

fects (i.e., when a voxel represents more than one tissue type) are likely, which impairs both identification and detectability. In order to distinguish the closely spaced nuclei in the brainstem ideally an isotropic resolution of 1 - 1.5 mm would be needed. A complication in high resolution imaging, however, is that the SNR rapidly drops as the resolution is increased. Increasing the resolution from $3 \times 3 \times 3$ mm³ to $1.5 \times 1.5 \times 1.5$ mm³ results in a 8-fold SNR drop as the amount of available signal scales with the voxel volume.

Another limitation is the minimum slice thickness that can be achieved. Excitation of a thin 2D slice requires a large slice selection gradient in combination with a narrow RF bandwidth. However, the gradient amplitude is hardware limited and can thus not be increased arbitrarily. The slice profile of the excited slice is determined by the time-bandwidth product (TBW) of the RF pulse. For thin slices impractically long RF pulses may be required to achieve the same time-bandwidth product [16].

3.4.2 Distortion

Due to its inferior location the brainstem is more prone to susceptibility artifacts than the cortical regions in the brain. Various structures with a different magnetic susceptibility are present in close proximity of the brainstem. The two nasal sinuses (e.g., the ethmoid and sphenoidal sinus) located anterior to the brainstem and the auditory canals on both sides lateral to the brainstem produce large field inhomogeneities across the brainstem. The field offsets introduced by these structures cause considerable distortions in single-shot EPI. Obviously the impaired image fidelity complicates accurate localisation of the small functional structures in the brainstem. Distortion artifacts scale with resolution in single-shot acquisitions. When imaging at higher resolution, more phase-encode lines are required, which lengthens the readout and therefore increases distortion artifacts. Distor-

tion artifacts are therefore an additional limiting factor on the maximum achievable resolution.

3.4.3 Physiological Noise

Physiological noise refers to artifactual signal fluctuations due to metabolic and vascular responses. Physiological noise deteriorates the temporal signal-to-noise ratio (tSNR) in the brainstem, which directly reduces the contrast-to-noise ratio of the BOLD signal. The main components of physiological noise are respiratory and cardiac in origin [65]. Respiration creates small magnetic field shifts, caused by gross magnetic susceptibility changes in the lungs [66]. Wowk *et al.* [67] have reported B_0 field shifts over the respiratory cycle that increase with decreasing distance to the chest. Deeper brain regions such as the brainstem are therefore more affected by respiratory-related noise. Additionally, cerebral blood volume and flow fluctuations introduce cardiac related noise [66]. Due to the vertebral arteries, which run alongside the brainstem, cardiac motion is a considerable source of artifact. Enzman [68] reports bulk motions due to cardiac pulsation of 0.5 mm in the brainstem compared to less than 0.05 mm in cortical regions.

3.5 Retrospective Correction of Physiological Noise

Several methods have been proposed that take advantage of the fact that most physiological fluctuations are quasi-periodic. One approach is to use temporal filters [69, 70] in order to suppress fluctuations at the cardiac and respiratory fluctuations. However, this type of correction fails when the spectra of the physiological fluctuations alias into the frequency band of the stimulus paradigm [71]. Other methods aim to estimate the physiological fluctuations from the acquired MR data by estimating the global phase of the k -space [67], or use principal

component analysis in the image domain [72]. Another approach that is commonly used is to estimate the physiological fluctuations based on independently acquired waveforms of the cardiac and respiratory cycle (e.g., with the aid of a pulse-oximeter and pneumatic bellows) [65, 71, 73].

In its most basic form, low order Fourier series are generated based on phase of the physiological cycles during the time of each acquisition to serve as confound regressors. These regressors are then fitted to the data to estimate the artifactual fluctuations. The correction can either be performed in k -space [65] or image space [71]. This image based correction method was called RETROspective Image CORection (RETROICOR, and subsequently the k -space counterpart was dubbed RETROKCOR). Typically the artifactual fluctuations are estimated voxel-wise (or point-wise in the case of RETROKCOR) by means of a linear regression using the low order Fourier terms as confound regressors [74].

Apart from the effects directly related to the phase in the respiratory cycle, breathing has also been shown to affect the BOLD signal by modulating the $p\text{CO}_2$ in the brain [75]. The signal modulations caused by changes in respiration rate can be modelled by a so-called respiration volume per time (RVT) regressor [76]. The RVT regressor can be computed from the respiratory trace by dividing the difference signal (i.e., the difference between maximum in- and expiration) by the period of the respiratory. Similarly, fluctuations in heart rate have been suggested as a source of signal variability [77], and corresponding regressors have been proposed [78, 79, 77]. Finally, motion effects can be added to the physiological model [80]. For more details on how these regressors are derived the reader is referred to Appendix A.

Chapter 4

Characterization of 3D Sequences for Brainstem FMRI: Preliminary Investigations

4.1 Introduction

The overall hypothesis of this thesis is that 3D, steady-state, FMRI methods may be superior to conventional GRE-EPI in the brainstem due to reduced distortion and lower sensitivity to physiological fluctuations. However, these benefits will need to be sufficient to overcome the reduced contrast levels of steady-state FMRI. This chapter presents a series of preliminary investigations conducted in the first year of my DPhil research that aim to study the performance of two related 3D FMRI methods (passband bSSFP and RF-spoiled SPGR) in the brainstem. Due to the significant departure from traditional GRE-EPI in terms of functional contrast and k -space acquisition, as well as the difficulties of imaging in the brainstem, each experiment aimed to isolate a limited range of relevant sequence properties (broadly: distortion and physiological noise). By dissecting sequence behaviour

in this manner, these experiments served to guide the direction of sequence development in years 2 and 3. As such, the investigations of this chapter are brief but informative. In these experiments a strong emphasis is put on the assessment of tSNR in order to characterise the sequences. Together with the functional sensitivity (i.e., relative signal change), the tSNR determines the CNR of the sequence as discussed in the introductory chapters.

Investigation I: This experiment aimed to study whether prior observations of low physiological noise sensitivity in 3D bSSFP FMRI might relate to spin-echo-like signal characteristics. SE-EPI and GRE-EPI were compared in terms of physiological noise sensitivity in 5 subjects. Results were partly inconclusive, but suggest similar characteristics for SE and GRE, indicating the bSSFP signal behaviour is not simply reflective of its spin-echo-like properties.

Investigation II: This experiment aimed to characterize the sensitivity of GRE-EPI, bSSFP and SPGR to different physiological noise sources in the cortex and brainstem. To separate different physiological effects, rapidly-sampled, single-slice images were acquired in 5 subjects. Results indicated that bSSFP has improved physiological noise characteristics, although it was shown that the physiological noise increases with the segmentation factor of the readout.

Investigation III: This experiment aimed to investigate the distortion properties of our 3D methods using the hippocampus as a model system. The hippocampus suffers from distortion but does not have particularly high levels of physiological noise. Functional results demonstrated the ability to achieve low-distortion hippocampal FMRI using bSSFP. However, a certain degree of spurious activation was also observed in regions that lie in the vicinity to CSF.

Investigation IV: In the last experiment 2D GRE, 2D bSSFP, and 3D bSSFP brainstem data were acquired to assess the temporal stability in the brainstem when the data are acquired with a 3D readout. However, large instabilities were

observed for the 3D acquisition, which were much larger than anticipated. This indicates that the readout plays a large role in how physiological instabilities are expressed. This finding formed the basis to design a comprehensive set of experiments to investigate the physiological noise characteristics of 3D acquisitions, which are presented in the next chapter.

4.2 The Effect of Spin Echo Contrast on Physiological Noise

Balanced SSFP produces spin echo-like contrast at $TE=TR/2$ [81], which could be a potential reason for the decreased physiological noise sensitivity compared to GRE-EPI [13]. This experiment aims to isolate the effect of the signal contrast from other sequence properties such as the steady-state spin dynamics present in bSSFP. SE-EPI was therefore used to serve as a model for bSSFP. This sequence produces spin echo contrast, but is not a steady-state sequence and therefore uses imaging parameters (such as the readout and TR) that are more comparable to GRE-EPI.

In the cortex SE sequences have been shown to have different noise characteristics compared to GRE sequences. For example, the physiological noise in SE-EPI does not scale with the image signal level [82] in the way that GRE-EPI does [83]. In the brainstem, however, the relative contribution of the physiological noise components may be different. Due to its inferior location the brainstem is expected to be more susceptible to respiratory related B_0 fluctuations and its different vasculature may have an impact on the cardiac related pulsations. Here we present a noise comparison study where we compare the physiological noise of GRE-EPI and SE-EPI in the brainstem. In order to investigate any differences between the brainstem and the cortical regions, data were acquired coronally and

cortical grey matter ROI was included in the analysis.

4.2.1 The Temporal Noise Model

In conventional fMRI various sources of noise contribute to the total variance of the fMRI time-series signal. In the temporal noise model proposed by Krüger *et al.* [83] the various noise sources are split into two groups: the raw image noise, σ_0 , and physiological noise, σ_p . Here σ_0 includes random noise generated by the scanner electronics, thermal noise, and other system imperfections, whereas σ_p describes the noise arising from a physiological source, such as cardiac and respiratory related fluctuations, but also the BOLD-like fluctuations. Assuming the two noise sources are statistically independent the total noise can be written as: $\sigma_{tot} = \sqrt{\sigma_0^2 + \sigma_p^2}$. The temporal signal-to-noise ratio (tSNR) is then given by:

$$\text{tSNR} = \frac{\bar{S}}{\sqrt{\sigma_0^2 + \sigma_p^2}} \quad (4.1)$$

where \bar{S} is the mean image signal intensity. At this point it is important to note that the physiological noise component, σ_p , is proportional to the image signal level, whereas σ_0 is not. Rewriting σ_p as $\lambda\bar{S}$, where λ is a constant, and defining SNR_0 , the signal to noise ratio in a single image as $\text{SNR}_0 = \bar{S}/\sigma_0$ we can write:

$$\begin{aligned} \text{tSNR} &= \frac{\bar{S}}{\sqrt{\sigma_0^2 + \lambda^2\bar{S}^2}} \\ &= \frac{\text{SNR}_0}{\sqrt{1 + \lambda^2\text{SNR}_0^2}} \end{aligned} \quad (4.2)$$

This equation describes the well known Krüger temporal noise model, proposed

by Krüger *et al.* in 2001 [83]. By exploring the limits of this function we find that, when λ goes to zero:

$$\lim_{\lambda \rightarrow 0} \text{tSNR} = \lim_{\lambda \rightarrow 0} \frac{\text{SNR}_0}{\sqrt{1 + \lambda^2 \text{SNR}_0^2}} = \text{SNR}_0 \quad (4.3)$$

In other words: when $\lambda = 0$, the tSNR is solely determined by SNR_0 , and if we were to plot tSNR against SNR_0 all points would lie on the unity line (tSNR= SNR_0). On the other hand, when $\lambda \neq 0$, we see that for infinite SNR_0 :

$$\lim_{\lambda \rightarrow \infty} \text{tSNR} = \lim_{\text{SNR}_0 \rightarrow \infty} \frac{\text{SNR}_0}{\sqrt{1 + \lambda^2 \text{SNR}_0^2}} = 1/\lambda \quad (4.4)$$

The SNR_0 -tSNR curve will approach an asymptotic limit, where the value of the asymptote decreases for increasing λ . The constant λ is thus a scaling factor that denotes the sensitivity of the sequence to physiological noise and predicts the tSNR of the data with any given SNR_0 .

By acquiring data at a range of SNR_0 values (e.g., by changing the flip angle, resolution, field strength, or receiver coil) λ can be determined by fitting Eq. 4.2 to the data [83]. This method has been effectively used to characterize the physiological noise sensitivity of different pulse sequences [13, 84], readouts [85], acquisition parameters [86], or hardware options (e.g., multi-channel array coils) [87].

4.2.2 Methods

Subjects and Imaging Parameters

Data were acquired in five healthy volunteers (3 male, 2 female) on a 3T Siemens TIM Trio system (Siemens Medical Solutions, Erlangen, Germany), using a 12-

channel head coil. All subjects were scanned with informed consent under a technical development protocol approved by the local ethics committee. For each subject five coronal scans using single-shot GRE-EPI and five scans using single-shot SE-EPI were acquired with varying flip angle. As the thermal noise is independent of the flip angles, the signal level will determine the spread of SNR_0 values in our experiments. The signal levels between GRE-EPI and SE-EPI were therefore matched as much as possible. Based on Bloch simulations lower flip angles were used for GRE-EPI than for SE-EPI and a slightly shorter TE was used for SE-EPI than one would for maximising BOLD contrast.

Data were acquired using the following parameters for GRE-EPI: TE=30 ms, $\alpha=15^\circ, 30^\circ, 45^\circ, 60^\circ,$ and 75° . For SE-EPI we used: TE=40 ms, $\alpha=30^\circ, 45^\circ, 60^\circ, 75^\circ,$ and 90° . Other parameters that were common between the scans included: TR=1500 ms, FOV=190 cm, matrix= 96×96 , 9 slices, 2 mm slice thickness, Bandwidth=1628 Hz/pixel, 120 volumes per run. The distance factor between the slices was set to 100% for both GRE-EPI and SE-EPI. The phase encode direction was set to Right-Left (RL) in order to prevent wrap around in the Superior-Inferior (SI) direction.

For registration purposes a GRE fieldmap (TE₁ = 4.92 ms, TE₂ = 7.38 ms, resolution = $2 \times 2 \times 2$ mm) and a GRE-EPI whole brain acquisition were acquired (99 coronal slices, $\alpha = 75^\circ$, other parameters identical to the time series acquisitions, but no inter-slice gap). T₁-weighted structural scans of each subject were acquired previously in a separate scan session.

Analysis

For these initial experiments all reconstruction was performed online by the default algorithms available on the scanner. The reconstructed magnitude images were exported offline for further processing and data analysis using FSL [88] and Matlab (MathWorks, Natick, MA).

First, all data were brain-extracted using FSL BET [89], after which the time-series data were corrected for motion. With FLIRT [90] transformation matrices from the time-series data sets to the subjects' structural were determined using the whole brain data sets as an intermediate step in the registration to achieve accurate registration. However, instead of transforming the time-series into one common space (e.g., standard space) we transformed the ROIs, which were defined based on the subjects' structural, to the individual time series. By performing the analysis in the native space of the time series we avoid any bias in our noise estimates caused by upsampling and interpolation of data (which would occur had the data been registered to the higher resolution structural data).

For our SNR_0 and tSNR analyses two ROIs were defined. The cortical grey matter mask was generated by segmentation of the structural image with FSL FAST [91], whereas the brainstem mask was obtained using FSL FIRST [92]. Both tools output voxelwise probability maps, which report the likelihood that a given voxel belongs to the structure in question. The probability maps were transformed to the time-series spaces, after which they were thresholded to 0.8 (meaning that the included voxels all had an 80% or higher probability of belonging to the structure). The tSNR was determined voxelwise by dividing the temporal mean by the temporal standard deviation and then averaging over the ROI. SNR_0 is defined by the mean value across the ROI divided by the background noise. Background ROIs were manually drawn outside the brain, carefully avoiding areas of ghosting. To obtain an accurate measure of σ_0 in the magnitude data, the background noise histogram was fitted to a Rician distribution using a Levenberg-Marquardt algorithm (the `nlinfit` function in Matlab) as previously described in [13]. The obtained tSNR values were plotted against the SNR_0 values and the physiological noise model as given in Eq. 4.2 was fitted to the GRE-EPI and SE-EPI data in order to determine λ .

4.2.3 Results

The *in vivo* results are given in Fig. 4.1. Here the tSNR is plotted against SNR_0 in order to determine λ . Fig. 4.1(a) shows the results obtained in the grey matter ROI and Fig. 4.1(b) shows the brainstem results. It is seen that, when only the flip angle is changed in order to modulate the SNR_0 , the achieved range of SNR_0 is fairly small for both GRE-EPI (red) and SE-EPI (blue). In both ROIs the maximum SNR_0 achieved in SE-EPI is reduced compared to GRE-EPI. In the grey matter the maximum SNR_0 is 28 for SE-EPI, whereas an SNR_0 of 41 was achieved with GRE-EPI. In the brainstem the maximum values for SNR_0 were 22 and 33 for SE-EPI and GRE-EPI, respectively.

Although the curves deviate from the unity line in both ROIs and for both sequences, the curves do not reach a plateau. This suggests that the SNR_0 values were not sufficiently high to sample the asymptotic curve of the SNR_0 -tSNR curve and thus estimate λ accurately. Fitting the Krüger model to the cortical SNR measures returned a $\lambda=0.034$ for the GRE-EPI data and $\lambda=0.041$ for SE-EPI. In the brainstem we found $\lambda=0.046$ for GRE-EPI and $\lambda=0.065$ for SE-EPI. The confidence intervals on these fits ranged between 15%-20%.

4.2.4 Discussion

Figure 4.2(a) shows the signal as a function of flip angle for GRE-EPI as predicted by Bloch simulations. The simulations predict a slightly lower, but fairly comparable range of signal levels for the scan parameters used in this study. In our data however a clear difference in the range of SNR_0 was observed between GRE-EPI and SE-EPI. In particular, the SNR_0 of the SE-EPI data obtained with the higher flip angles is considerably smaller than the GRE-EPI data; an effect that is particularly noticeable in the brainstem. Another observation that was made

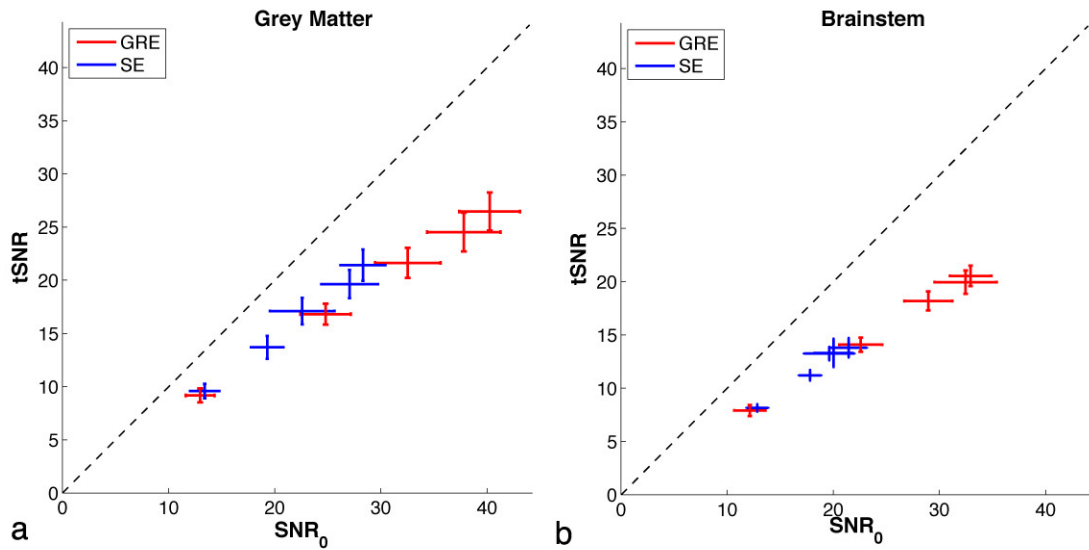


Figure 4.1: tSNR versus SNR_0 plots showing GRE-EPI (red) and SE-EPI data (blue) in the grey matter (a) and brainstem (b). The data are presented as the mean across the six subjects with error bars denoting the standard deviation in SNR_0 (horizontal bars) and tSNR (vertical bars). The dashed line indicates unity ($\text{tSNR} = \text{SNR}_0$). These plots show that, for the tested SNR_0 values, all the data lie in the linear part of the curve, where the asymptotic limit has not yet been reached.

during initial pilot scans was that the signal levels in SE-EPI are influenced by the inter-slice gap. For a distance factor of 100% or 200% (i.e., 1–2 slices gap between adjacent acquired slices) the signal in SE-EPI was comparable. When the distance factor was set to 0%, however, the SE-EPI signal dropped by 22%. This indicated a potential problem with the refocusing pulses used in this vendor-supplied sequence. The dependence of the signal on distance factor can occur when the refocusing pulse has a broader slice profile than the prescribed slice thickness and was saturating the adjacent slices.

Upon inspection of the RF pulses (the excitation pulse and refocusing pulse) it was found that both pulses are SINC pulses with a time-bandwidth product (TBW) of 2. This means that the slice excitation profile will not be entirely homogeneous across the slice. The bandwidth of the refocusing pulse was $1\frac{2}{3}$ wider than the bandwidth of the excitation pulse. The increased thickness of the refocusing profile explains why the signal was reduced when a slice gap of less

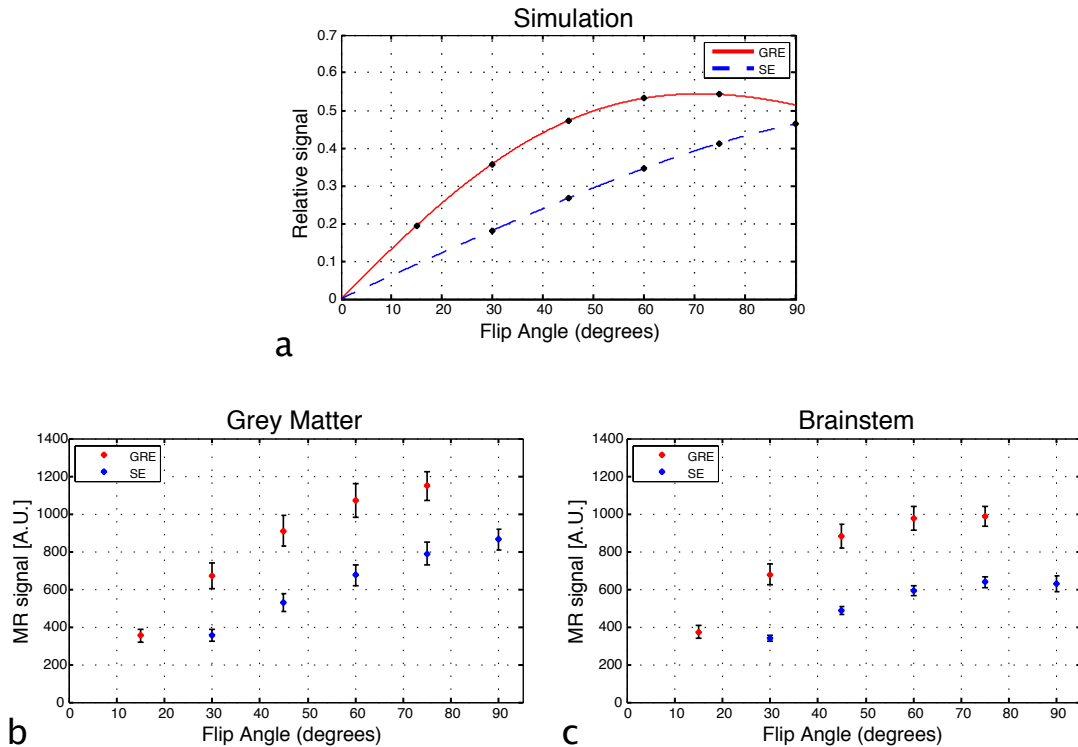


Figure 4.2: Simulated signal levels for GRE-EPI and SE-EPI (a) compared to *in vivo* signal levels in the grey matter (b) and brainstem (c). The T_1 , T_2 , and T_2^* values used in the simulations were: 1330 ms, 110 ms, and 45 ms [21]. The simulated scan parameters were $TR=1500$ ms, $TE_{GRE}=30$ ms, and $TE_{SE}=40$ ms). It is shown that for SE-EPI the signal levels are more reduced compared to GRE-EPI than predicted by our simulations. In the brainstem both sequences show a reduction in signal level at high flip angles.

than 100% was used. Most likely the bandwidth of the refocusing pulse is set wider than the excitation pulse in order to mitigate the effects of the poor refocusing profile in slice selective 180° SINC pulses [16]. However, it has been demonstrated that in order to achieve an acceptable refocusing profile with these type of pulses the bandwidth of the refocusing pulse needs to be set to at least 3 times the bandwidth of the excitation pulse [93]. The refocusing profile achieved with the default pulses is therefore not likely to be flat. This in combination with the fact that the 180° was surrounded by crusher gradients might well explain the reduced signal in SE-EPI compared to the GRE-EPI data. The fact that the observed effect is larger in the brainstem might be due to inhomogeneities in the B_1 -excitation field. This, however, was not further investigated.

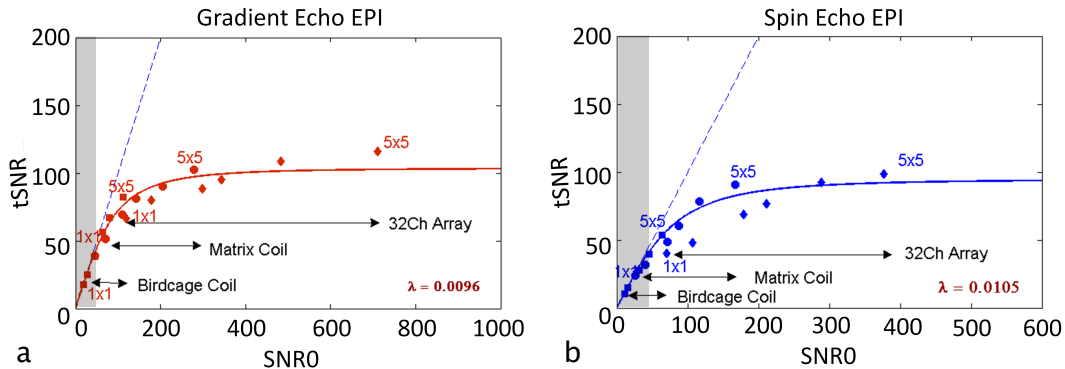


Figure 4.3: Results of fitted Krüger SNR model to GRE-EPI (a) and SE-EPI (b) data. It is seen that the wide range of SNR₀ obtained by varying resolution and receiver hardware in this data is sufficient to reach an asymptote in tSNR. For reference the SNR₀ range obtained in our experiment is marked by the grey areas. (Figure reproduced from [94] with permission from the author)

Most of the issues described above could be addressed in a follow-up experiment if we wish to obtain reliable estimates of λ . For example, the range of SNR₀ could be increased by changing the voxel size or acquiring data at varying field strengths, or the TR could be increased to allow more time for recovery of the longitudinal magnetisation. In order to mitigate problems with the refocusing profile we could design a more refined refocusing pulse with a Shinnar-Le Roux (SLR) algorithm to replace the standard 180° SINC pulse.

However, shortly after these experiments were conducted, Triantafyllou *et al.* presented a similar study in which GRE-EPI was compared to SE-EPI at various field strengths [94]. This study acquired axial data covering the cortical grey matter. Considerably higher SNR₀ values were achieved by varying the receiver coils and resolution. Maximum SNR₀ values of 1500 were achieved for GRE-EPI and 800 for SE-EPI, which was sufficient to reach the asymptotic limit and therefore obtain a more reliable estimate of λ (Fig. 4.3). The following λ values were found: $\lambda=0.0096$ for GRE-EPI, and $\lambda=0.0105$ for SE-EPI. The estimated λ values from our data are considerably higher than the results obtained by Triantafyllou *et al.*.

The likely reason for this discrepancy is that the original model does not fit

multi-channel data accurately at low SNR_0 , which can be attributed to a bias in the SNR_0 calculation caused by the noise covariance between the receivers in multi-channel coils [87, 95]. When the noise covariance between the channels is non-zero, the conventional method that was used to estimate SNR_0 will introduce a positive bias in SNR_0 . The overestimation is clearly depicted in Fig. 4.3, which demonstrates that for multi-channel data the low SNR_0 points all deviate from the fitted model. Thus, if the model is fitted to the low SNR_0 data only, an overestimation of SNR_0 results in an overestimation of λ . An elegant method, which has been adopted recently [87, 95], is to acquire a separate noise reference scan (i.e., a measurement with no RF excitation). The reference scan is then used to correct the SNR_0 value - taking into account the noise covariance, effective receiver bandwidth, and noise rectification in magnitude images [96]. At the time that this experiment was conducted, however, a separate noise measurement was not performed.

Nevertheless, the similarity between the λ values for GRE-EPI and SE-EPI as shown by Triantafyllou *et al.* (Fig. 4.3) suggests that the reduced sensitivity to physiological noise for bSSFP [13] is not driven by the spin echo character of the sequence.

4.3 Physiological Noise in 2D balanced SSFP and SPGR

Based on the evidence above that GRE and SE signals exhibit similar sensitivity to physiological noise, we now compare GRE-EPI directly with balanced SSFP. However, creating a range of SNR_0 values for both sequences without inherently changing a range of other properties is non-trivial when comparing sequences with such distinctly different contrast mechanisms. Balanced SSFP does not have a simple relationship between TE or flip angle and SNR, and the presence of

bands generally complicate SNR calculations. Further, the Krüger approach does not elucidate the source of physiological noise.

Therefore in this set of experiments we will compare the sequences by means of power spectrum analyses. These analyses are performed on rapidly acquired data in order to critically sample the cardiac and respiratory cycles. We show that the frequency spectra can be used to extract SNR_0 and $t\text{SNR}$, which we will use to determine the ratio of physiological noise to thermal noise, as well as decompose signal fluctuations into different sources.

4.3.1 Time-Series Power Spectrum Analysis

Time-series power spectrum analysis is a powerful signal processing technique, which is commonly used in a variety of applications, such as the analysis of audio waves, EEG, and astrophysics. A power spectrum shows how much information is contained at each of the frequencies up to the Nyquist frequency (i.e., half the sample frequency, $f_{nyq} = f_s/2$). The power spectrum is thus a concise representation of all the periodic modulations that are part of the time-series. This allows one to easily identify and quantify any (quasi)-periodic fluctuations in the data set. For example, for signal that is fluctuating at the cardiac frequency a clear peak would be seen around the cardiac frequency ($f_{card} = 0.8 - 1.3$ Hz). When the signal is not described by a perfect sinusoid the power spectrum will also contain peaks at multiples of this baseline frequency, called the higher harmonics. For a signal at frequency f , which can be written as $S(f) = a * \sin(2\pi f) + b * \cos(2\pi f)$, the power of the signal represents the square of the amplitude $P(f) = |S(f)|^2$, which is related to the standard deviation of the signal by: $\sigma = \sqrt{P(f)}/2$.

4.3.2 Comparing 2D bSSFP and SPGR to GRE-EPI

In a healthy volunteer the heart rate during rest ranges between 0.8–1.3 Hz. The respiratory rate is considerably lower (approximately 0.2–0.35 Hz during rest). In order to temporally resolve these fluctuations a sampling rate of at least 2.5 Hz is required, or 5 Hz when we wish to resolve the second order harmonics as well. This means that the data need to be acquired with a volume acquisition time of 200 ms or less.

Data acquisition

GRE-EPI, RF-spoiled SPGR, and passband bSSFP data were acquired from five healthy volunteers on a 3T Siemens TIM Trio system (Siemens Medical Solutions, Erlangen, Germany). The GRE-EPI acquisition is representative of conventional single-shot fMRI acquisitions, while SPGR uses the same acquisition as the bSSFP sequence. This will enable us to disentangle effects due to sequence contrast (SPGR vs. bSSFP) and effects due to readout strategy (GRE vs. SPGR).

Three time-series acquisitions consisting of a single coronal slice (matrix = 128×128 , resolution = $2 \times 2 \times 2.5$ mm) through the brainstem and motor cortex were acquired every 152 ms. A total of 2000 volumes were acquired, which resulted in a scan duration of five minutes per time-series. GRE-EPI parameters were: $\alpha = 28^\circ$, TR/TE = 152/30 ms, 1860 Hz/pix. BSSFP and SPGR parameters were: $\alpha = 30^\circ$, TR/TE = 9/3.5 ms, 1860 Hz/pix, 8 lines per TR. For the bSSFP sequence the RF increment was set to place as much of the slice as possible in the bSSFP pass band. The SPGR acquisition was identical to SSFP, but included RF spoiling (117° quadratic phase increment) and gradient spoiling (2 cycles per voxel). During each acquisition heart and respiration rate were monitored using pneumatic bellows and a pulse oximeter to determine the frequency bands of the physiological cycles during each scan.

Analysis

A linear trend was removed before Fourier transforming each voxel time course to generate its power spectra. Fig. 4.4 shows a cartoon representation of a power spectrum alongside a real example of a voxel time series. The integral over the entire power spectrum reflects the total noise present in the data, and consists of a white noise component (the spectral baseline) and additive physiological noise (the peaks at the respiratory and cardiac frequency and their harmonics). To enable direct comparison of the noise components across the three sequences, we normalize the noise variance by the square of the temporal mean, μ :

$$\frac{\sigma_{tot}^2}{\mu^2} = \frac{\sigma_0^2 + \sigma_r^2 + \sigma_c^2}{\mu^2} \quad (4.5)$$

here σ_x (with $x = 0, r, \text{ or } c$) represents the white, respiratory, or cardiac noise, respectively, and σ_{tot} the total time-series noise. The individual SNR components, including SNR_0 can then be obtained by $\text{SNR}_x = \mu / \sigma_x$.

In order to calculate the thermal noise, the spectral baseline was estimated and integrated over the entire power spectrum. This baseline was then subtracted from the spectrum such that the remaining spectrum only contained the (additive) physiological noise components. The respiratory and cardiac noise are calculated by integration of the power spectrum over the respiratory and cardiac frequency bands, which were determined from the physiological recordings (pulse-ox and bellows). The frequency bands were estimated by fitting a Gaussian distribution to the frequency spectra of the physiological recordings. The fits were performed in Matlab using `nlinfit` and initialised based on the average values 0.8–1.3 Hz and 0.2–0.35 Hz for the cardiac and respiratory rate, respectively. The boundaries were set to 2σ on either side of the centre frequency of the estimated Gaussians.

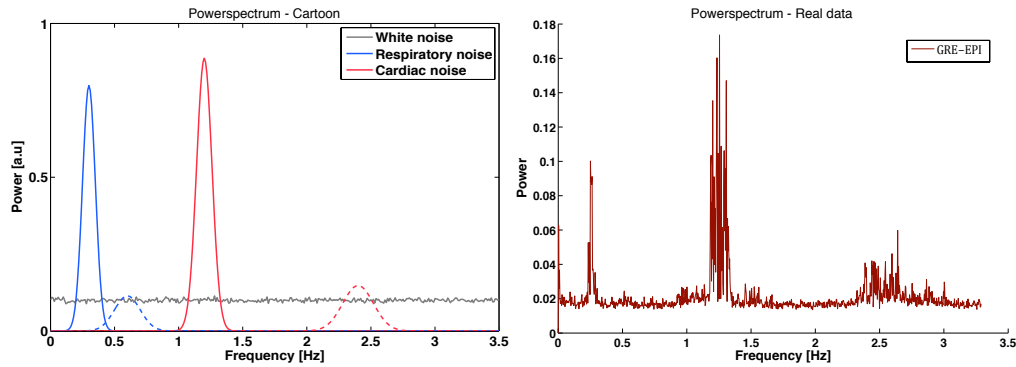


Figure 4.4: (a) A cartoon representation illustrating the white various additive noise components. Thermal noise is shown in grey. The respiratory and cardiac noise are shown in blue and red, respectively. Higher harmonics are shown as dashed lines. (b) shows an example power spectrum from the acquired data.

With these values in hand we can perform an estimate of λ by inserting the SNR_0 and $t\text{SNR}$ values into equation 4.2, although it should be mentioned that we are effectively only using a single point on the SNR_0 - $t\text{SNR}$ curve, so caution should be taken when interpreting those values.

For the ROI analysis, four ROIs were manually drawn for each of the subjects. The ROIs included: the medulla, the pons, the hippocampus, and grey matter. To prevent a bias between the sequences introduced by the fact that the ROIs were drawn manually, the data sets were first registered to each other after which the ROI was drawn on each of the data sets separately. The intersection between the three manually drawn ROIs was selected and transformed back into the native space of each of the three acquisitions.

Results and Discussion

Maps representing the contribution of the individual noise components in a representative subject are shown in Fig. 4.5(a), including the total noise variance ($t\text{SNR}^{-1}$) map. The left column shows the temporal mean of each of the acquisitions. The σ_0/μ reflects the normalised thermal noise and thus $1/\text{SNR}_0$. Due to the reduced read-out efficiency (i.e. duty cycle) of GRE-EPI and the spoiling of

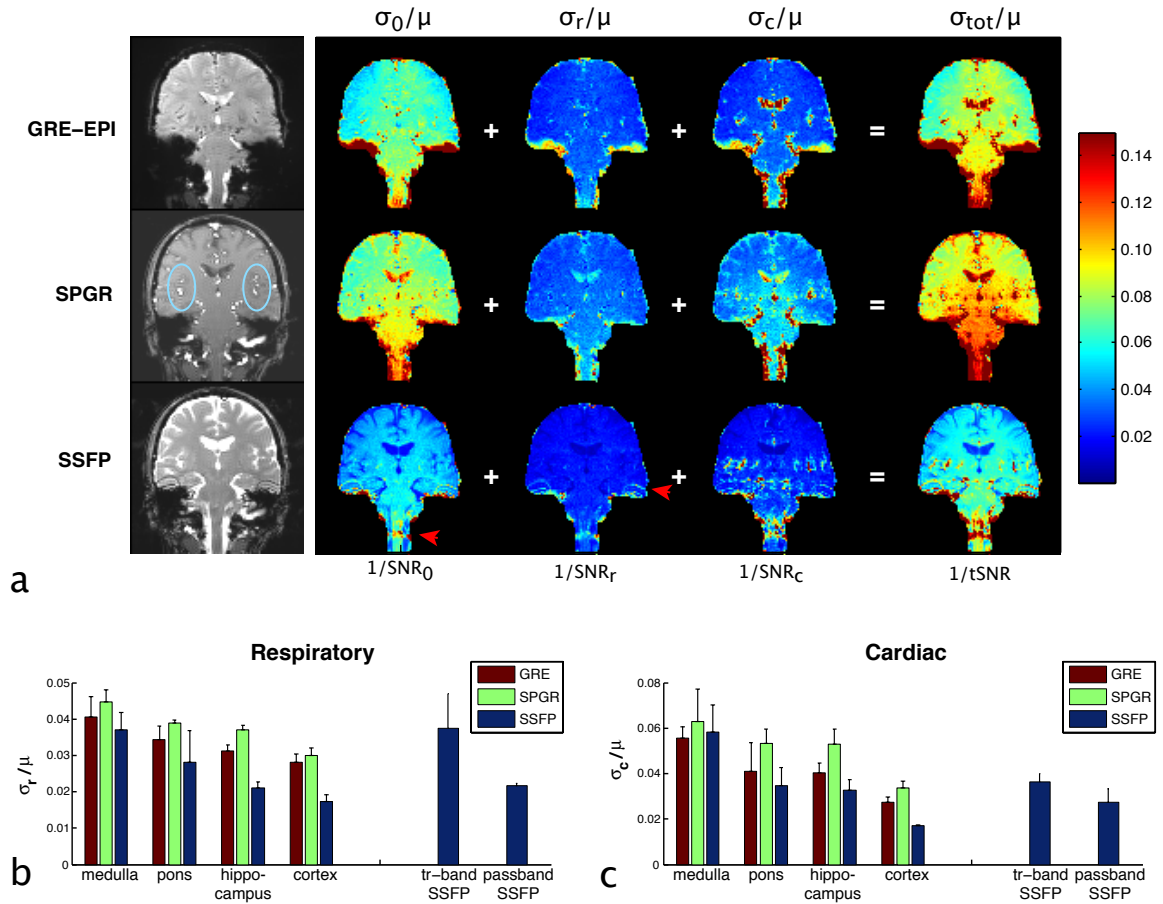


Figure 4.5: (a) Maps showing the decomposed noise measures, extracted from the power spectra using Eq. 4.5, for a representative subject. All sequences demonstrated elevated levels of cardiac noise in the CSF space of the Sylvian fissure (denoted by the blue circles in the SPGR image). Also shown here is the increased sensitivity to respiratory noise in the transition bands in bSSFP (red arrow). The group level results ($n=5$) are presented in (b) and (c). Overall, bSSFP demonstrated the lowest sensitivity to respiratory and cardiac noise.

transverse magnetisation in SPGR, these sequences have a lower SNR_0 compared to SSFP, which is reflected by the elevated σ_0/μ (i.e., the level of white noise in the normalised power-spectrum). In this study shimming was set up to place as much of the entire FOV in the passband of the bSSFP profile. Due to the increased resonance offsets around the auditory canals and the lower part of the medulla several bands were present (the transition band portion of the bSSFP signal profile). When data are acquired at $\alpha = 30^\circ$ the signal in the transition bands is considerably lower than in the passband region (see Fig. 2.6(b)). This is reflected

by the increased σ_0/μ (red arrows).

The regions of the transition bands also show increased fluctuations at the respiratory frequency (σ_r/μ), which is caused by small B_0 field fluctuations caused by bulk susceptibility changes from varying lung volume during the respiratory cycle [58]. Unlike the flat passband portion of the bSSFP signal profile, signal near the transition band is highly dependent on the local resonance frequency. Small changes in the B_0 field will therefore cause the signal to fluctuate with the respiratory cycle.

The σ_c/μ maps for all sequences show large cardiac fluctuations in regions of CSF surrounding the brainstem, the ventricles, and the CSF spaces of the Sylvian fissure (light-blue circles in the SPGR reference image). This effect can be attributed to CSF pulsatility, which is known to be correlated to the cardiac cycle [97].

The ROI analyses, which show the average results over five subjects, shows increased respiratory noise for the more inferior regions of the brain (Fig. 4.5(b)). This trend is visible for all sequences. The smallest σ_r/μ is observed for bSSFP, which is consistent across all ROIs. Fig. 4.5(c) shows that σ_c , like σ_r , is also increased in the inferior regions of the brain. Again, the noise variance was smallest for bSSFP, across all ROIs, except for the medulla where GRE-EPI slightly outperforms bSSFP. Inspection of the individual tSNR maps revealed that for two subjects bands were present in the medulla as a result of the large shim volume used in this study. An additional ROI that covers the transition bands of the bSSFP images illustrates that transition band bSSFP (tr-band bSSFP) is much more sensitive to physiological noise, in particularly to respiratory effects

From the extracted SNR_0 and tSNR values the λ values were calculated for each acquisition. In the pons we calculated the λ values to be 0.059, 0.069, and 0.048 for GRE-EPI, SPGR, and bSSFP, respectively. The λ values for the cortex

were: 0.045, 0.049, and 0.029. Although these values are still higher than the values reported by others [13, 94, 84], they are comparable to the values found in the previous experiments (Section 4.3), which was conducted with the same slice orientation and comparable SNR_0 .

Despite having the same contrast as GRE-EPI, SPGR shows increased noise variability in both ROIs. This indicates that the multi-shot character of the segmented EPI readout introduces a source of instability. Previous work has reported increased sensitivity to physiological noise in 2D segmented readouts [98]. The increase was attributed to the fact that, in multi-shot sequences, data are sampled at different phases in the cardiac or respiratory cycle, but combined into a single image. The results further indicate that the reduced sensitivity to physiological noise in bSSFP may be mitigating these effects.

The Effect of Segmentation

In order to test whether segmentation has an effect on the physiological noise sensitivity of bSSFP, a second experiment was conducted. For this experiment three subjects were scanned, acquiring four Single-slice bSSFP data sets with varying amounts of segmentation. The matrix size for these scan was set to 96×96 and a total number of 1000 volumes was acquired. Other scan parameters were kept identical to the parameters used in the previous experiments. The segmentation factor ranged from 6 lines per TR (16 shot) to 16 lines per TR (6 shot). In order to investigate the effect both passband and transition band bSSFP the z-shim was deliberately offset to create multiple bands in the SI direction. The data were processed and analysed as before and results are shown in Fig. 4.6.

Fig. 4.6(a) shows the σ_r/μ and σ_c/μ maps along with the raw images for each of the acquisitions. The raw images show that the distortion becomes less with increasing number of shots, as the readout per excitation becomes shorter. Com-

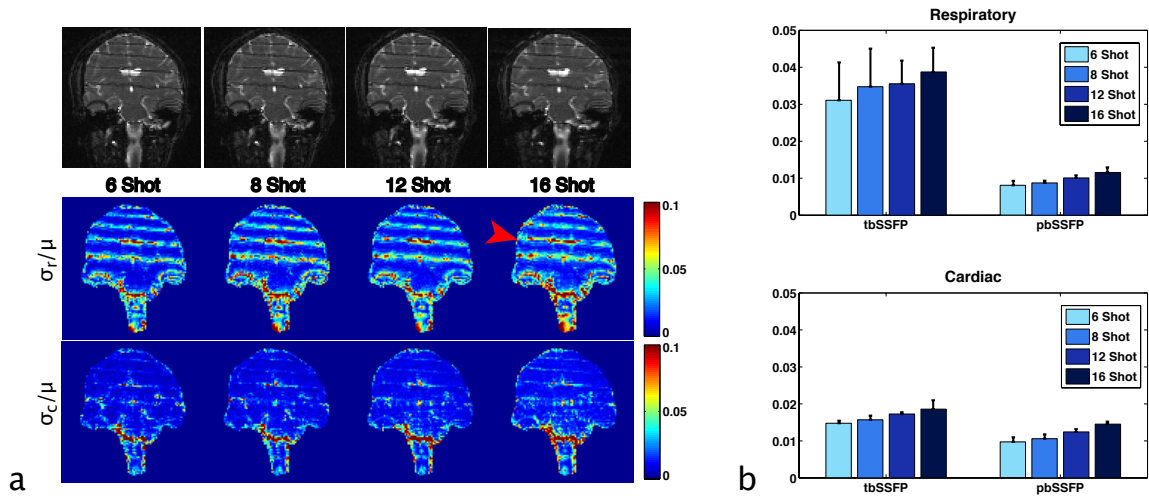


Figure 4.6: Figure showing the raw images (a), the derived noise maps (b), and ROI results (c) for data acquired with different amounts of segmentation in the EPI readout.

paring the relative contributions of σ_r and σ_c , we see that for passband bSSFP the noise introduced by respiratory and cardiac are comparable. For transition-band bSSFP on the other hand σ_r/μ is considerably larger than σ_c/μ (also shown by the ROI results displayed in Fig. 4.6(b)). The amount of physiological noise increases with the amount of segmentation that is applied suggesting a tradeoff between distortion and physiological noise. In both transition-band and passband bSSFP σ_r and σ_c increase with the number of shots, although the differences are fairly small.

4.4 3D balanced SSFP; FMRI of the Hippocampus

So far most studies reporting the advances in bSSFP FMRI development have been using simple functional tasks (e.g., visual and motor) to demonstrate activation [10, 13, 11, 56]. These areas produce large and robust signal changes. It is important, however, to test the capability of bSSFP to FMRI when using more elaborate tasks that produce more localised areas of activation, as these are the type of tasks ultimately used in brainstem FMRI.




Figure 4.7: The pulse sequence diagram (a) and k -space trajectory of the stack-of-segmented EPI readout (b) as used in this study. After slab selective excitation, segmented EPI planes are acquired in the k_x - k_y -dimensions; with the 3D phase encode direction in k_z . In this particular example 8 lines are acquired per segment. At the end of the readout all gradients are rewound.

As an initial assessment of the functional contrast of bSSFP under more elaborate conditions we conducted a memory task experiment that activates the hippocampus. Although in the scope of this thesis the hippocampus served as an initial test bed for our sequence before acquiring brainstem FMRI data, it has potential as a real application for bSSFP FMRI.

This project was conducted in collaboration with the Medical Imaging group of the Norwegian University of Science and Technology in Trondheim. This group seeks to develop a protocol for high resolution, low distortion, hippocampal FMRI in order to investigate the functional differentiation between the various sub-areas of the hippocampus. The study consisted of an assessment of the functional sensitivity of bSSFP using a memory-encoding paradigm. Whole brain (single time-point) images were also assessed on distortion and compared to GRE-EPI.

4.4.1 3D Stack-of-segmented EPI

For this study the stack-of-segmented EPI readout is used as briefly presented in Chapter 2. Here, the readout is presented more in detail and in the context of a bSSFP acquisition. The pulse sequence diagram of a bSSFP acquisition that uses the 3D stack-of-segmented-EPI readout is shown in Fig. 4.7(a). The k -space trajectory is shown in Fig. 4.7(b). The illustration shows the segmented EPI planes

stacked along k_z . In this diagram 8 EPI lines are acquired per segment. After slab selection the dephaser along G_y determines which segment in the EPI-plane is collected. The G_z dephaser determines the position along k_z from which data are collected. At the end of each TR all imaging gradients are rewound (this effectively brings the trajectory back to the centre of k -space, but this is for illustration purposes not shown in (b)). In its most straightforward form, the readout order takes place in a linear fashion. Starting with the most negative k_z -location and ending with the most positive k_z . For each EPI-plane all the segments are acquired before moving on to the next. This readout has been successfully applied for high resolution bSSFP FMRI of the visual cortex at 1.5 Tesla [10].

4.4.2 Methods

Functional bSSFP data were acquired in 3 healthy volunteers using a 3T Siemens TIM Trio system (Siemens Medical Solutions, Erlangen, Germany). Oblique coronal 3D bSSFP images were acquired perpendicular to the hippocampal long axis with the frequency-encoding in the superior-inferior (SI) direction and scan parameters: TR/TE = 8.0ms/3.8ms, $\alpha = 30^\circ$, matrix $96 \times 96 \times 32$, resolution $2 \times 2 \times 2$ mm, 1860 Hz/pixel, 6 lines per TR, and volume TR, $T_{\text{vol}} = 4.1$ s. The shim was carefully localised to the hippocampal region to place the both hippocampi in the passband region of the bSSFP signal profile. The functional paradigm involved a memory task, in which the subject was presented with 9 alternating blocks (24s) of novel (N) and familiar (F) scene pictures. The subject was instructed to memorize all images and to encode each image by indicating whether a building was present in the picture by using a button-box. The task was presented across two runs, which each lasted for 7.5 minutes (112 volumes). The two scan sessions were individually analysed using a standard GLM approach and then combined using fixed effects averaging of the t-statistics to generate the activation maps on a

single subject level.

To demonstrate the reduced distortion in the bSSFP images, whole brain (non-functional) bSSFP and GRE-EPI images were acquired in a single subject as well as a T_1 -weighted structural scan. The scan parameters were carefully chosen such to be representative of a functional experiment. The GRE-EPI image was acquired using the following parameters: single-shot EPI, resolution $2 \times 2 \times 2.5$ mm, matrix 96×96 , 60 slices, 1860 Hz/pix, $\alpha = 90^\circ$, TR = 4990 ms, TE = 30 ms. To reduce the distortion as much as possible the data were subsampled along the phase encode direction by a factor of 2 using GRAPPA [44]. The bSSFP imaging parameters were: TR/TE = 8.0ms/3.8ms, $\alpha = 30^\circ$, matrix $96 \times 96 \times 80$, resolution $2 \times 2 \times 2$ mm, 1860 Hz/pixel, 6 lines per TR. For the structural, an inversion-prepared 3D MPRAGE sequence was used with TE/TR/TI = 4.6ms/3300ms/1100ms, matrix $256 \times 176 \times 224$ and 1mm^3 isotropic resolution.

4.4.3 Results

In this section, we first demonstrate the improved (static) image quality in bSSFP compared to single-shot GRE-EPI. Next, we will present the preliminary functional results that were obtained in the three subjects that were scanned. The results will be discussed at individual subject level.

Image Quality

Example slices of the raw images are shown by Fig. 4.8. For both acquisitions the data were acquired coronally with the readout direction in SI in order to avoid signal aliasing from outside the field-of-view. The EPI blip direction was placed along the Right-Left (RL) direction, so the largest distortion will manifest in the RL direction of the image. When comparing bSSFP to GRE-EPI it is seen that the distortions are considerably larger in the GRE-EPI data, due to the single-

shot nature of the latter acquisition. The susceptibility offsets of the nasal sinuses cause strong local field gradients in the inferior part of the frontal lobe and the anterior part of the brainstem. This results in signal dropout in these areas in GRE-EPI due to the relatively long TE ($\sim 30\text{ms}$) that is typically used in GRE experiments. In bSSFP no such signal dropout is observed, although bands of low signal arise due to the dip in the signal profile at particular frequency offsets ($\pm \frac{1}{2T}$ Hz, see Fig. 2.6(b)). However, in bSSFP the location of the bands can be changed by changing the RF phase increment. A method which combines the data from two acquisitions with varying RF phase increment has been used for whole brain bSSFP FMRI [99, 11]. By collecting one image with a 0° phase increment and one with a 180° phase increment, all the off-resonance frequencies will be acquired in the passband in at least one of the acquisitions. The images can then be combined with a maximum intensity projection (MIP) or sum-of-squares to produce an image without bands [100]. The disadvantage of this method, however, is the time penalty involved for a second acquisition. For studies that target a specific structure a better approach is to localise the shim to the given structure to maximise the field homogeneity in that region and accept banding artifacts outside the region of interest.

Fig. 4.9 shows the registration of the acquired GRE-EPI and bSSFP images (red outline) to the subject's high resolution anatomical scan. Due to the distortion of the GRE-EPI the registration is considerably worse for GRE-EPI than for bSSFP. Although both images show good alignment in the superior areas of the brain, the GRE-EPI data clearly show poor alignment around the medial (blue arrows) and inferior (green arrows) temporal lobes. The frontal areas of the brain are also affected. In the edge-maps this is most clearly visible by the misalignment around the ventricles. Conversely, the bSSFP image shows excellent alignment throughout the brain. The exquisite anatomical detail and reduced distortion of




Figure 4.8: Coronal, sagittal, and transverse planes displaying the GRE-EPI (**a**) and balanced SSFP (**b**) data as acquired. The distortion is shown by the skew of the image in the Right-Left direction (which was the EPI-blip direction during acquisition). From the coronal view it is clearly visible that the distortion is much less in the bSSFP images compared to GRE-EPI. In areas of strong field inhomogeneities GRE-EPI shows severe signal dropout (white arrows). In bSSFP banding occurs in these regions.

the bSSFP image will give more confidence in identifying, for example, specific subfields of the hippocampus and may improve the statistical analysis of FMRI data at group level.

Functional Results

Fig. 4.10 shows the activation maps (z -statistic) for each subject in three in three perpendicular slices overlaid on the raw bSSFP images (threshold: $z > 2.3$, cluster $p > 0.05$). All subjects showed hippocampal and/or parahippocampal activation, although the intensity and extent varied. Subject C showed the strongest activation in both hippocampal and parahippocampal regions. The activation was less strong in subjects A and B, with subject B showing only minor activation in the parahippocampal gyrus. Also shown in this figure are the dark bands typical for bSSFP images. Subject C and B further showed bilateral cerebellar activation for




Figure 4.9: Functional to structural co-registration. High resolution T_1 anatomical data are overlaid with (a) the edge-maps (red outline) of the registered GRE-EPI data, and (b) balanced SSFP data. The images were acquired coronally with the EPI blip direction in RL. Due to the distortion present in the GRE-EPI data the registration in the hippocampus is inaccurate (blue arrow) and other areas in the frontal regions of the brain (green arrows). The bSSFP data on the other hand show excellent registration throughout the brain. Note that the apparent misregistration at the front of the brain is not due to distortion, but due to a reduced field of view in the bSSFP.

the $N > F$ contrast (blue). This may indicate the involvement of the cerebellum, which is consistent with previous reports in the literature [101].

Artifactual activation (false positives) was seen in subject A and C, mainly in the regions of the transition bands, but also near CSF boundaries. This may be due to temporal fluctuations of the CSF signal caused by flow effects or motion that correlated with the stimulus.

Recall that the GLM is a least-squares fit to the data, described by: $Y = X\beta + \epsilon$, where Y is the data, X the regressor, β the parameter estimate, and ϵ the residual error. One summary measure of the GLM fit that is useful to inspect is σ_{res}^2 , the temporal standard deviation of the residuals. Inspection of the σ_{res}^2 images showed that a significant amount of the activation that survived the threshold also had high σ_{res}^2 values. This can be an indication that the β estimate is driven by a

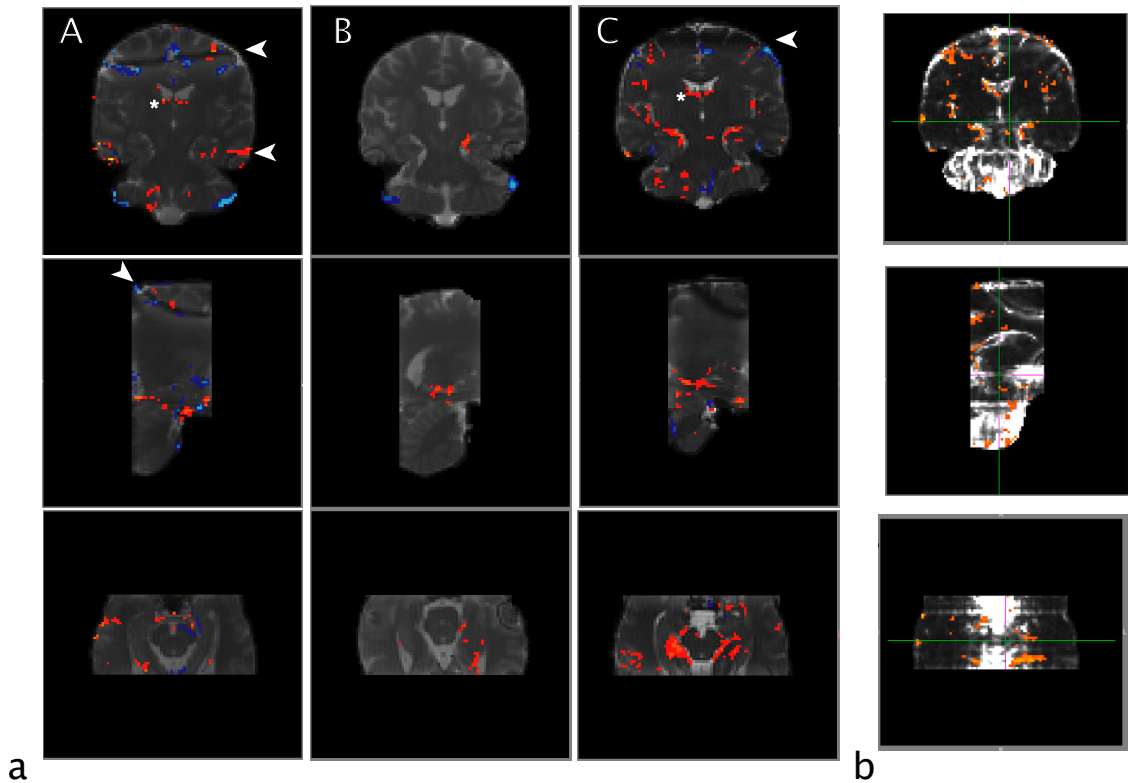


Figure 4.10: (a) z-statistic activation maps per subject overlaid on the raw bSSFP images (threshold $z > 2.3$, cluster significance $p > 0.05$. Red-yellow activation denotes the $F > N$ contrast, whereas blue represents $N > F$. In two subjects false positives were found in proximity of the transition bands (arrows), but also near the lateral ventricles (asterices). (b) σ_{res}^2 map of subject 3 denoting the error between the model and the data. Areas of high error appear as a ghosting pattern in the RL direction, whereas the appearance is blurred in the AP direction.

few large signal fluctuations that happen to coincide with the stimulus paradigm. Visual inspection of the time courses outside the hippocampus confirmed this. The hippocampal activation, however, is also in close proximity of voxels with high error. Although visual inspection of the time courses assured that most of the activation was real, this is an issue that one should be aware of. Furthermore the σ_{res}^2 image shows large errors in the inferior regions of the brain. The appearance of the errors is peculiar. In the RL direction (i.e., the EPI-blip direction of the readout) a ghosting pattern is visible, whereas the errors appear to be smeared in the Anterior-Posterior (AP) direction (i.e., the direction in which the EPI planes are stacked). This indicates that the nature of the readout determines the appear-

ance of image artifacts (ghosting and blurring in this case) that are fluctuating over time. Since the large fluctuations are mostly in the inferior regions this may be a potential problem for brainstem fMRI.

4.5 Temporal Stability in the Brainstem using 3D Readouts

Because the hippocampal fMRI data showed large σ_{res}^2 values in the pons and cerebellum it was decided to conduct resting experiments (i.e., without stimulus paradigm) with the FOV covering the entire brainstem and assess the temporal stability in 3D bSSFP. Although the physiological noise assessment of 2D data suggested beneficial physiological noise characteristics in the brainstem, it is important to assess whether this remains when we acquire the data with a full 3D readout.

4.5.1 Methods

Five-minute resting state scans were conducted in a single subject. 2D GRE-EPI, 2D bSSFP and 3D bSSFP data were collected on a 3T Siemens TIM Trio system (Siemens Medical Solutions, Erlangen, Germany). Scan parameters are similar to the previous studies. For the 2D data, single coronal slices were acquired every 152 ms with a resolution of $2 \times 2 \times 2$ mm. For the 3D data identical scan parameters to the hippocampus study were used, but with the slice orientation fully coronal to include the entire brainstem (75 volumes).

4.5.2 Results and Discussion

Fig. 4.11 shows the improved temporal stability of 2D bSSFP compared to 2D GRE-EPI, which is one of the motivations for using bSSFP FMRI in the subcortical areas that are more prone to physiological noise. In the cortical regions the temporal stability is further improved for the 3D variant, which can be attributed to the increased signal available in 3D readouts. For regions around the ventricles and the brainstem, on the other hand, the tSNR is considerably reduced. In the brainstem the areas of instability are highly localised. The increase may be related to the fact that the brainstem is surrounded by the major cerebral arteries and pulsatile CSF. In 2D single-slice imaging most of these structures are located outside the FOV and therefore do not contribute to the signal. The presence of those structures in the 3D FOV may introduce additional fluctuations. In multi-shot imaging each segment is collected at a different time and then combined to a single image, which increases the sensitivity to physiological noise. Fig. 4.6 showed a small increase in physiological noise when the number of shots increased. However, the time scale at which a single volume is collected is significantly different for 3D readouts (i.e., a few seconds vs. 150 ms). Physiological signal fluctuations during the acquisition of a single volume may therefore be much larger and thus have a much stronger impact on the image quality in 3D readouts compared to 2D. For some voxels in the brainstem the instabilities were as high as 30% of the mean signal. The current sequence is therefore not a viable option for brainstem FMRI and further development is required before attempting brainstem FMRI using bSSFP.

4.6 Discussion

In this chapter two candidate sequences 3D FMRI sequences, 3D bSSFP and SPGR, were assessed based on a range of imaging properties important for brainstem




Figure 4.11: $1/tSNR$ maps of 2D GRE-EPI, 2D bSSFP and 3D bSSFP. 2D bSSFP shows superior temporal stability to GRE-EPI. When the data are acquired with a 3D readout, however, large signal instabilities arise. Although the $tSNR$ is larger in the superior regions of the brain, the 3D acquisition exhibits large signal fluctuations in the brainstem, which are not present in the 2D data. Similar results were obtained in SPGR data, although not shown here.

FMRI.

First, we compared the physiological noise sensitivity of SE-EPI to GRE-EPI in order to investigate whether the reduced physiological noise of bSSFP in cortical regions could be attributed to the spin echo contrast of the sequence. For this set of experiments the Krüger model was used to estimate the contribution of physiological noise by fitting a temporal noise model to a range of SNR_0 values. By assuming that physiological noise scales with signal intensity unlike thermal noise, this model estimates the theoretical limit of $tSNR$ at $SNR_0 = \infty$ (given by $1/\lambda$). However, with the current hardware setup and scan parameters we did not achieve a large enough range of SNR_0 in our data to demonstrate asymptotic behaviour (i.e., all the data points resided in the linear part of the curve), which resulted in a poor fit of the model to the data.

When using the Krüger model the fairest comparison is to create a range of

SNR₀ by altering the flip angle (e.g., increasing the resolution could confound the results when the spatial correlation of the physiological noise is different between the sequences). The vastly different signal dependency on flip angle in bSSFP, however, makes this approach non-trivial.

By means of power spectrum analysis we were able to directly compare the physiological noise sensitivity of bSSFP and SPGR to GRE-EPI. The method allows extraction the same noise measures (i.e., σ_0 and σ_p), but with the added benefit that the physiological noise component, σ_p , could be decomposed into cardiac noise and respiratory noise. An estimate of λ was also obtained, by plugging the determined SNR₀ and tSNR values into Eq. 4.4. It should be noted, however, that the accuracy of λ is even more dependent on the SNR₀ as only a single point on the curve is used for the estimation (ideally one would only accept an SNR₀ for which the asymptote of the SNR₀-tSNR curve has been reached as this gives the most accurate estimate).

The results showed that bSSFP has reduced physiological noise in both the cortical regions and the brainstem. SPGR however performed worse than GRE-EPI. The observed λ values are considerably larger than the values previously reported by Miller *et al.* [13] and Triantafyllou *et al.* [94], but are comparable to the values we found in the first experiments presented in this chapter. We believe the discrepancy is due to the fact that both studies sample the tSNR-SNR₀ curve in the low SNR₀ region due to the relatively high resolution at which the studies were conducted.

Closer inspection of Fig. 4.3 shows that for 12- and 32-channel acquisitions the data are not well described by the Krüger model in this region of the curve, which may be the reason that our λ estimates are positively biased. This observation is supported by a recent paper by Triantafyllou *et al.* [87], in which an improved temporal noise model was presented for multi-channel data.

FMRI experiments of the hippocampus were conducted and the levels of distortion around the hippocampal regions were assessed in bSSFP as well as GRE-EPI. It was demonstrated that, even at a moderate resolution, the distortions were considerably less in the 3D segmented bSSFP readout compared to single-shot GRE-EPI. These properties are important, as the brainstem is located in regions of high susceptibility gradients. The distortions caused by the field inhomogeneities are the limiting factor in the achievable resolution in sequences that require a single-shot readout. It was further shown that bSSFP is able to pick up signal changes from more elaborate tasks, but the extent of functional activation varied between subjects. Spurious activation was present in two subjects. Inspection of the σ_{res}^2 images (the summary image denoting the variance of the residuals) demonstrated increased signal fluctuations unrelated to the task in areas of CSF when the data are acquired with a 3D acquisition. Resting state scans demonstrated large temporal instabilities in the brainstem. The instabilities are highly localised suggesting that the signal fluctuations are driven by structured artifacts, which fluctuate from one volume to the next. The source of these effects and possible methods for mitigating them will be the focus of the following research chapters.

4.7 Conclusion

In this chapter the image properties of two candidate sequences, bSSFP and SPGR, were compared to conventional GRE-EPI. It was demonstrated that 3D bSSFP using a 3D stack-of-segmented EPI readout has superior image quality in terms of signal dropout, distortion, and SNR_0 . Furthermore it was found that bSSFP is able to achieve functional contrast in tasks other than simple motor or visual tasks that have been reported previously. False positive activations, however,

were found in regions of with increased signal fluctuations. Although bSSFP demonstrated reduced physiological noise sensitivity when a 2D readout was used, large signal instabilities were observed in data acquired with a 3D readout. The levels of the signal fluctuations are large enough that fMRI studies with the current implementation of the sequence are not possible. Therefore, in order for 3D techniques to be viable for brainstem fMRI methods are needed that correct or prevent these signal fluctuations. In the next chapter a thorough characterization of the physiological instabilities is provided and various correction methods are explored.

Chapter 5

Physiological Noise: Characterization and Correction Methods

5.1 Introduction

As shown in Chapter 4 large temporal signal instabilities are observed in the inferior regions of the brain when the data are acquired with a 3D readout. Figure 4.10 showed that the way the instabilities are manifested is dependent on the readout order in which k -space is collected. We therefore hypothesise that the instabilities are associated with the multi-shot character of the 3D acquisition. Any errors between the segments in k -space result in instabilities in image space. Because the high levels of variability appear to arise mainly from regions of CSF and blood vessels we believe that they are caused by physiological fluctuations (i.e., cardiac pulsatility and respiratory effects).

This chapter describes experiments and simulations that were conducted in order to identify the source of the physiological instabilities and their effects in 3D acquisition schemes. In the second part of the chapter we investigate the potential of various prospective (e.g., real-time) correction methods. Because prospective

correction methods often involve considerable pulse sequence implementation, the proposed methods are first assessed by simulations. The simulation results provided a valuable insight into the formation of the artifacts and served as a good indicator for which correction method to pursue.

5.1.1 Background

It is well known that fMRI time-series data are compromised by physiological fluctuations related to the cardiac and respiratory cycle [102, 76, 69, 103], with particular impact for multi-shot acquisitions [98]. Apart from the pulsatile movement of blood and CSF, the brain undergoes gross motion due to intra-cranial pressure changes [68, 97, 104]. These physiological effects lead to changes in the measured MR signal. Depending on the source of the artifacts (e.g. respiratory or cardiac pulsation) and the type of pulse sequence (e.g. anatomical, diffusion or fMRI), various correction methods have been introduced. Navigator techniques have successfully been applied in fMRI to correct global fluctuations from respiration [105, 106], but are impracticable to correct the more localized cardiac fluctuations, as high-resolution navigators would be needed. Similarly, retrospective correction methods that operate in k -space are able to effectively correct for low order spatial fluctuations, such as respiratory effects [67, 65]. Spatially localized cardiac fluctuations on the other hand are easier to correct for in image space [71]. Finally, cardiac gating has been used to reduce cardiogenic noise [63]. Cardiac gating, however, has the disadvantage that it produces a variable acquisition rate, which can be cumbersome or prohibitive for some fMRI stimulus paradigms, but also reduces the efficiency of the sequence. Furthermore, the non-stationary TR associated with some gating techniques introduces temporal signal fluctuations due to variable T_1 recovery, which complicates post-processing of the data [107]. Another possible method to reduce physiological noise is to order the acquisition

in k -space in real-time based on the current phase in the respiratory or cardiac cycle [108, 109, 110]. In order to ensure that each line in k -space is acquired with the desired physiological phase (e.g. point in the cardiac or respiratory cycle), these methods generally alter the acquisition rate and thereby introduce similar issues to conventional gating.

It is clear that each of the correction methods has their advantages and disadvantages and often correct for a specific source of signal fluctuation. In order to determine which type of method is required we first need to characterize the instabilities.

5.2 Regression of Physiological Processes with 2D data

Power spectrum analysis presented in Chapter 4 showed that the signal in 2D rapidly acquired data is fluctuating at both the respiratory and the cardiac frequency. In this experiment we aim to investigate to what extent we can explain the fluctuations in the data using external physiological recording (i.e., pulse plethysmograph data and pneumatic bellows). We use this investigation to gain more insight in the underlying complexity of the signal fluctuations and how well they can be described with simple models. Regression of data acquired with a short volume acquisition time (T_{vol}) will encapsulate the effects before they become distributed in k -space (i.e., the data mimic “instantaneously” acquired data). This should provide a more accurate insight into the relative contributions of the cardiac and breathing cycle to the fluctuations in the signal.

5.2.1 Methods

All *in vivo* experiments presented in this chapter were performed on a 3 Tesla Siemens Trio system (Siemens Healthcare, Erlangen, Germany) using a 12-channel

head receive coil. Simulations, image reconstructions and statistical analyses were performed in Matlab (MathWorks, Natick, MA). As we are aiming to characterize the intrinsic signal fluctuations (rather than under task conditions), all presented data were acquired without an explicit task.

Single-slice bSSFP and SPGR time series were acquired in one subject. In order to temporally resolve physiological fluctuations, the acquisitions were restricted to a single coronal slice covering the brainstem and motor cortex. Both the 2D-bSSFP and 2D-SPGR time series were acquired with the following scan parameters: $\alpha=30^\circ$, TR/TE=12/6 ms, FOV=192×192×2.5 mm, Matrix=96×96, Bandwidth=1860 Hz/pixel, 8 lines per TR, $T_{\text{vol}}=156$ ms, 1500 volumes. For both acquisitions the shim volume was targeted to the brainstem. For the bSSFP acquisition the RF increment was set to place the brainstem in the passband of the bSSFP signal profile. For the SPGR acquisition RF spoiling (quadratic phase increment=117°) and gradient spoiling (2π phase per voxel) were performed. Heart and respiration rate were monitored during data acquisition using a pulse oximeter and pneumatic bellows to correlate the voxel time courses to the physiological waveforms.

The magnitude and phase time-course data were processed separately to generate z -statistic maps representing correlations with the recorded cardiac and respiratory waveforms. Pre-processing was minimized to keep the data as close to its raw form as possible. The phase data for each voxel were unwrapped in the time domain. High-pass temporal filtering (high-pass cut-off=25s) was performed on both the magnitude and phase time series to remove any baseline drift. Visual inspection revealed little to no bulk motion, so motion correction and spatial smoothing were omitted in order to detect respiratory and/or cardiac-related motion. The cardiac and respiratory regressors were represented by low order Fourier series as used in the RETROICOR method [71]. The first and second or-

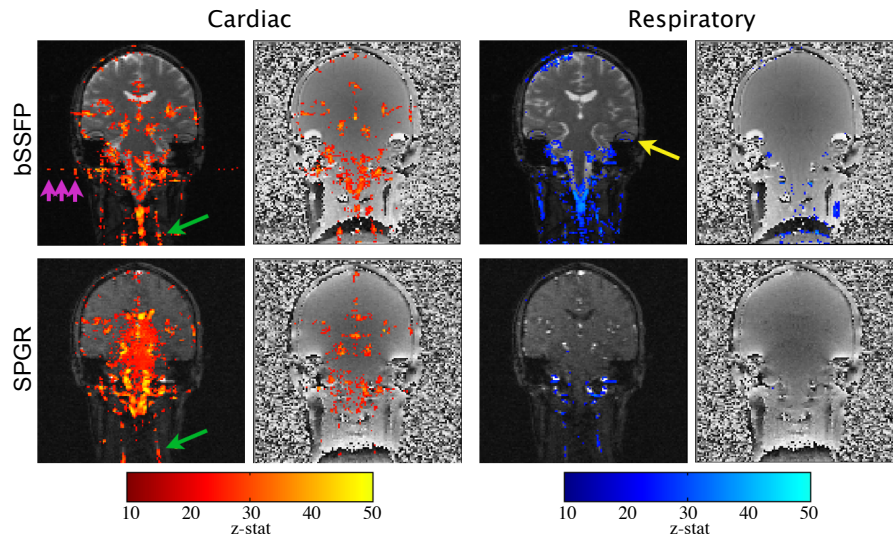


Figure 5.1: Results from the regression analysis of 2D bSSFP and SPGR data with the cardiac and respiratory regressors obtained by physiological monitoring (bellows and pulse oximeter). The z-statistic maps show the spatial distribution of cardiac (red) and respiratory (blue) correlations in the magnitude and phase time-series images. Arrows denote cardiac correlations located in the carotid arteries (green) and what appears to be ghosted signal from blood vessels in the brainstem (purple), and respiratory correlations due to shifting of bSSFP bands (yellow).

der Fourier series were then correlated to the magnitude and phase time courses using GLM regression with FSL FEAT [88].

5.2.2 Results

Fig. 5.1 shows the correlation of the magnitude and phase time-course data with the physiological correlations (displayed as z -stat maps overlaid onto the corresponding raw image). Both bSSFP and SPGR show strong correlations with the cardiac cycle in the signal magnitude and phase.

In the bSSFP data the effects of cardiac pulsation are confined to the large arteries (green arrows) and CSF. In the magnitude data a few voxels that are highly correlated to the cardiac waveform produce fluctuating ghosts (purple arrows) in the left-right direction. This direction is the EPI phase-encode direction. Such ghosts are not visible in the SPGR magnitude data. The cardiac-induced magnitude fluctuations in the SPGR data, however, are less spatially confined and

include the subcortical structures such as the pons, midbrain, and the thalamus. The phase data, on the other hand, show a very similar spatial pattern in both bSSFP and SPGR.

Respiration effects are much less pronounced than the cardiac effects in both data sets. In SPGR respiration effects are almost absent. In bSSFP correlations are visible in regions of the bSSFP transition bands (yellow arrow). The largest correlations are observed in the neck region inferior to the medulla where off-resonance gradients are large and multiple bands are present. It is interesting to note that, when only the first-order terms are included (data not shown), the statistical maps are highly similar to the ones shown here. This indicates that the fluctuations are well described by simple sinusoids.

5.2.3 Discussion

At this point, we are prepared to hypothesize on the origin of the signal instabilities in SPGR and bSSFP and will discuss the commonalities and differences between 2D acquisition and 3D acquisitions.

The regression results demonstrate that physiological instabilities are largely driven by cardiac pulsation and are much more pronounced in the inferior regions of the brain compared to the cortical areas. Respiratory effects were much less intense than cardiac effects and only of significance in the bSSFP data. Visual inspection indicated that the respiratory fluctuations in the bSSFP that are localised to regions of bSSFP bands are caused by shifting of the bands due to B_0 changes with inhalation (i.e., field shifts due to changes in the size of the lungs). Another source of instability in the bSSFP data could be respiratory related motion, as strong correlations were seen both very inferior (in the neck region) and the in most superior part of the brain.

By far the strongest correlations with the cardiac cycle are localised to the

CSF and major blood vessels, which suggests that the instabilities are driven by flow effects. There are various mechanisms by which flow can introduce signal fluctuations: 1) inflow effects, 2) outflow effects, and 3) in-plane flow effects. We begin by considering the familiar scenario of 2D imaging that excite thin planes, and then discuss differences in 3D imaging.

Inflow Effects

Inflow effects are signal variations caused by the inflow of fresh magnetisation from outside the excited slice. Both SPGR and bSSFP are steady-state sequences. As discussed in Chapter 2 a steady-state is an equilibrium state where M_z is less than M_0 , but does not change from one TR to the next. When fresh magnetisation enters the imaging slice its longitudinal component will be larger than the steady-state magnetisation. The signal of inflowing spins therefore tends to be higher than that for the spins that are in steady-state [111]. Due to pulsatility the amount of inflowing spins will vary each TR creating fluctuations in the signal.

Outflow Effects

Outflow effects are signal contributions by flowing spins that have already left the imaging slice in a preceding TR. Even outside the slice, spins with a residual transverse component may still contribute to the total signal even if they no longer experience additional RF-pulses [112]. In SPGR this effect is not observed as the gradient spoiler at the end of each TR prevents spins that don't experience the subsequent RF-pulse from contributing to the signal. In bSSFP, however, all imaging gradients are rewound at the end of the TR, which means that even the spins that don't see the next RF-pulse will contribute to the signal when the next set of imaging gradients is played out. This difference can be better appreciated by considering the k -space trajectory that residual magnetisation experiences in

a subsequent TR in the event that it does not experience the RF pulse (i.e., has flowed out of the slice). For bSSFP the imaging gradients take the readout back to the centre of k -space at the end of each TR and thus the same k -space trajectory is traversed in the subsequent TR regardless of whether an RF-pulse is experienced. In the case of SPGR, on the other hand, the gradient spoiler at the end of the TR pushes the magnetisation away from the centre (e.g., $2 * k_{max}$) and mainly freshly excited spins (which have coherent phase and thus start at the centre of k -space) will contribute to the signal in the subsequent TR.

In-plane Flow

Finally, there can be signal instabilities caused by flowing spins that move within the FOV without leaving the imaging plane. When the imaging gradients are not compensated for flow, moving spins will experience a different phase evolution based on their changing position during the application of the gradients. When the motion is coherent (either due to flow or bulk motion) this will result in a mean phase offset compared to the static tissue, which may result in fluctuations being seen in the phase data [16]. In bSSFP phase discrepancies can also lead to a modification of the steady-state in bSSFP or even cause signal cancellation [112, 111]. Other magnitude effects can arise when the amount of displacement within a voxel varies; for example due to the varying flow velocities when there is laminar flow. In our data in-plane flow effects are seen in regions where the flow is parallel to the imaging plane such as the CSF surrounding the brainstem and the carotid arteries (Fig. 5.1).

Interpretation in the Context of 3D imaging

In the 2D data presented here, the signal fluctuations in regions of CSF where the flow is perpendicular to the imaging slice are likely enhanced by the inflow and

outflow effects. This effect will clearly be reduced in 3D acquisitions as a larger slab is excited repeatedly and only a small portion of inflowing or outflowing spins contribute to the total signal. In-plane flow effects, however, will still be present. In our 3D acquisition the third dimension is encoded by phase encoding. Spins flowing along the third dimension may thus experience phase perturbations due to the phase encode gradients as discussed above. The 3D readout thus effectively exchanges in- and out-flow effects for in-plane (or better “in-slab”) flow effects. To be confident that this effect is causing the instabilities we see in 3D data (shown in Section 4.5) it is important to investigate this effect more thoroughly.

5.3 Characterization of the Effects of Flow

As indicated by the previous experiment a major portion of the instabilities appear to originate from pulsatile CSF flow. Particularly in bSSFP data, where the CSF signal level is very high, disturbances in the CSF signal have profound effects on adjacent tissue. In the following sections we aim to characterize the flow sensitivity of our sequence and investigate whether the phase perturbation effects as described by Bieri and Scheffler [113] can explain the signal instabilities we see in our 3D bSSFP data.

Firstly, the first-order moments are calculated for our stack-of-EPI readout with commonly used parameters to model the sensitivity to flow and motion at various time points during the readout. Secondly, an *in vivo* experiment is described in which the effect of pulsatile flow was mimicked on static tissue by altering the phase increment of the RF pulses over the course of the acquisition. This experiment should answer the question of whether flow-induced phase discrepancies at the end of the TR are large enough to induce disturbances in the

steady-state signal as described by Bieri and Scheffler [113].

5.3.1 Quantification of the First-order Gradient Moments in a bSSFP Acquisition with a 3D Stack-of-Segmented EPI Readout

As discussed in Chapter 2 in a homogeneous field the amount of phase accrual experienced by a given spin depends on the area under the gradient waveform and its position (Eq. 2.8). This equation, however, only assumes static spins for which $X(t) = x_0$. We can extend this equation to include motion with as:

$$x(t) = x_0 + v_0 t + \frac{1}{2} a_0 t^2 + \dots \quad (5.1)$$

where x_0 is the initial position, v_0 is the velocity, and a_0 the acceleration of the spin. The amount of phase accumulated at time t caused by a gradient along the same direction is then given by:

$$\varphi(t) = \gamma \int_{t_0}^t G(u) x(u) du \quad (5.2)$$

$$= \gamma \int_{t_0}^t G(u) \left(x_0 + v_0 u + \frac{1}{2} a_0 u^2 + \dots \right) du \quad (5.3)$$

In order to describe the effect of the gradients on static spins and moving spins independently we define so-called *gradient moments*, with the n^{th} moment given by [16]:

$$m_n(t) = \int_{t_0}^t G(u)u^n du \quad (5.4)$$

such that we can write:

$$\varphi_0(t) = \gamma m_0(t)x_0 \quad (5.5)$$

$$\varphi_1(t) = \gamma m_1(t)v_0 \quad (5.6)$$

$$\varphi_2(t) = \gamma \frac{1}{2} m_2(t)a_0 \quad (5.7)$$

The zeroth-order moment, $m_0(t)$, is a measure that characterizes the amount of phase accrual static spins at a given location experience by the gradient (Note that, for bSSFP the net gradient, and thus m_0 , is zero at the end of each TR). The amount of flow-induced phase accrued by a flowing spin is dependent on the size of m_1 and the flow velocity (Eq. 5.6). The total phase a spin accumulates due to the gradient waveform is the summation of the contributions from each of the moments. For simplicity we will only consider the zeroth- and the first-order moments in further discussions.

Methods

First, we will set out to calculate the magnitude of the m_1 moments for a typical bSSFP sequence with a stack-of-segmented EPI readout [10], which is the 3D extension of the 2D segmented EPI readout used in our regression analysis. In this 3D readout data are acquired starting from the lowest k_z -plane (with k_x being the frequency encode direction and k_y the EPI blip direction, Fig. 5.2(a)). Once all the EPI segments in that plane are collected, the next k_z plane is acquired by

adjusting the 3D phase encode gradients. This is repeated until all planes in k_z are collected (Fig. 5.2(b)). Rephasing gradients at the end of the readout ensure a zero gradient area (m_0) in all three directions prior to the next RF-pulse.

Gradient waveforms were generated using dedicated simulation software within the Siemens programming environment. The following parameters were used for the simulation: TR/TE=14/7 ms, Matrix=96×96×16, Resolution=2×2×2 mm, Bandwidth=1860 Hz/pixel, 8 lines per TR. These parameters are representative of a typical bSSFP fMRI acquisition. The gradient waveforms were subsequently read into Matlab (Mathworks, Natick, MA). The first-order gradient moments were calculated following Equation 5.4 and inspected at various time points in the TR. In particular we examined the m_1 moments at the time that echoes are formed to assess the phase of each of the collected k -space lines. For example, if flow is pulsatile, a non-zero m_1 at the echo times will cause a periodic phase modulation across k -space, which results in image artifacts. Further we examined the m_1 just prior to the next RF pulse, because phase perturbations at $t=TR$ can cause disturbances in the steady-state signal.

Results

Fig. 5.2 shows the k -space trajectory, gradient waveforms, and first-order moments for the simulated acquisition using our stack-of-segmented EPI readout. Panel (a) shows the k -space trajectory of a segmented EPI plane, with the different shots colour-coded in different tones of blue. Panels (c) and (d) show the x-gradients, G_x , and corresponding $m_{1,x}$ overlaid for each of the TRs, and panels (e) and (f) show G_y and $m_{1,y}$, respectively. The G_x waveforms are identical between the shots, whereas the G_y waveforms differ in the magnitude of the pre- and de-phaser blips. The result on the first-order moments is that for $m_{1,x}$ the value is constant between the TRs ($m_{1,x}=-20 \text{ mT} \cdot \text{ms}^2/\text{m}$), whereas $m_{1,y}$ varies

between -8 and $-16 \text{ mT} \cdot \text{ms}^2/\text{m}$, depending on which segment is acquired. The gradient waveforms that encode for the different k_z planes are shown by panel (g). The colour coding corresponds to the planes shown in panel (b), which shows a side-view of the stacked EPI planes along k_z . It is seen that the first order moments in the z -direction, $m_{1,z}$, show the largest variation ranging between -50 and $42 \text{ mT} \cdot \text{ms}^2/\text{m}$ for the different k_z planes. It is worth pointing out that the gradient waveforms each have varying amounts of first-order moments at the echoes of each of the lines (denoted by the red x marks). For the gradients in the x -direction, $m_{1,x}$ is $-20 \text{ mT} \cdot \text{ms}^2/\text{m}$ for all uneven lines, and $-17 \text{ mT} \cdot \text{ms}^2/\text{m}$ for the even lines. The first-order moments in the y -direction follows a staircase pattern for which moments of the lines closest to the centre of k -space are -8 and $-16 \text{ mT} \cdot \text{ms}^2/\text{m}$. $m_{1,z}$ finally, is constant for during each EPI plane, but slowly increases from -4 to $4 \text{ mT} \cdot \text{ms}^2/\text{m}$ as we collect the various planes along k_z . These variations in the first-order moments can cause phase discrepancies between the lines in k -space, which may result in ghosting and blurring in image space. This will be discussed more extensively in the following sections.

5.3.2 The Effects of Flow-induced Phase perturbations on the Steady-state signal *in vivo*

As discussed in Chapter 2 in bSSFP the magnetisation of each spin isochromat (i.e., ensemble of spins at a specific resonance frequency) follows a unique trajectory during the TR. In the steady-state the relaxation and precession processes are exactly cancelled by the RF pulse every TR. This means that the magnetisation of each isochromat needs to follow the same trajectory from one TR to the next in order to maintain the steady-state (i.e., the amount of relaxation, the flip angle, and the precession relative to the RF phase, needs to be constant).

If spins are moving within the imaging plane during a readout with non-zero

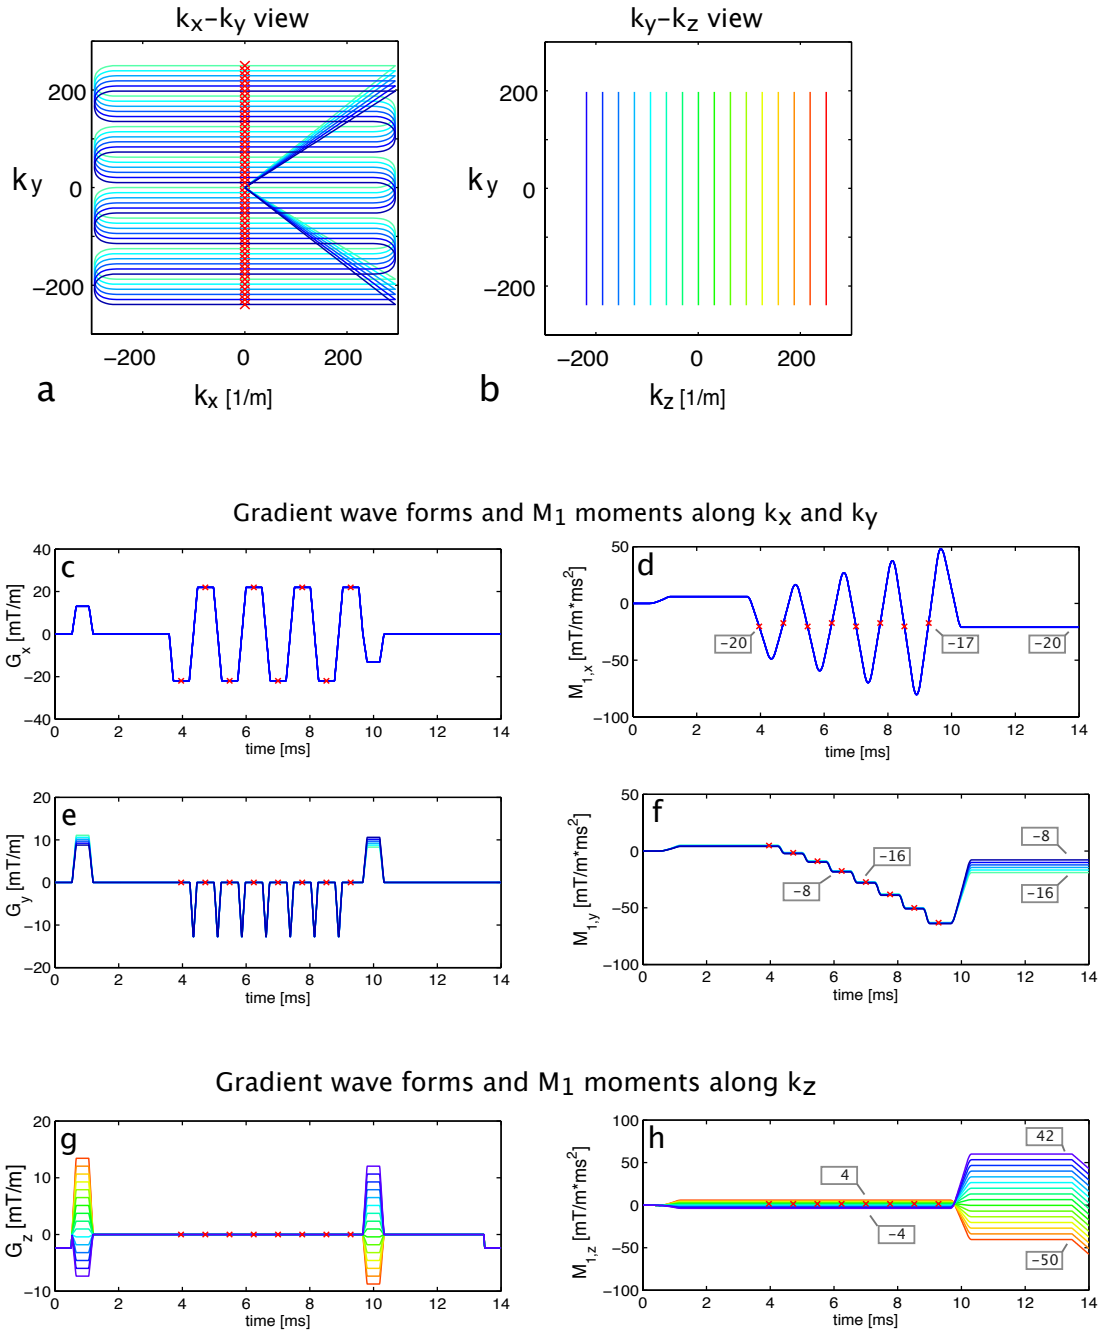


Figure 5.2: The gradient waveforms (G_x , G_y , and G_z) with their corresponding m_1 -moments for the stack-of-segmented EPI readout using parameters typically used in fMRI protocols. Two side views of the k -space trajectory are shown in panel (a) and (b). The m_1 -moment values [mT \cdot ms 2 /m] are listed at the echo times and at the end of each TR in panel (d), (f), and (h)

m_1 they will accrue a certain additional phase depending on the m_1 at the end of the TR and their velocity (Equation 5.6). In practice the flow induced phase, φ_1 , can vary over time if either the flow velocity changes (e.g., pulsatile flow) or if m_1 changes from one TR to the next. Variations in phase accrual due to varying velocity or m_1 can lead to instabilities in the signal, as the steady-state signal conditions are no longer met. Bieri *et al.* reported significant signal loss for acquisitions with linear phase increments above 1° . Foxall [114], on the other hand, has demonstrated that when the changes in phase are small enough only a small modulation of the steady-state is seen, which is stable over time.

CSF flow in the brain, and particularly around the brainstem, is pulsatile, so the flow velocity and therefore phase discrepancy will vary semi-periodically with a frequency of approximately 1 Hz (the average heart rate at rest). The aim of the following experiment is to estimate the magnitude of the phase perturbations and to test whether the non-constant phase evolutions are large enough to result in the signal instabilities we see in our 3D data.

Methods

The velocity of CSF has been investigated extensively in the literature. A number of studies has reported maximum velocities around 10 mm/s [115] in the cerebral aqueduct (which connects the third and fourth ventricle). Linninger *et al.* found that flow velocities in the SI direction in the pontine cistern (the CSF space anterior to the pons) ranged from -5 mm/s to 14 mm/s over the course of a cardiac cycle [116]. By inserting these values and the determined $m_{x,1}$ value of $-20 \text{ mT} \cdot \text{ms}^2/\text{m}$ into Equation 5.6 we estimate φ_1 to vary from -1° to 5° degrees over the course of one cardiac cycle. This calculation uses the first-order moment in the RO-direction, because this is the direction, which for brainstem imaging we would place along SI in order to avoid wrap-around problems. The additional

advantage of placing the RO-direction along SI is that m_1 at the end of each TR will be constant, so the only variation in phase will be induced by the pulsatility of the flow.

To test the effect of a periodic variation of phase *in vivo*, but in a controlled manner, we decided to test the effect of a slowly varying phase on static (non-flowing) brain tissue, by externally introducing a slow varying phase. This was achieved by modifying the existing sequence to allow a sinusoidal variation in the phase of the RF-excitation. Note that, in terms of the steady-state signal it does not matter whether the additional phase is induced by flow or externally by the RF pulse. Only the amount of precession 'relative' to the direction of the RF pulse is important as explained in Section 2.4.

Three data sets were acquired in a healthy volunteer on a 3T Siemens Trio system using the following parameters: TR/TE=12/5.7 ms, $\alpha=30^\circ$, matrix = $96 \times 96 \times 24$, res= $2 \times 2 \times 2$ mm, 8 lines per TR, $T_{vol}=3.5$ s, 30 volumes per run. After a dummy period of 6 seconds (i.e., a period in which RF pulses are applied, but no data are collected) the first 15 volumes were acquired with a constant RF-phase, after which the RF-phase was varied in a sinusoidal fashion at a frequency of 1 Hz. The amplitude of the sinusoid varied between the runs and was set to $\varphi_{sin}=0.5^\circ$, $\varphi_{sin}=1.5^\circ$, and $\varphi_{sin}=2.5^\circ$ (i.e., for the $\varphi_{sin}=2.5^\circ$ data set the phase was varied sinusoidally between -2.5° and 2.5° at a frequency of 1 Hz). The acquisition was localised to the most superior region of the brain as this is generally least confounded by motion or pulsatility. In order to ensure the analysis only included static tissue and not CSF, subject specific white matter ROIs were defined by segmentation of the subject's structural with FSL FAST [91], which were subsequently transformed into time-series data space using FSL FLIRT [90]. The data sets were assessed on temporal stability by inspection of the voxel time courses and the standard deviation over time (tSD).

<i>data set</i>	φ_{sin}	Flow sim OFF			Flow sim ON			
		<i>Mean</i>	<i>tSD</i>	<i>tSNR</i>	φ_{sin}	<i>Mean</i>	<i>tSD</i>	<i>tSNR</i>
$\varphi_{sin}=0.5$	0°	148.2	3.6	43.5	0.5°	148.2	3.6	42.5
$\varphi_{sin}=1.5$	0°	148.1	3.6	43.5	1.5°	148.3	4.2	38.1
$\varphi_{sin}=2.5$	0°	148.7	3.6	43.2	2.5°	148.8	5.0	32.9

Table 5.1: Mean Signal, tSD, and tSNR within the grey matter ROI for each data set. “Flow sim OFF” reports the measures taken from the first 15 volumes when the RF increment is kept constant over time. The measures under “Flow sim ON” are taken during the second half of the run, when the RF increment is varied to simulate pulsatile flow

Results

Table 5.1 lists the mean signal, tSD and tSNR within the WM ROI during the first 15 volumes (when the RF-phase is kept constant) and the last 15 volumes (when the RF-phase is varied). It shows that the mean signal and tSD is the same across all three data sets for the first 15 volumes. The tSD increases with increasing RF-phase variation while the mean signal remains unaffected.

Fig. 5.3 shows the normalised tSD and the RAW voxel time courses within the white matter ROI shown in blue in Fig. 5.3(a). The RAW signal time courses show that the signal remains essentially unaffected by an RF-phase variation of 0.5° (Fig. 5.3(b)). The other two data sets clearly show an increase in signal fluctuations as soon as the RF-phase starts to vary (from volume 15 onwards). This is confirmed by the increased tSD in Table 5.1 and the maps shown in Fig. 5.3(c). The tSD maps further show that the increase in tSD is largest in the frontal regions of the brain. It is known that the steady-state signal modulation due to phase variations are dependent on local field inhomogeneities [113]. Field inhomogeneities caused by the nasal sinuses might explain why the tSD increase is larger in the frontal regions of the brain.

Discussion

The *in vivo* experiments presented here show that pulsatile flow can introduce phase variations large enough to cause instabilities in the steady-state signal.

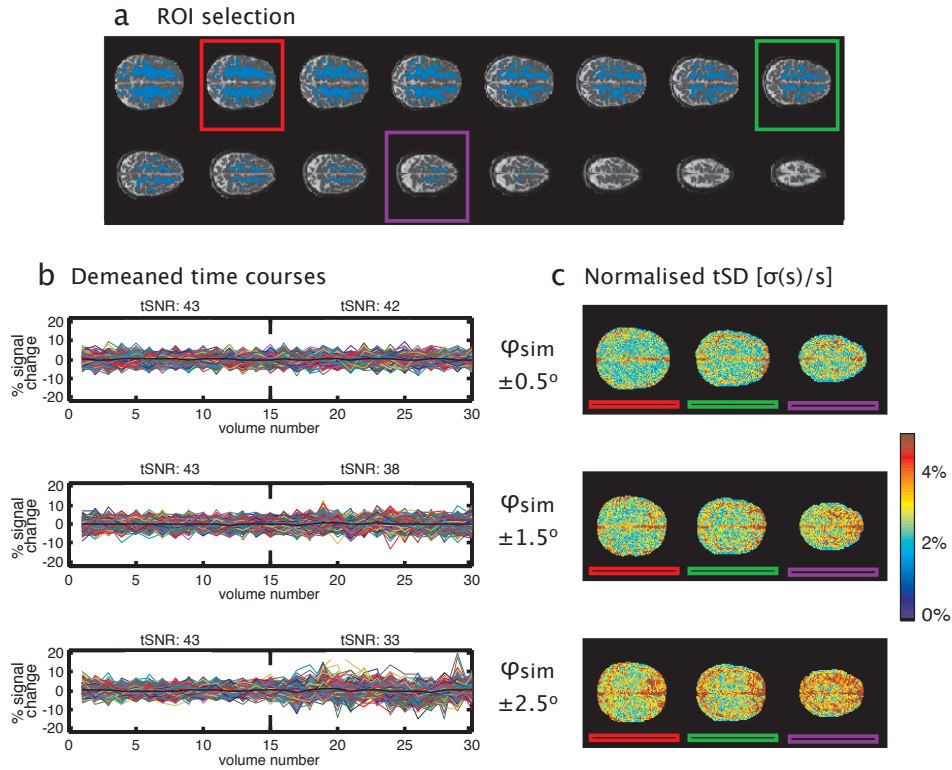


Figure 5.3: Results showing the effect of a slowly varying RF increment mimicking flow-induced phase accrual. Panel (a) shows the ROI selection covering the white matter in blue. Panel (b) shows the individual time courses of voxels within the WM ROI. It is seen that the stability decreases when the RF increment is varied (volumes 15–30). The effect is most clearly observed for the $\varphi_{sim} = \pm 2.5^\circ$ data set. Panel (c) shows the tSD over the last 15 volumes (when the RF increment is varied) normalised against the mean signal in the WM ROI. The tSD increases across the brain when the RF increment is varied.

Based on the gradient waveforms of our acquisition we calculated the first order moment in the readout direction to be $20 \text{ mT} \cdot \text{ms}^2/\text{m}$. For a velocity range of 20 mm/s over the course of the TR, the total phase variation was estimate to be approximately 5° . The *in vivo* data showed a considerable drop in tSNR of 24% when the RF phase is varied sinusoidally with an amplitude of $\pm 2.5^\circ$ at a frequency of 1 Hz, suggesting this mechanism could indeed be the source of the signal instabilities. Note, however that this experiment is not intended to provide quantitative results about the magnitude of this effect in CSF. The T_2 of CSF is significantly longer than that of WM tissue, which is likely to result in greater instabilities, as the transverse magnetisation from a larger number of preceding

TRs will contribute to the signal in the current TR. Moreover, the actual flow dynamics during the cardiac cycle are likely to be more complex than the simple sinusoid used in this experiment. This experiment is therefore merely intended to serve as a test to investigate whether the steady-state is significantly disrupted when the phase is varied a few degrees at a frequency of 1 Hz. The fact that we see instabilities arise in WM is a strong indication that those same effects will be present in CSF. The subsequent sections will consider two potential approaches to addressing the instabilities: first by spoiling CSF signal in bSSFP and second by reordering the k -space acquisition.

5.4 Spoiling CSF in Balanced SSFP

For this first of two correction methods we pursued the pragmatic approach of spoiling the CSF signal in bSSFP. In bSSFP the signal of CSF is much larger than that of GM and WM. Any ghosting of CSF into the tissue due to steady-state instabilities will therefore result in large signal fluctuations in the regions of interest. By spoiling CSF we hope to prevent CSF from reaching a steady-state in order to avoid the instabilities caused by disruptions of the steady-state condition (due to phase errors at $t=TR$).

5.4.1 Introduction

One of the limitations of steady-state acquisitions is that contrast manipulation is more difficult. Spins need to reach steady-state after application of the preparation pulses (i.e., the use of an inversion pulse to null CSF). Although various catalization schemes have been developed that speed up the approach to steady-state [117], the schemes are often only effective for on-resonance spins and still require a certain amount of time before the signal is truly stable from one TR to

the next.

In this section we explore the possibility of spoiling the CSF signal by intentionally introducing phase variations between subsequent TRs. As demonstrated in the previous section the steady-state signal formation is highly dependent on the amount of phase accrual during the TR and small variations in RF phase can lead to substantial signal variations. Various RF phase cycling schemes have been proposed in the past to alter the properties of a sequence. Zur *et al.* [30] introduced in 1991 the use of a quadratic RF phase increment to spoil the coherences between the spins and achieve pure T1 contrast. Scheffler *et al.* later investigated the effect of alternating RF phase variations. In this method the sign of the phase accrual is reversed every other TR ($\varphi, -\varphi, \varphi, -\varphi, \dots$). In non-balanced SSFP sequences, the alternating RF-phase causes spoiling of the magnetisation. In balanced SSFP the effect of an alternating RF increment will set-up two different steady-states [118].

Here we propose to include bipolar gradients into a bSSFP sequence in order to induce phase offsets in the flowing spins, without affecting the static spins. This is similar to the method by Overall *et al.*, in which alternating bipolar gradients were used to set up two different steady-states; one for flowing spins and one for static spins, in order to improve the contrast in angiography [119]. Our aim, however, is not to set up a steady-state for the flowing spins, but to prevent the flowing spins from reaching a steady-state (i.e., spoil the signal of flowing spins). We therefore explore the use of (semi-) random bipolar gradients instead, which will hopefully have a similar effect on flowing spins as RF-spoiling has on static spins [30].

5.4.2 Methods

By changing the size of the bipolar gradients each TR we have control over how much phase is induced to the flowing spins. We can therefore create a phase

cycling scheme that only affects the flowing spins and not the static spins. In order to test which phase cycling scheme would be effective we have conducted a set of simulations.

The influence of various phase cycling schemes on the bSSFP signal formation was investigated by numerical Bloch simulations. In these Bloch simulations 3×3 rotation and relaxation matrices were applied iteratively to a set of 100 magnetisation vectors to analyse the approach to steady-state in a given voxel, the mean steady-state signal, and the stability. All simulations were performed using the relaxation parameters for CSF at 3 Tesla ($T_1=4000$ [120], $T_2=500$ [121]). Other parameters included: TR=12 ms, TE=6 ms, $\alpha=30^\circ$. The simulations considered 500 repetitions (TR=12ms) to ensure the steady-state is reached. Five different phase cycling schemes were tested on a fully balanced (i.e., no simulated dephasing) acquisition. These cycling schemes included: 1) no flow induced phase, 2) randomly varying phase between -10° and 10° , and 3) perfectly alternating phase between -3° and 3° . A perfectly alternating phase, however, requires a constant flow as the amount phase accrual is determined by m_1 (which we control with the bipolars), but also the flow velocity. We therefore included two extra simulations, which simulated alternating phase of $\pm 3^\circ$ and $\pm 10^\circ$, but with a slight perturbation to mimic non-stationary flow. Plots of the simulated phase cycling schemes are shown along with the results in Fig. 5.4.

5.4.3 Results and Discussion

Fig. 5.4 shows the approach to the steady-state for on resonance spins when various phase cycling schemes are used. For reference the top panel shows the approach when the RF phase is kept constant. The inset shows a constant φ_{bipol} of zero over the complete time course (i.e., when no additional phase is induced by the bipolar gradients). Panels (b)–(e) show the steady-state approach for various

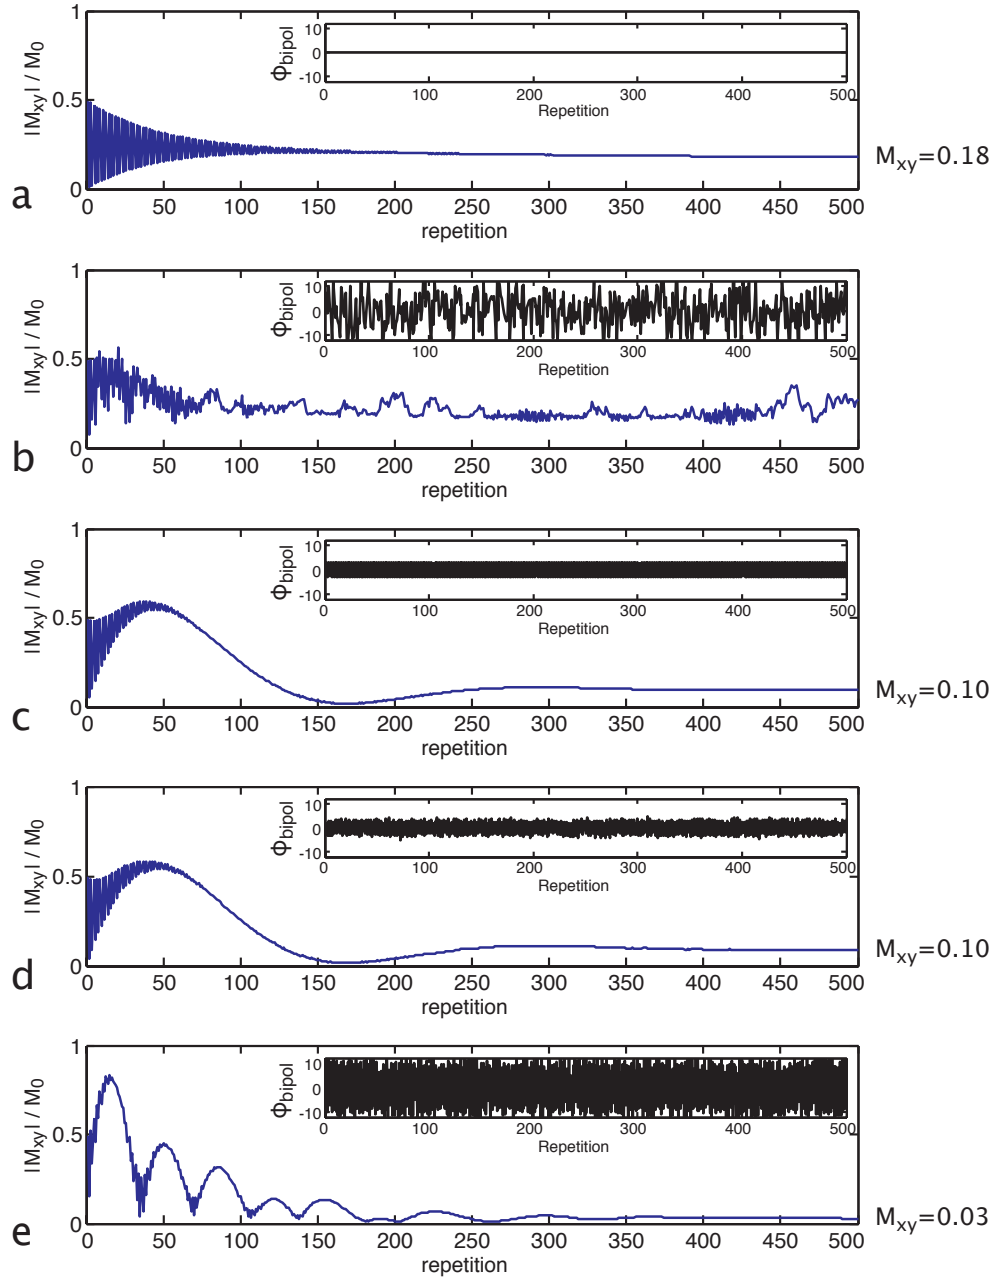


Figure 5.4: The approach to steady-state for on-resonance spins with various RF cycling schemes. The phase cycling schemes are shown as insets. The evolution of the magnetisation is shown in blue. (a) shows the approach when a constant phase of $\varphi_{bipol}=0$ is used. In Panel (b) a random phase variation between -10° and 10° is applied. (c)–(e) depict a phase cycling scheme in which the phase is alternated between subsequent TRs. Alternation between -3° and 3° are shown in (c) and (d), whereas (e) shows $\pm 10^\circ$ alternations. (d) and (e) are simulated with slight imperfections of the achieved phase alternations.

phase cycling schemes. When the phase is varied randomly between -10° and 10° from one TR to the next (Panel (b)) the signal evolution is very erratic and a stable steady-state is not reached within 500 repetitions. When the sign of the phase accrual is alternated between -3° and 3° every other TR (Panel (c)) we see that the signal does evolve to a steady-state with reduced signal ($M_{xy}/M_0=0.10$ compared to the reference signal ($M_{xy}/M_0=0.18$). When the amount of phase between which is alternated is not perfect (modelling slight imperfections due to non-constant flow velocities) we see that a steady state is still achieved (Panel (d)) with similar amplitude. The simulations show that the approach to the steady-state is not affected very much by small imperfections in the alternating phase cycling scheme. If we are able to achieve an alternating phase of $\pm 10^\circ$ (Panel (e)), the signal is reduced to $0.03M_0$. This suggests a 6-fold reduction in CSF signal, which would be a very acceptable level of spoiling.

Fig. 5.5 plots the signal profiles for three of the phase cycling schemes investigated. The mean signal over the last 100 repetitions is plotted in blue and the error bars (red) denote the standard deviation of the signal over this period. When the phase is varied randomly between -10° and 10° the signal varies considerably across the passband region of the profile. The large standard deviation suggests a steady-state has not been reached for any of the off-resonance frequencies (the centre of the profile denotes on-resonance. Figures 5.5(b)) and 5.5(c)) show the signal profiles for the simulations with (imperfect) the phase cycling schemes with alternating phases of $\pm 3^\circ$ and $\pm 10^\circ$, respectively. For reference these plots are overlaid with the non-spoiled balanced SSFP profile (solid black line) and the signal that we would obtain for a non-balanced RF-spoiled sequence (dashed line). Panels (b) and (c) show that the signal profile only demonstrates a sharp dip at the centre of the profile. This observation is confirmed by work that was recently presented, where the same phenomenon was observed and ascribed to

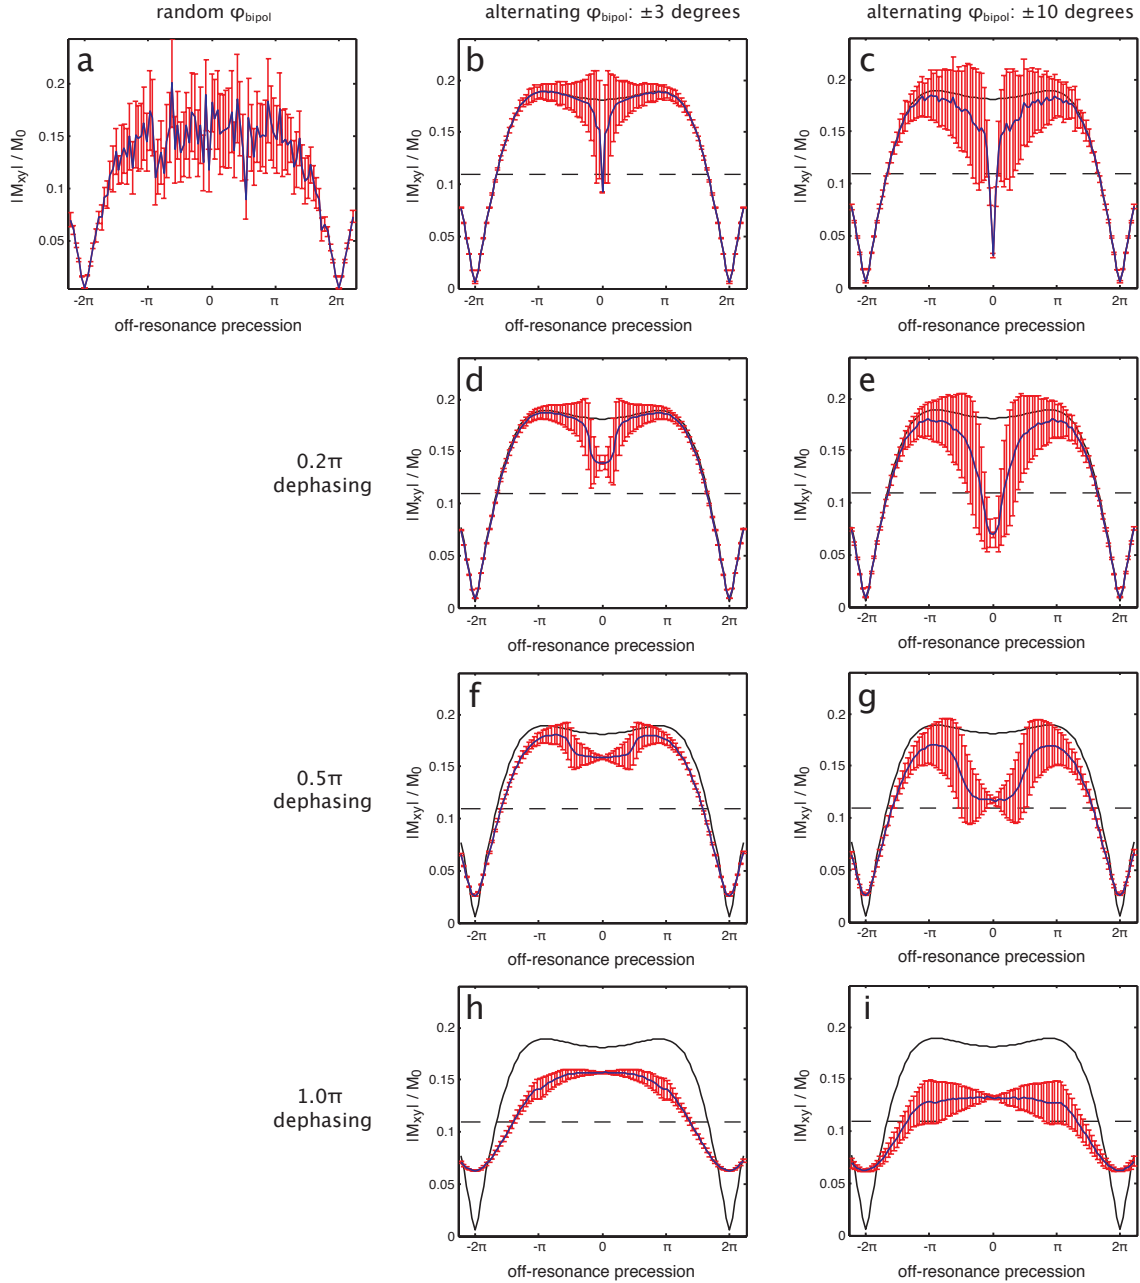


Figure 5.5: Frequency profiles of the three RF phase cycling methods. Panels (a)–(c) show the profiles obtained using the random and alternating RF phase cycling schemes when dephasing is assumed to be zero. Panels (d)–(i) show the profiles when increasing amounts of dephasing are simulated. These plots show that spoiling is only achieved at the centre of the passband unless a large amount of dephasing is achieved, which limits the applicability of this method.

singularities in the bSSFP signal equation [122]. The narrow dip will obviously limit the applicability of the technique, because even with careful shimming we will not be able to reduce the field inhomogeneities around the brainstem enough to confine all the spins to this narrow frequency band in the centre of the profile. Panels (d)–(i) show the signal profiles obtained the same phase cycling schemes are used, but with additional dephasing simulated. It is reasonable to assume that a certain amount of dephasing will occur due to the nature of the flow (i.e., a certain amount of laminar flow is to be expected). Moving down the columns an increasing amount of dephasing is simulated. The effect of dephasing is that the range of off-resonance frequencies on which spoiling has an effect widens. The amount of spoiling is effectively averaged over a range of off-resonances. As we step through the increasing amounts of dephasing (moving down the columns) we see that the signal profile starts to resemble the non-balanced RF-spoiled case (dashed line), which has a constant signal for all off-resonance frequencies. This is consistent with work by Ganter, who analytically showed that for spoiled GRE sequences to reach a homogeneous steady-state strong crusher gradients were needed in combination with RF phase cycling [123].

To ensure an even response across the signal profile a large amount of flow-induced dephasing as shown in panels (h) and (i) would be desirable. However, it is unlikely that we can ensure a dephasing of π in each CSF voxel as the flow profiles around the brainstem are highly variable [124] and will fluctuate over time due to pulsatility. We therefore conclude that this technique is not likely to be robust *in vivo* and decide not to pursue this approach.

5.5 Readout Synchronization Simulations

In this section a second method to reduce the signal instabilities is explored. In the previous approach we attempted to improve the stability by removal of the unstable signal component (flowing CSF). Here we take a different approach. Instead of focussing on the dynamics behind the signal formation we focus on the effects of the readout. We investigate how the signal fluctuations map across k -space and how this is expressed into image space. Based on simulations that use real data, we try to come up with alternative acquisition strategies, which are more forgiving to fluctuations in the measured MR signal.

5.5.1 Theory

In our readout the 3D k -space data are combined from multiple segments. Temporal fluctuations in the MR signal over the course of the acquisition cause discrepancies between the different segments, which result in modulations across k -space. Depending on the shape of the modulations (e.g., smooth, erratic, or periodic) the effect in image space will be different. An intuitive way to describe this is to picture the modulation as a complex filter with which k -space is multiplied. The effect in image space is then described by a convolution of the moving tissue with the Fourier transform of this filter, which we will refer to as the point-spread function (PSF). If the modulation is erratic across k -space, its PSF will be broad, dispersing the signal up to several voxels away from its original location. In general, any abrupt signal changes, but more specifically periodicity, leads to ghosting artifacts in the image domain. When, in a time-series acquisition, the pattern of the modulations across k -space differ from one volume to the next, the artifacts in image space (e.g., ghosts) will also vary from volume to volume creating large temporal signal instabilities.

We propose to order the acquisition relative to the cardiac or respiratory cycle to induce a smooth modulation across k -space. Rather than disruptive ghosting artifacts, this is expected to introduce a minor blurring of the image, which is a subtler and more acceptable artifact.

5.5.2 Methods

The results obtained by our regression analysis (Section 5.2) suggest that the dominant source of physiological variation is the cardiac cycle, which affects both signal magnitude and phase. In this section, we present simulations of several potential synchronization schemes for 3D data acquisition based on these observations. Highly realistic physiological effects are achieved by using the 2D single-slice bSSFP data presented above to simulate the formation of a 3D k -space volume in 3D readouts.

The basic principle behind these simulations is depicted on a numerical phantom in Fig. 5.6. Our simulation begins with the rapidly acquired 2D time-series images (with complex data), which are transformed into 2D k -space matrices. Next, 3D k -space volumes are simulated from the 2D k -space time series by concatenating adjacent time points to form the third dimension (k_z). An inverse 3DFT is applied to each 3D k -space volume to obtain a time series of 3D images. Thus, if we start from n_t time points and simulate 3D volumes with n_z slices, the output is a time series of n_t/n_z 3D volumes. For a temporally stable signal, k -space will be constant along k_z and the inverse transformation of the simulated 3D k -space would correspond to signal for a single slice in the middle of the imaging volume, with no signal outside this central slice (Case I, Fig. 5.6(b)). In the presence of temporal instability, however, the signal along k_z will fluctuate, and introduce dispersed signal into adjacent slices that varies from volume to volume (Case II). This is the hypothesized source of the increased signal instabilities in regions that

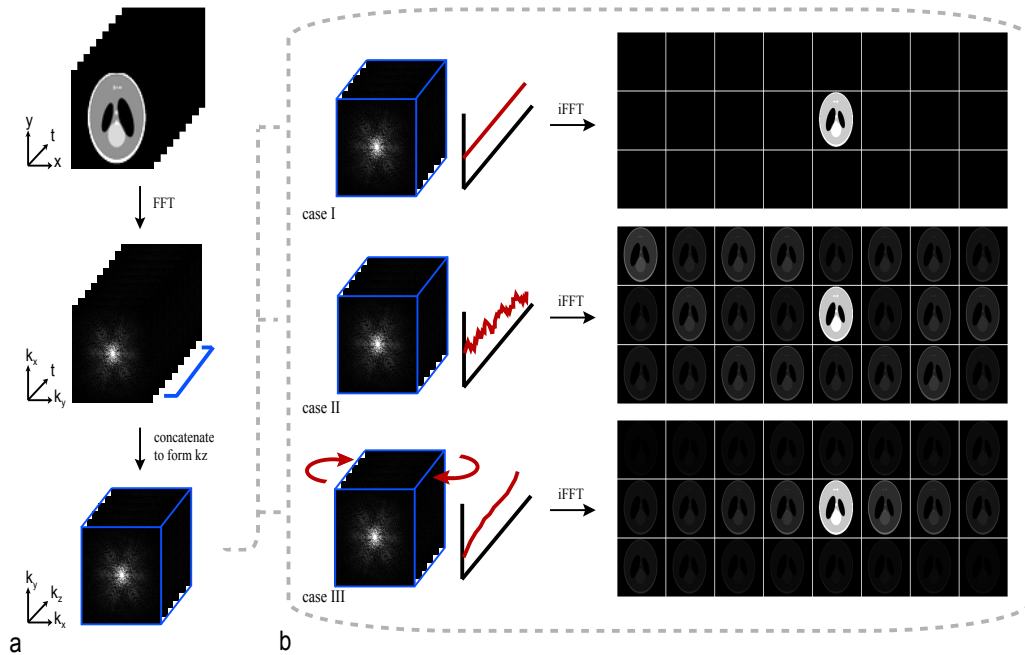


Figure 5.6: Simulations of 3D acquisitions using rapidly acquired 2D data (demonstrated conceptually on a numerical phantom). (a) 3D time-series data were generated by concatenating 2D data in k -space. (b) Case I; when the signal is temporally stable the resulting images will only contain signal in the centre slice. Case II; in the presence of temporal fluctuations from the cardiac cycle the signal varies along k_z , and the signal disperses outside the centre slice. Case III; by reordering the partitions with respect to the cardiac cycle, a smooth function can be imposed along k_z , which reduces the artifacts.

are in close proximity to blood and/or CSF. If we fill in successive k_z -planes at consecutive time points, we simulate our standard acquisition scheme and expect to see signal instabilities similar to standard 3D FMRI.

We use this framework to simulate two methods for stabilizing signal: 1) synchronization of k -space acquisition order to the position in the cardiac cycle and 2) rejection of data acquired during systole. Cardiac synchronization was simulated by reordering the k_z -planes within each volume relative to the cardiac cycle (Case III). Systolic rejection was simulated by discarding any data acquired within 100 ms of systole before concatenating the 2D matrices in the third dimension. The resulting 3D time series (generated with 15 k_z -planes) were assessed based on the signal fluctuations that were present outside the central slice, which is by definition artifactual.

It should be noted that these simulations only consider the artifacts that would arise in the k_z -direction, whereas the ghosting artifacts in the 2D data (Fig. 5.1) occur along the k_y - (blip) direction. In our 3D segmented-EPI acquisition, both of these axes are phase encoded over multiple TRs and therefore are susceptible to physiological artifacts. It was chosen to simulate the k_z -direction because, as the slowest axis of acquisition in the standard scheme, it is expected to have greatest physiological sensitivity.

5.5.3 Results and Discussion

Fig. 5.7 shows simulated signal instabilities for uncorrected bSSFP data and the two correction methods outside the central slice (where no signal should be present in the case of perfect temporal stability). The figure shows the temporal standard deviation (tSD) normalized to the mean signal intensity in the centre slice. We observe signal fluctuations throughout the brain, which can be attributed to thermal noise. In regions of CSF, however, signal fluctuations are significantly higher in the uncorrected data (Fig. 5.7(a), top row). Following cardiac synchronization, these fluctuations are reduced significantly in most areas of physiological noise. When systolic rejection is performed in conjunction with synchronization, a further reduction is observed in some areas, although the additional benefit is small compared to the initial improvement obtained by synchronization alone.

Fig. 5.7(b), shows the average tSD per slice in CSF voxels. This figure shows more clearly the small but consistent additional improvement when systolic rejection is performed. It also demonstrates that most of the signal instabilities arise in Slices 6 and 10 when no correction is performed. This can be explained by the fact that the periodicity of the cardiac cycle creates a PSF with peaks away from the centre (slice 8), which will introduce aliasing in the image due to convolution

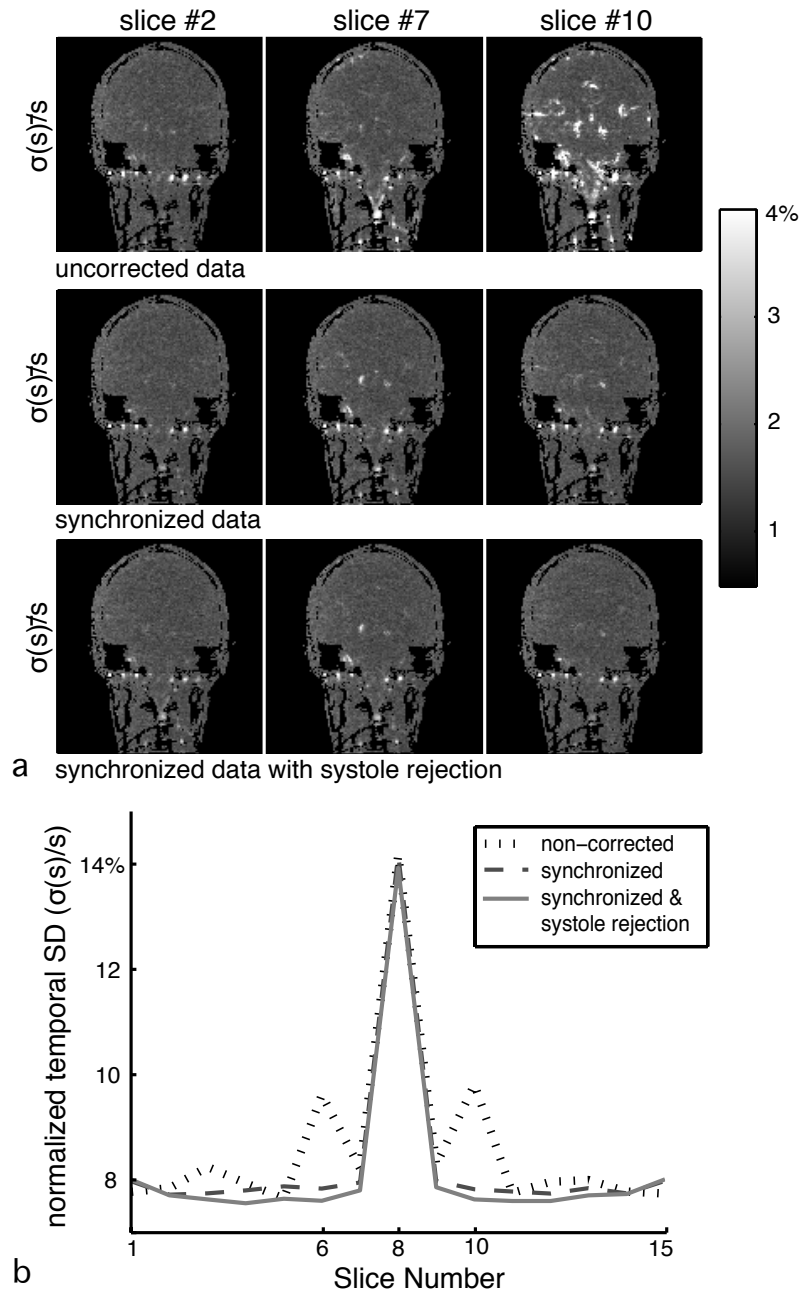


Figure 5.7: Results from the simulated 3D data sets. (a) Spatial maps of the temporal signal fluctuations in slices 2,7, and 10, normalized against the mean brain signal in the 8th (centre) slice ($\sigma(s)/S$). (b) Normalized signal fluctuations per slice averaged over the CSF ROI for uncorrected data and the two correction methods tested. In the uncorrected data (Panel (a), top row; Panel (b), grey line) the signal dispersal caused by pulsatile CSF increases signal variability significantly compared to the baseline fluctuations caused by thermal noise. When the partitions are reordered to the cardiac cycle, these fluctuations are considerably reduced, almost to the same level as thermal noise. When reordering is combined with removal of data that was acquired during systole, the variability is reduced even further, although the additional stability from systolic rejection is small.

with this PSF. The exact location of these peaks depends on the ratio between the cardiac rate and the volume acquisition rate. For the specific parameters and data used in these simulations, most of the signal to be dispersed lies two slices away from the centre. When the k_z -planes are synchronized to the cardiac cycle the PSF has a more benign shape, dispersing less signal into adjacent slices. Moreover, the shape will be constant over time, thus reducing the tSD in slices outside the central slice. For a 3D acquisition cardiac synchronization could therefore be a successful method to reduce the signal instabilities in the brainstem that are originating from adjacent (fluctuating) CSF signal.

5.6 Conclusions

In this chapter the source of the instabilities in 3D multi-shot sequences was discussed and various correction methods were explored. Regression analyses based on simultaneously acquired physiological recordings (i.e., pulse-ox and bellows) showed that the instabilities were mostly confined to regions of CSF and correlated highly with the cardiac cycle. In a second experiment it was shown that flow-induced phase discrepancies can indeed cause considerable disruptions of the steady-state signal. Although spoiling the CSF signal in bSSFP proved to be cumbersome, an alternative method in which the k -space is reordered based on the cardiac cycle showed promising results in simulated data. An additional benefit of this kind of approach over signal spoiling is the fact that the method is not dependent on the signal formation dynamics and is therefore applicable to bSSFP as well as SPGR sequences. For this reason cardiac synchronization was selected as the method to implement *in vivo*. In the following chapter the sequence implementation is discussed and *in vivo* experiments are conducted to assess the effectiveness this method.

Chapter 6

Cardiac synchronization

6.1 Introduction

Based on the results presented in Chapter 5 we propose a prospective method for synchronizing the k -space acquisition to the cardiac cycle, which does not alter the volume acquisition time. This is achieved by combining real-time scheduling of the 3D k -space acquisition with a parallel-imaging reconstruction. In this chapter, several important practical considerations that were not encountered in the simulations are discussed and a specific method is presented. We show that this method is able to reduce temporal instabilities in the brainstem by up to 45% in 3D bSSFP and up to 20% in 3D SPGR time-series data.

6.2 Theory

As discussed in Section 5.5 synchronizing the acquisition relative to the cardiac or respiratory cycle alters the modulation induced by physiological signal fluctuations across k -space. In conventionally ordered acquisitions pulsatile signal fluctuations during the acquisition can induce quasi-periodic phase and magnitude

modulations in the k -space data, which will lead to ghosting artifacts. By ordering the acquisition relative to the cardiac or respiratory cycle a smoother modulation across k -space can be created. Rather than disruptive ghosting artifacts, this creates a minor blurring of the image, which is subtler and is a more acceptable artifact. This principle was first demonstrated in 2D spin-warp anatomical acquisitions [108, 109], where it was shown that synchronization with respect to the cardiac or respiratory cycle reduced the intensity of the ghosting artifacts in the images. These methods were dubbed COPE (cardiac ordered phase encoding) and ROPE (respiratory ordered phase encoding), respectively. Later the same principle was applied to reduce temporal fluctuations in a 3D spiral gradient-echo fMRI sequence [110]. In a time-series acquisition synchronization ensures that the k -space modulation is similar for each volume. Consequently the artifacts in image space will not fluctuate from one volume to the next and therefore the temporal stability is improved.

The disadvantage of the previously proposed synchronization methods, however, is that periods of waiting (i.e. dummy scans) are required during the acquisition, if the current point in the physiological cycle does not correspond to the k -space lines that are left to be acquired (Fig. 6.1). The time it takes to acquire a single volume is thus not fixed if we are to impose the strict requirement that every k -space line is collected with a precise phase in the physiological cycle. This introduces a variable volume acquisition time (T_{vol}), which can be problematic for fMRI paradigms and analysis. We overcome this limitation by combining real-time synchronization with a tailored parallel-imaging reconstruction as discussed below.




Figure 6.1: Schematic showing the readout order of a 2DFT readout with and without synchronization (denoted by stars and squares, respectively). Here, the non-synchronized acquisition would take 3.5 seconds to acquire one 3D volume. When cardiac synchronization is performed, the order of the phase encode (PE) lines is determined by the relative position in the cardiac cycle, or cardiac phase. In order to collect every PE line at the desired cardiac phase, the synchronization inevitably has to skip acquisition for some TR-periods, as seen in the delays between the PE lines toward the end of the acquisition period. If we wish to avoid increases in scan time, we must either acquire non-ideal PE lines or accommodate gaps in k -space.

6.3 Description of the Method

All *in vivo* experiments presented in this chapter were performed on a 3.0 Tesla Siemens TIM Trio system (Siemens Healthcare, Erlangen, Germany) using a 12-channel head receive coil. All subjects had no known neurological deficit and were scanned with informed consent under a technical development protocol approved by the local ethics committee. Simulations, image reconstructions and statistical analyses were performed in Matlab (MathWorks, Natick, MA). As we are aiming to characterize the intrinsic signal fluctuations (rather than under task conditions), all presented data were acquired without an explicit task.

The simulations presented in the previous chapter suggest that ordering the k -space acquisitions with respect to the cardiac cycle will significantly reduce signal fluctuations in inferior brain areas. Unlike retrospective corrections [71, 65, 67], this prospective correction must be implemented in the sequence itself. Real-time measurements of the cardiac cycle (e.g., with a pulse oximeter) are fed into the sequence and used to decide which k -space data to acquire in a given TR.

6.3.1 Synchronization Method

Without loss of generality, we present our method in the context of a 3D stack-of-segmented EPI readouts (Fig. 6.2(a)–(b)) [10], although the method should be compatible with most 3D trajectories. This trajectory has both a blip direction (k_y) and a 3D phase-encode direction (k_z). Fig. 6.2(c) shows example acquisition ordering methods (assuming matrix= $96 \times 96 \times 24$, 12 segments per EPI plane, 8 k_y -lines per segment) along with colour-coded diagrams showing the k -space distribution of cardiac phases for the different methods. The cardiac phase is defined as:

$$\varphi_c[t] = t_{trig}[t]/T_{RR} \quad (6.1)$$

where φ_c is the cardiac phase, t_{trig} is the time that has elapsed since the last trigger and T_{RR} is the current estimate of the duration of the current cardiac cycle. Our goal is to determine the order of segment acquisition that enforces smooth variation of φ_c across k -space. In general, the segment that is acquired in a given TR can be referred to by a segment index, n_{seg} , that gives its location in 3D k -space:

$$\begin{aligned} n_{kz} &= \left\lceil \frac{n_{seg}}{N_{blip}} \right\rceil \\ n_{blip} &= (n_{seg} - 1) \% N_{blip} + 1 \end{aligned} \quad (6.2)$$

where n_{kz} denotes which k_z -plane to acquire, and n_{blip} which segment within the given k_z -plane. N_{blip} is the number of segments per EPI plane. For non-synchronized acquisitions, n_{seg} always increments linearly from 1 to N_{seg} (the total number of segments per volume), starting the acquisition with the first segment at $k_z=1$, and ending with the eighth segment at $k_z=24$ for the above matrix. The

periodic pattern across k -space in the non-synchronized data (Fig. 6.2(c), method 0) reflects that, in this particular example, k -space was collected over a period of 3-4 cardiac cycles. Depending on the current heart rate, this pattern will change from one volume to the next creating time-varying ghosts as discussed above.

To achieve a smooth mapping of cardiac phase across k -space, we let: $n_{seg} = \varphi_c \times N_{seg}$, which is related to the location in k -space via Eq. 6.2. Since the cardiac phase must be determined in real-time, the duration of the current cardiac cycle (T_{RR}) at time t is estimated by the mean duration of the 10 preceding cardiac cycles. During the acquisition an array is filled to track the segments that have been acquired to avoid double acquisitions in subsequent cardiac cycles.

One major consideration in this type of acquisition is that we will inevitably encounter the situation in which the ideal k -space segment determined by the current cardiac phase is no longer available (e.g. when the same cardiac phase was encountered in the preceding cardiac cycle). We must therefore decide whether to acquire a non-ideal k -space segment, which may introduce artifacts, or skip acquisition in the current TR and wait for the ideal cardiac phase to become available for one of the remaining segments (see Fig. 6.1). These periods without data collection not only reduce the efficiency of the sequence, but also introduce a variable frame rate (the rate at which image volumes are acquired). This is because the time that is spent waiting will vary per volume depending on the current heart rate and the cardiac phase at which the volume acquisition commences.

When a fixed frame rate is desirable, such as in fMRI, the most straightforward solution would be to avoid waiting periods by simply collecting the nearest available segment when the optimal segment, n_{seg} , has already been acquired. This is similar to the method proposed by [110] in the specific case that the number of dummy scans is set to zero in order to avoid an increase in acquisition time. This approach, however, is sub-optimal as this will lead to a portion of the

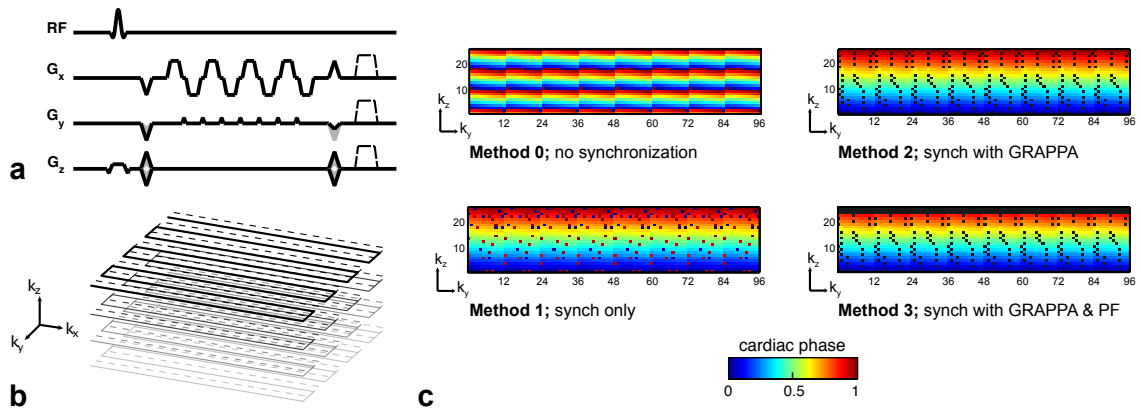


Figure 6.2: The pulse sequence diagram (a) and k -space coverage of the stack-of-segmented EPI readout (b) as used in this study. After slab selective excitation, segmented EPI planes are acquired in the k_x - k_y -dimensions; with the 3D phase encode direction in k_z . At the end of the readout all gradients are rewound for bSSFP whereas gradient spoilers are applied for SPGR (dashed lines) as well as RF spoiling. (c) shows side views of the readout, with the k_x -direction perpendicular to the page. The colour-coding represents the cardiac phase at which each line is acquired. Without synchronization (Method 0) a periodic pattern is visible. When the readout is synchronized (Method 1), the modulation across k -space is smoother, but some segments are acquired with an incorrect cardiac phase. Method 2 avoids these corrupted segments by carefully placing these segments at locations that can be reconstructed with GRAPPA (black points in the subfigure). Method 3 extends the technique by also utilizing a Partial Fourier reconstruction to reject all the data acquired during systole (at one edge of k -space).

segments being collected at incorrect locations in k -space. At the start of the acquisition this won't be a problem as the pool of available segments is large and, if the optimal segment is already acquired, the distance to the nearest available segment will be small. As k -space is steadily filled in, however, the probability of finding a segment, which is still left unacquired, within acceptable range decreases. This thus leaves a portion of the segments acquired with a cardiac phase that is significantly different from optimal. We will refer to this method as "synch-only" (Fig. 6.2(c); Method 1).

We propose to remove the corrupted segments and fill in the missing data with a parallel-imaging reconstruction, such as Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA) [44]. This should reduce the abrupt signal changes in k -space introduced by the incorrectly assigned segments and thus im-

prove image quality. In our implementation, a customized GRAPPA reconstruction is used to replace any segments that are acquired more than 30 ms from their optimal location in the cardiac cycle. This method (Fig. 6.2(c); Method 2) will be referred to as “synch+GRAPPA” in the remainder of this chapter.

One potential further advantage of re-ordering the k -space data with respect to the cardiac cycle is that data acquired during systole (the most unstable period in the cardiac cycle) is confined to one edge of k -space. The data that are acquired around systole can thus be rejected and replaced using a partial Fourier (PF) reconstruction. This formed the third method we investigated, called “synch+GRAPPA+PF” (Fig. 6.2(c); Method 3).

6.4 Sequence Implementation

This section will describe the various steps that were needed to instantiate the real-time feedback loop in the sequence. A generic conceptual description is given rather than a detailed description of the sequence code with the aim that this section will apply to most scanner systems. Therefore, prior knowledge with regards to sequence programming is assumed.

The 3T Siemens Trio system that the sequence is equipped with a physiological measurement unit, which allows the recording of physiological signals (e.g., ECG or Plethysmograph) by the scanner. In order to allow for real-time synchronization the first step was to create a feedback channel between the scanner and the operating system that controls the sequencing. To achieve this our custom sequence used libraries supplied as standard vendor product.

Once this link was established, the cardiac phase and the corresponding segment could be calculated within the sequence code. These calculations were done in the inner most loop of the sequence (i.e., the portion of the code, which calls

a single TR). The shapes of the gradient waveforms to acquire a specific k -space line were calculated using the existing libraries for EPI readouts. The reordering of the lines was achieved by overwriting the segment number in the call to this set of functions based on Eq. 6.2. Since the ordering of the segments was no longer controlled by the native functions within the sequence, a separate lookup table was created to prevent the same k -space lines from being acquired multiple times during one volume acquisition. This approach was later extended to also allow “corrupted” segments to be placed on pre-defined locations that would aid the GRAPPA reconstruction (see Section 6.5.1 below).

By default the computer that controls the sequencing will initialise all the pulses needed for an entire scan as fast as possible after the scan is started. This means that for a long time-series scan all the pulses are calculated during the first few seconds and stored temporarily until they are needed by the scanner. In order to determine the readout order in real time, however, the sequence needed to be modified, so that the timing calculations were performed only one TR ahead.

For each acquired k -space segment information, such as the current cardiac phase and the acquisition order, was passed onto the data header. This information was later used at the reconstruction stage to identify “corrupted” segments that needed to be replaced with our GRAPPA reconstruction.

6.5 Sequence Optimization

This section describes a few initial experiments that were performed during sequence development, which were crucial for the design of the final synchronization strategy and image reconstruction.

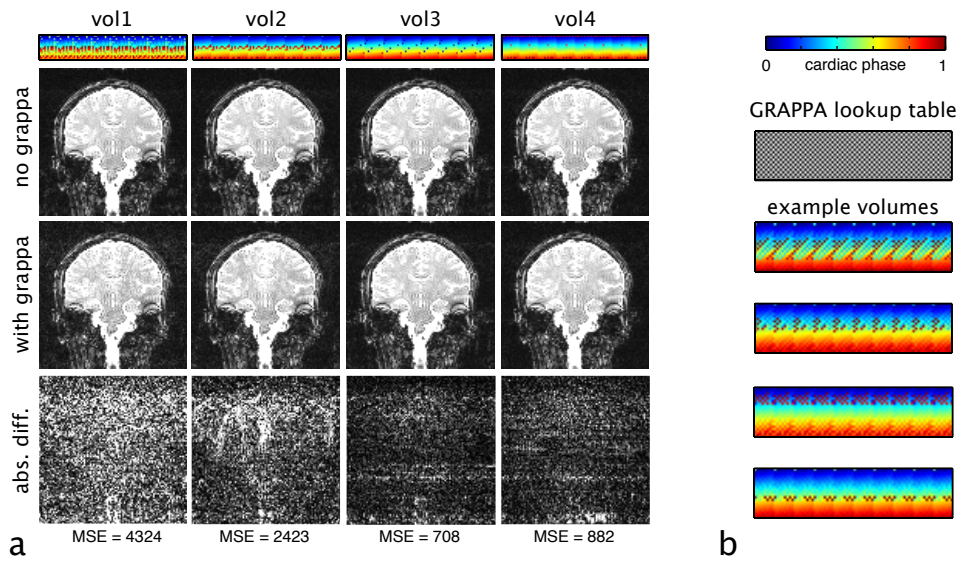


Figure 6.3: Example data acquired with a version of the sequence that did not utilize a second lookup table are shown in (a). The top row of images shows the reconstruction without replacement of corrupted segments (i.e., no GRAPPA). The middle row shows the reconstructed images when GRAPPA is used to replace the corrupted segments. The difference maps (windowed to $1/10^{th}$ of the images) is shown in the bottom row. Panel (b) shows the improved distribution of the corrupted lines by using a second lookup table (the white entries in the top panel) as soon as the segment cannot be placed within 30ms away from its optimal location. This ensures that each corrupted data point is surrounded by non-corrupted data points, which aids the GRAPPA reconstruction.

6.5.1 Optimization of the Reordering Strategy

In Methods 2 and 3 we utilize a GRAPPA approach to retrospectively replace the segments in k -space that could not be acquired with the correct cardiac phase during the acquisition. The quality of the GRAPPA reconstruction will depend on how the “corrupted” segments are distributed across k -space. GRAPPA requires neighbouring k -space points to fill in the missing data. If the un-corrupted data lie too far away from the data that need to be replaced, the GRAPPA reconstruction will become ill-conditioned, resulting in a poor reconstruction. Fig. 6.3(a) shows a few example volumes of the initial reordering strategy that was employed during the development of the sequence. In this more rudimentary version of the sequence the segments were placed as close as possible to their ideal location according to Eq. 6.2 regardless of where that nearest location in k -space would

be. At the reconstruction stage the segments that would lie more than 30ms away from their ideal location would be identified as corrupt and replaced by GRAPPA. This however could cause clustering of the segments that needed to be replaced. Replacing the segments with this type of reordering scheme is thus analogous to reconstructing a randomly sub-sampled data set (i.e., with a locally varying acceleration). We therefore attempted a variable GRAPPA kernel approach. In this approach the support of the GRAPPA kernel was varied based on the local availability of source points (non-corrupted segments) in order to include as many points as possible. However, this approach still resulted in a sub-optimal reconstruction for the volumes, in which the corrupted segments were closely spaced (as seen by the noise enhancement in the background in Volumes 1 and 2 in Fig. 6.3). The sequence reordering was therefore optimized by including a decision criterion that tells the sequence to use a second lookup table for choosing a segment to acquire if the current segment cannot be acquired within 30 ms from its ideal location. The lookup table only has entries for even lines on the even k_z -planes, and uneven lines on the uneven k_z -planes, which forms a checkerboard-like pattern. This ensures that all the corrupted data in k -space are placed in a checkerboard-like pattern in the k_y - k_z -plane and thus always surrounded by uncorrupted data points in both dimensions. Fig. 6.3(b) shows the lookup table and a few examples of the distribution of segments in k -space for data acquired with this extended functionality. It is shown that the corrupted data are always surrounded by uncorrupted data points in both the k_y - and k_z -dimension. This aids the GRAPPA reconstruction (for reconstruction examples also see Fig. 6.4). Using this synchronization strategy has the additional benefit that a fixed GRAPPA kernel can be used for each volume (i.e., the same kernel can be used for all locations in k -space), which simplifies the implementation of the reconstruction.

6.5.2 Optimization of the Image Reconstruction

The properties of the GRAPPA kernel, such as kernel size and dimensionality, can have a distinct impact on the quality of the reconstructed images [45, 125]. In 2D acquired data the reconstruction is performed on a slice-by-slice basis. In this scenario, it is theoretically sufficient to use a 1D kernel with points defined along the subsampled dimension, but generally a 2D kernel is used, which also contains a few points along the other dimension for interpolation. In 3D k -space data we have an extra dimension at our disposal when we construct our GRAPPA kernel. This allows us to take advantage of the coil sensitivity variations in two dimensions and still use the points along the third dimension for interpolation [126]. However, arbitrarily increasing the number of kernel points may result in a less accurate estimation of the kernel weights depending on the size of the Auto-Calibration Scan (ACS) data [45]. The ultimate performance of the reconstruction will therefore be a trade-off between the extra information provided by including additional data and the conditioning of the kernel estimation.

To optimize the GRAPPA reconstruction we assessed the performance of two 2D kernels and one 3D-kernel. The three GRAPPA reconstructions were run on the non-synchronized SPGR and bSSFP data with scan parameters as described Section 6.6. The 12-channel head coil that was used in these experiments has coil sensitivity variations will therefore be along the AP and RL-direction. Due to its geometry there is no variation in coil sensitivity along the SI direction. The data sets were therefore acquired in the following orientation:

- k_x = Readout direction = SI
- k_y = EPI blip direction = RL
- k_z = Partition direction = AP

The k -space data was fully sampled during the scan, but were subsampled for the

GRAPPA reconstruction. Before subsampling the centre $15 \times 15 \times 20$ ($k_x \times k_y \times k_z$) portion of the first 10 volumes was taken to serve as the ACS scan. After that the data were subsampled by a factor of 2 in the same checkerboard pattern as shown in Fig. 6.3(b), as these are the positions at which corrupted lines would appear in the synchronized data sets. The three kernels tested were: 1) the “kxky-kernel”; a standard 3×7 2D kernel in k_x - k_y , 2) the “kykz-kernel”; a 7×3 2D kernel in k_y - k_z with source points in a checkerboard pattern as described in the methods section, and 3) the “3D-kernel”; a $3 \times 7 \times 3$ 3D kernel in k_x - k_y - k_z , which is the extended version of the kykz-kernel and includes three adjacent points along k_x .

The reconstructed images were compared to the fully-sampled reference image and the Mean Square Error (MSE) was calculated by the sum of squares of the intensity difference between the GRAPPA image and the fully sampled reference image divided by the sum of squares of the reference image over the entire field of view:

$$MSE = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N |x_{i,j} - x_{i,j}^{FS}|^2 \quad (6.3)$$

where $x_{i,j}$ is the pixel from the GRAPPA reconstructed image and $x_{i,j}^{FS}$ is the pixel from the fully-sampled reference image. $M \times N$ is the total number of pixels in the image.

As Fig. 6.4 shows the best reconstruction is achieved with the 3D kernel. By inspection of the difference maps it is seen that the 3D kernel is able to remove most of the remaining aliasing seen in the difference images for the kykz-kernel (most easily seen in the bSSFP data). Quantitatively the 3D kernel also exhibits the smallest MSE for both data sets. The improved reconstruction quality does come with a time penalty, though. The average total times for kernel estimation

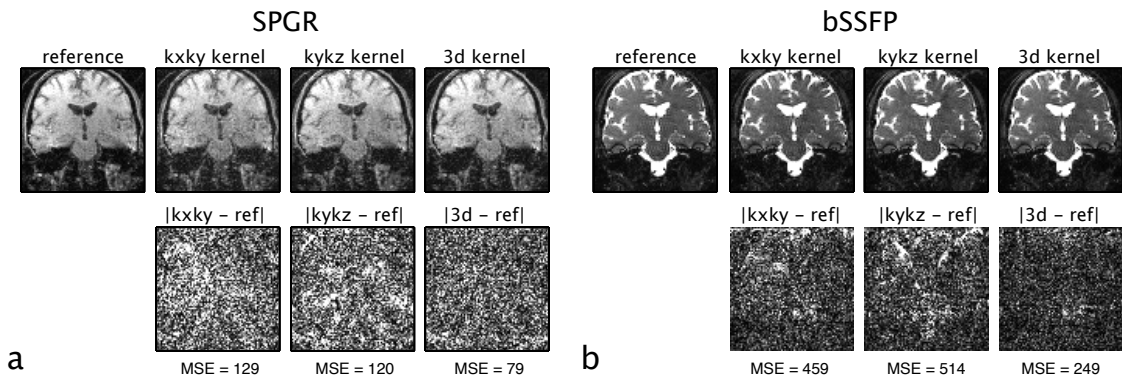


Figure 6.4: GRAPPA reconstruction results using three different kernels with SPGR (a) and bSSFP (b) data. The top images show the fully-sampled reference data along with the GRAPPA reconstructed images. The bottom row shows the absolute difference images between the reference image and the GRAPPA images. The windowing for the difference images was set to $1/10^{th}$ of the corresponding reference image. For both sequences the best reconstruction is achieved using the 3D kernel.

were 5.1s, 3.7s, and 100.0s for the kxky-, the kykz, and the 3d-kernel, respectively. For reconstruction (i.e., applying the estimated weights to the undersampled data in all time-points) the durations were 71s, 63s, and 198s. For this study the reconstruction was done off-line and there was no time constraint. We therefore chose to reconstruct the images using the 3d-kernel to achieve the best possible reconstructions. When online reconstruction is desired, however, this might need to be reconsidered.

6.5.3 Trigger Delay

The cardiac cycle consists of two distinct phases: systole, the time during which the heart's ventricles contract, and diastole, the quiescent phase of the cardiac cycle [59]. Due to the large pressure wave associated with systole [127, 128], this period is generally considered as the most unstable period in MR imaging. For the effectiveness of our method it is important that we synchronize the acquisition in such a way that unstable systole data are indeed confined to the edge of k -space and not placed at the centre. This will benefit the stability of Methods 1–3, and

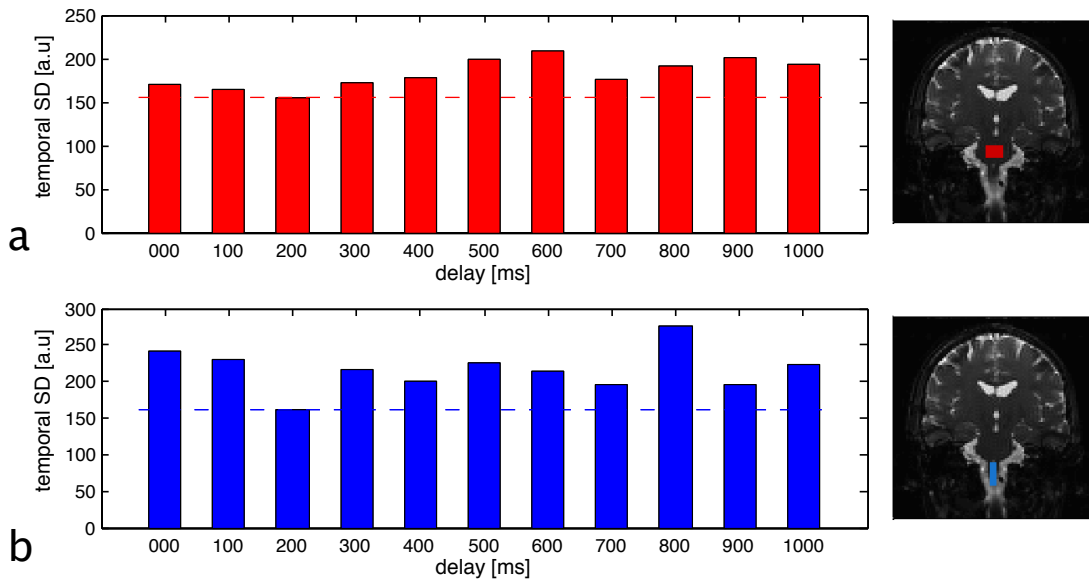


Figure 6.5: ROI results for the pons (a) and the medulla (b) comparing the temporal standard deviation (in scanner units) between data sets acquired with different trigger-acquisition delays. The delay is denoted in milliseconds. 11 delays were tested ranging from 0–1000 ms (D100–D1000). For both ROIs the 200 ms delay is the most stable.

is a requirement for Method 4. This means that we have to determine the timing delay between the trigger recorded by the pulse-oximeter and the arrival of the cardiac pulsations at the location of the brainstem.

In this pilot study, 11 time-series acquisitions with scan parameters similar to the parameters presented for the optimized method, but with varying trigger-acquisition delays, were acquired in a single subject. The delays ranged from 0 to 1000ms with 100ms increments between the scans. Note that “delay” in this context does not refer to a period of waiting as it generally does in cardiac gated studies. The delay is implemented into the sequence by simply adding the delay value to t_{trig} . The result is that the modulation across k -space will be shifted in the positive k_z -direction. For some delays the centre of k -space will be collected during systole and we expect these data sets to show the largest variability, whereas datasets in which the systole data are collected at the edge of k -space will be less affected. As shown in Fig. 6.5 ROI analysis of the pons and medulla showed the

lowest temporal variability when a delay of 200ms was used. This is consistent with the combined results presented by Gupta *et al.* [124] and Shkumat *et al.* [128]. Gupta *et al.* reported peak CSF flow changes in the pontine cistern approximately 200 ms after ECG trigger detection, whereas Shkumat *et al.* demonstrated a delay of 400 ms between ECG and pulse-oximeter trigger detection. In addition this paper showed that the trigger accuracy of pulse-oximeter measurements was comparable to ECG triggering.

6.6 Assessment of Cardiac Synchronization in 3D SPGR and bSSFP

The developed synchronization method was tested with 3D-bSSFP and 3D-SPGR acquisitions in four healthy volunteers. In each subject four resting scans were acquired: 1) non-synchronized bSSFP, 2) cardiac synchronized bSSFP, 3) non-synchronized SPGR, and 4) cardiac synchronized SPGR. The non-synchronized readout was acquired using the standard acquisition scheme mentioned above (acquiring all k_y -segments in a given k_z -plane before moving on to the next k_z -plane). The cardiac synchronized data were acquired with the final (optimized) version of the sequence and reconstructed following each of the three methods described in Fig. 6.2 (i.e., synch-only, synch+GRAPPA, and synch+GRAPPA+PF). So in total, eight reconstructed data sets were generated for each subject (four bSSFP and four SPGR data sets).

The scan parameters for all sequences were as follows: $\alpha=30^\circ$, TR/TE=12/6 ms, FOV=192×192×48 mm, Matrix=96×96×24, Bandwidth=1860 Hz/pixel, 8 lines per TR, $T_{vol}=3.5s$, 60 volumes. The shim volume was placed to cover the pons and medulla. The SPGR acquisitions used RF spoiling (quadratic phase increment=117°) and gradient spoiling (2π phase per voxel). For the bSSFP acquisitions the RF in-

crement was set to place as much of the brainstem as possible in the passband of the bSSFP signal profile. To determine the readout order in real-time, cardiac triggers were acquired using a pulse oximeter and fed back into the sequence. The heart rates across scans for each of the subjects were: 1.05 ± 0.04 Hz, 0.88 ± 0.05 Hz, 1.14 ± 0.02 Hz, and 1.19 ± 0.06 Hz (mean \pm sd). Standard T_1 weighted structural scans were also acquired in each subject.

Image reconstruction was performed off-line using custom in-house code in Matlab (Mathworks, Natick, MA). The data acquired with the conventional readout order were phase corrected and Fourier transformed into image space, after which a sum-of-squares algorithm was used to combine the individual coil images. For the GRAPPA reconstruction a $3\times 7\times 3$ GRAPPA kernel was defined with 30 source points distributed in a checkerboard pattern in three adjacent k_y - k_z -planes. The kernel weights were calculated from auto-calibration data comprising the $24\times 24\times 12$ points at the centre of k -space taken from 5 adjacent volumes in time (to improve SNR of the kernel estimation). A 5/6 Partial Fourier reconstruction was performed to discard systole data at the edge of k -space (replacing 4 out of 20 k_z -planes, which corresponds to approximately 150ms in the cardiac cycle). Partial Fourier reconstruction often simply zero-fills missing lines. This causes image blurring that would bias our SNR comparisons. A projection onto convex sets (POCS) algorithm [129] was therefore implemented to perform the partial Fourier reconstruction. The POCS algorithm is an iterative approach to synthesize missing lines, and thus does not introduce blurring. More information on the implemented POCS algorithm that was implemented can be found in Section 2.5.2.

The correction methods were assessed by calculating the temporal SNR (tSNR) within four atlas-based ROIs: 1) brainstem, 2) hippocampus, 3) thalamus, and 4) cortical grey matter. After motion correction, each of the time-series was reg-

istered to the communal midpoint space with FLIRT [90]. All further analyses were performed in this midpoint space to match the effects of interpolation during registration in all acquisitions. Additional transformations for the time series to anatomical scans, and anatomical scans to the MNI152 brain atlas template were also determined and used to register the atlas ROIs to the time series midpoint space. The Harvard-Oxford probabilistic atlas for subcortical structures, as distributed by FSL [88, 130], was used to define the subcortical masks. The probability maps were thresholded to 0.95 to define a conservative mask for each structure (i.e., only voxels with a probability greater than 95% were included). The grey matter mask was generated by segmentation of the structural image with FAST [91]. The grey matter mask was defined by thresholding the grey matter probability map to 0.85. Voxels with a probability $>10\%$ of being CSF were excluded from the mask to avoid voxels that contained CSF.

6.6.1 Results

Fig. 6.6 shows the temporal SNR in four ROIs averaged over four subjects. The data without cardiac synchronization are compared to the three reconstruction methods for the cardiac synchronized data. Inspection of the tSNR of non-synchronized data (Figs. 6.6(a)–(b)) shows that bSSFP generally has good temporal stability in the cerebral regions (tSNR > 30), which is considerably higher than the tSNR observed in the SPGR data (tSNR = 15-20). In the brainstem, however, the bSSFP data show a greatly reduced tSNR (comparable to SPGR) due to large, highly localized, signal fluctuations.

When comparing the three correction methods it is found that synch+GRAPPA consistently outperforms synch-only and synch+GRAPPA+PF in both the SPGR and bSSFP data. For each ROI synch+GRAPPA shows greater tSNR than the other two correction methods. The improved performance of the synch+GRAPPA

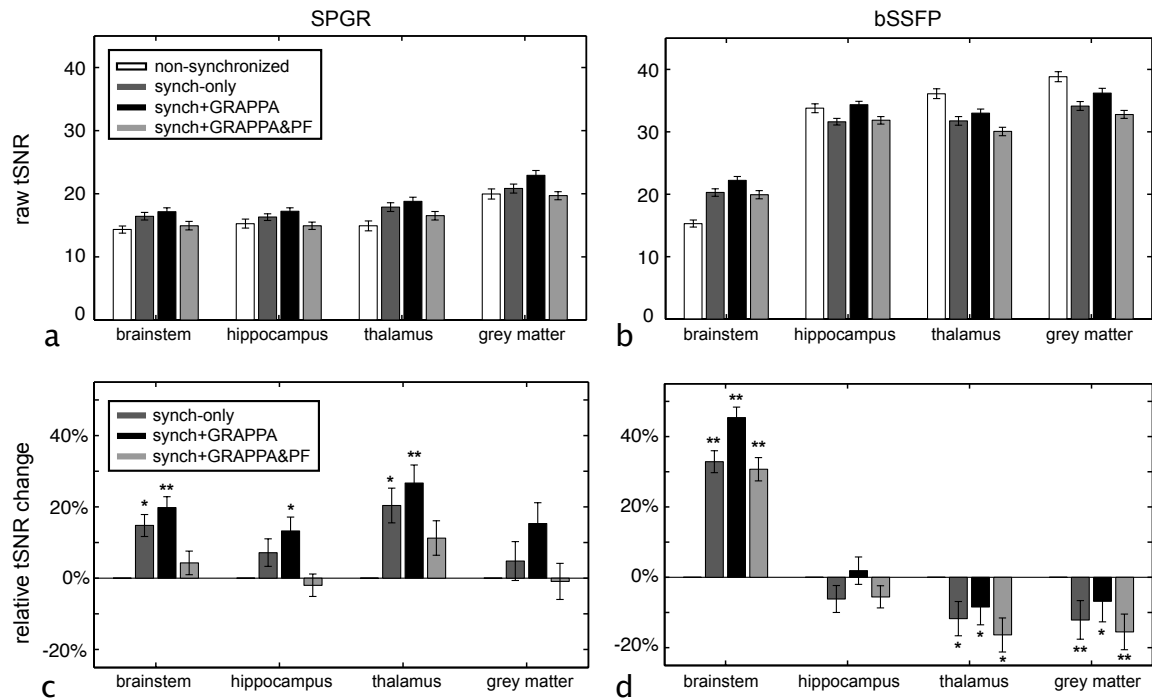


Figure 6.6: Results comparing synchronized 3D data reconstructed in three different ways (methods 1, 2 and 3) to the conventional (non-synchronized) acquisition. Panels (a) and (b) show the raw tSNR within each ROI averaged over four subjects. Panels (c) and (d) demonstrate the relative tSNR change of the correction methods versus the non-synchronized data. The results show that synch+GRAPPA consistently outperforms the other correction methods. For SPGR, synch+GRAPPA shows an improvement over all ROIs (15-25%). For bSSFP, a large improvement is achieved in the brainstem (45%) at the cost of a loss in tSNR in the thalamus and cortical grey matter (-10%). Significant tSNR differences are indicated by * ($p < 0.05$) or ** ($p < 0.01$); error bars denote the standard deviation between the subjects.

method over the synch-only method reflects the expected effect of removing the abrupt signal changes across k -space by replacing the corrupted segments with a GRAPPA reconstruction. When partial Fourier is used however (Method 3) a consistent reduction in tSNR of approximately 10% is observed compared to Method 2.

Fig. 6.6(c) and 6.6(d) show the relative tSNR change of the correction methods compared to the non-synchronized data. Focusing on synch+GRAPPA and ignoring the synch-only and synch+GRAPPA+PF methods for now it can be seen that synchronization improves the temporal stability significantly in the SPGR data for all three subcortical ROIs. The tSNR in the cortical grey matter for SPGR is

also increased with synchronization, although not significantly. bSSFP shows the greatest tSNR improvement of 45% in the brainstem when synchronization is performed (i.e., the tSNR is 1.45 times greater when synchronization is performed using synch+GRAPPA). No improvement was observed in the other ROIs. In the hippocampus no significant difference is seen, whereas the thalamus and the cortical grey matter show a reduction of approximately 10% compared to the non-synchronized sequence.

Fig. 6.7 and 6.8 show tSNR maps of the central five slices of non-corrected data and corrected data for SPGR and bSSFP, respectively. The results of two representative subjects are shown to demonstrate the reproducibility of the data. The top row shows the temporal mean for anatomical guidance, the centre row and the bottom row show the tSNR maps of the non-synchronized data and synchronized data (synch+GRAPPA), respectively. Zoomed insets of the central slice are shown on the far right. The tSNR maps of uncorrected SPGR data (Fig. 6.7) clearly show the spatial extent of the physiological fluctuations, which include the hippocampal regions and the thalamus. Note that the spatial extent of the reduced tSNR of these 3D data closely resembles the spatial pattern of the voxels that correlated with the cardiac waveforms in our 2D regression results (Fig. 5.1). When synchronization is performed the tSNR is most significantly improved in these subcortical regions resulting in a more homogenous tSNR map. The effect of synchronization is also seen by visual inspection of the raw time series.

Fig. 6.8 demonstrates that the physiological fluctuations in bSSFP are highly localized to the brainstem. At the same time, the brainstem shows the biggest improvement when synchronization is performed. The cortical regions of the brain show a small, spatially homogeneous, drop in tSNR, which reflects the 10% drop in tSNR in the grey matter shown in Fig. 6.6. In one subject spatially localized reductions of tSNR are visible in the white matter, which resemble ghosts of

the lateral ventricles (encircled). Subject C is the only subject in which this was observed.

6.7 Discussion

6.7.1 Characterization of Signal Instabilities

The signal instabilities in 3D data are predominantly localized to the inferior regions of the brain, which are known to be more prone to respiratory effects and cardiac pulsation [68, 66]. The results presented in this and the previous chapter suggest that the main contributor is cardiac pulsation (see Fig. 5.1). The signal instabilities in the bSSFP data are very structured and localized to the regions surrounding CSF (i.e., fluctuating CSF ghosts). This can be explained by the fact that the bSSFP signal is a summation of coherent transverse magnetisation formed by several preceding excitations. For a stable steady-state signal it is important that the phase coherence is preserved between consecutive RF pulses. Phase perturbations caused by spins moving through gradient fields can severely disrupt the steady-state signal [113, 111]. This is most pronounced in long T_2 species, such as CSF.

The signal instabilities in SPGR, on the other hand, are more spatially extended and include most of the subcortical structures. In SPGR, spoiling of the transverse magnetisation prevents any build up of coherence signal over multiple TRs, and thus the signal instabilities are less localized to CSF. Nevertheless, the magnetisation is sensitive to phase disruptions introduced by physiological fluctuations. The results presented here are in line with a previous study [131] in which RF spoiled and non-RF spoiled 3D EPI, which acquired each EPI plane in a single-shot (TR=65ms), were compared. Like the bSSFP data presented here, the instabilities in non-RF spoiled 3D EPI were most pronounced in tissues with




Figure 6.7: SPGR tSNR maps in two subjects showing the central five slices. A zoomed inset of slice 13, showing the brainstem and subcortical structures, is shown on the right. In both subjects a clear increase in tSNR is observed in the brainstem, the hippocampus, and thalamus, when synchronization is performed.




Figure 6.8: bSSFP tSNR maps in two subjects showing the central five slices. A zoomed inset of slice 13, showing the brainstem, is shown on the right. In both subjects a clear increase in tSNR is observed in the brainstem when synchronization is performed. However, in subject C some artifactual ghosting of the ventricles (encircled) was visible in the white matter in the segmentation direction (k_y).

long T_2 , such as CSF. The signal fluctuations in the subcortical structures were attributed to motion induced by nearby pulsatile arteries.

6.7.2 Current Implementation of the Correction Method

Synchronizing the readout to the cardiac cycle reduces the temporal variability in bSSFP and SPGR time-series acquisitions. Our results are in line with a previous method that synchronized a 3D spiral GRE sequence [110]. In that previous method, however, a portion of the acquisition had to be devoted to dummy scans, to ensure effective synchronization. When the number of dummy scans per volume is fixed (to ensure a fixed frame rate), one either has to base the number of dummy pulses on the worst-case scenario (which will significantly reduce the temporal resolution), or allow some segments to be acquired with the incorrect phase (and thus risk a reduction in the efficacy of artifact removal). In this chapter we have shown that GRAPPA can be employed to overcome this problem by replacing those segments that are not acquired with the correct cardiac phase. This allows data collection with the same temporal resolution as data acquired with a conventional (non-accelerated) readout.

Limitations

The effectiveness of cardiac synchronization depends on the regularity of the heart rate during the acquisition. As the cardiac phase needs to be determined prospectively, the current duration of the cardiac cycle (T_{RR}) is estimated by the average period of 10 preceding cardiac cycles. An incorrect estimation of the cardiac phase at the start of a segment acquisition causes the current segment to be assigned to the wrong location in k -space, which may reduce the effectiveness of the method. This problem has been identified in the original COPE paper [108], and simulations showed that for normal heart rate variability image degradation due to

timing errors was negligible [108]. However, when imaging patient populations with increased heart rate variability the effectiveness of the method would have to be re-evaluated (e.g. in a separate pilot study). We foresee a potential method to minimize the effect of increased heart rate. At the end of each acquisition one could re-evaluate which segments were assigned to an incorrect cardiac phase by comparing the estimated cardiac phase of each segment at the time of acquisition with the true cardiac phase (once all the triggers have been recorded). The additional corrupted segments can then be included in the GRAPPA reconstruction.

One important factor to consider with 3D acquisitions in general is the inability to define a single time point at which the image was acquired. This can be important in fMRI data analysis if one wishes to precisely assess signal timings. Whereas a given slice in 2D single-shot EPI can be attributed to a narrow window of tens of milliseconds, 3D acquisitions are effectively blurred over several seconds. Previous work with conventional ordering has noted that the majority of signal comes from the central portion of k -space, and suggested that this time can be considered representative of the time that the 3D volume is acquired [131]. While such an interpretation should always be undertaken with caution, defining a time that is representative of the 3D volume is problematic for synchronized data, as the point at which central k -space is acquired is not predefined and may vary from one volume to the next.

Partial Fourier

In all four subjects synch+GRAPPA consistently outperformed the other two synchronization methods. The lack of improvement from partial Fourier rejection of systolic data may at first seem to contradict our simulation results (see Section 5.5). However, simulations produced full k -space volumes containing systole-free data, whereas our real-time synchronization method reduced the k -space data

by 5/6 using partial Fourier reconstruction. Due to the reduced amount of data available, a drop in image SNR of approximately 10% is expected for a 5/6 partial Fourier reconstruction. This amount is similar to the tSNR reduction we see in our data when partial Fourier is used. Any improved stability gained by rejecting systole data would appear to be outweighed by the increase in thermal noise inherent to the loss of data in the partial Fourier reconstruction. It has been demonstrated previously that the SNR efficiency in 3D acquisitions can be improved when partial Fourier is employed to speed up the volume acquisition rate [132]. Here, partial Fourier fulfils the task of replacing corrupted data and such a speed up is not obtained.

Synchronization Induced Ghosting

One subject, Subject C, exhibited some degree of ghosting in the left-right direction in a number of volumes, which resulted in localized regions of reduced tSNR lateral to the ventricles when synchronization was performed (black circle in Fig. 6.8). The direction in which the ghosts are visible corresponds with the direction in which the EPI planes are segmented and the location of the ghosts is determined by the segmentation factor. A possible explanation for this effect could be that large jumps along the k_z -direction from one TR to the next increase the sensitivity to flow induced instabilities for spins that are flowing in the AP (k_z) direction [113]. With synchronization large jumps along k_z are often made near the end of the acquisition of a volume to collect the last few segments. This rapid switching of the gradients can lead to large $m_{1,z}$ differences, and thus flow-induced phase variations between TRs. Respiration related pulsations could be another possibility of the increased signal fluctuations. Reordering of the readout with respect to the cardiac cycle might therefore scramble the k -space modulation caused by respiration effects. Our 2D regression results (Fig. 5.1), however,




Figure 6.9: Images showing the normalised tSD ($\sigma(s)/S$) for synchronized and non-synchronized bSSFP. The results are shown for a short TR and long TR acquisition to assess the effects of Eddy currents. When the data are synchronized some of the instabilities within the dark bSSFP band are aliased into the background. This effect however is observed for both data sets, which indicates that the effect is not caused by Eddy currents.

do not seem to support this, so further investigation is needed. If the instabilities are indeed respiratory related, one way to compensate for respiratory effects would be to add 1D or 2D navigators to the sequence [133, 106], which compensate for respiration induced B_0 -field changes [58], or use retrospective correction techniques [65, 71]. Both methods are expected to be fully compatible with the synchronization technique presented here, and are topics of further investigations.

Eddy Currents

Similar to pulsatile flow, time-varying eddy currents can cause disruptions of the steady-state signal in balanced SSFP, resulting in signal instability. It has been previously shown that the large jumps in k -space associated with non-linear encoding schemes, such as centric or random orderings, can create such rapidly changing eddy currents, creating image artifacts [134]. As shown by Fig. 5.2 the G_x and G_y gradient waveforms are highly similar for each of the segments (the G_x remains constant for every TR, whereas G_y only shows small variations in the

pre- and de-phaser blips). The G_z gradients on the other hand span a much larger range of gradient strengths ($-9mT/m$ to $13mT/m$) in order to encode for the different planes in k_z . The large jumps along the k_z direction in k -space might therefore generate eddy-current induced artifacts in our cardiac synchronized readout.

To test the effect of time-varying eddy currents in our readout, phantom experiments were performed and analysed for temporal stability. Synchronized¹ and non-synchronized data were acquired with varying TE and TR. For the short-TR data set (TR=13,TE=5.9) the time between the pre-phaser and the readout, and the time between the de-phaser and the subsequent RF-pulse was set as short as possible (< 1 ms) to maximise the sensitivity to Eddy currents. In the long-TR data set (TR=24,TE=12), the dead times before and after the readout gradients was 6 ms. If eddy currents are introducing instabilities then we expect the instabilities to be more pronounced in the short-TR data compared to the long-TR data. The normalised tSD (normalised against the mean signal within the phantom) is shown for both data sets. The non-synchronized data show the largest instabilities around the dark bSSFP bands. Visual inspection of the time series showed that this was due to a slow drift of the bands (caused by gradient heating) over the course of the experiment. At short TR we see that to a certain degree these instabilities are aliased into the background when the data are cardiac synchronized (Fig. 6.9). The magnitude of the aliased instabilities, however, was small ($< 1.5\%$ of the mean signal). A reordering technique, in which consecutive phase encoding steps are paired, has previously been suggested to minimize the effects of eddy-currents [134]. In this technique adjacent phase encode lines are grouped in pairs, to be acquired consecutively, such that the phase error produced by the first phase encode step is immediately cancelled by the preceding phase encode step. Our synchronization method can be extended to use the pairing technique

¹To perform cardiac synchronization a cardiac trace (period = 900 ms) was simulated using a physiological measurement simulator on the scanner

by ensuring that each time an even number of segments is acquired per k_z -plane, before moving on to a different location along k_z . It is expected that this might also reduce the effects of the non-zero first-order gradient moments in k_z [113]. However, this has not been implemented at this stage.

6.8 Conclusion

in this chapter we have demonstrated the effectiveness of a cardiac synchronized readout in reducing the temporal fluctuations caused by cardiac pulsatility. In data acquired with a conventional readout order, the periodicity of flow leads to periodic errors in k -space, which subsequently results in ghosting of blood and CSF into adjacent tissue. In order to improve the temporal stability of 3D sequences, the order of the readout was synchronized in real-time to the cardiac cycle. Fixed volume frame rate was ensured by combining the synchronization technique with a customized parallel-imaging reconstruction. This method significantly improved the temporal stability in areas that were affected by cardiac-related signal fluctuations. In the brainstem the tSNR was increased by 45% in the bSSFP data, although a small reduction in cortical areas (which generally have high tSNR) was detected. For SPGR all subcortical structures showed improved temporal stability when synchronization was performed. This method thus offers improved temporal stability without the costs of additional scan time and may therefore open the possibility to use of segmented 3D readouts for fMRI of subcortical structures, in order to mitigate the effects of distortion and signal dropout.

Chapter 7

Retrospective Correction Strategies

7.1 Introduction

The aim of this chapter is to consider both prospective and retrospective corrections in order to optimize the temporal stability in the brainstem. Specifically, each method will be assessed independently as well as their combined use.

Retrospective corrections methods that use independent physiological measures to model the physiological fluctuations have been shown to reduce physiological fluctuations significantly [71, 65, 76, 77]. Initially this approach was proposed as a k -space based correction. Hu *et al.* [65] argued that k -space is the most natural domain to correct for physiological fluctuations as this is the domain in which the MR signal natively is acquired. Glover *et al.* [71], however, showed that when the correction is performed in image space instead, better performance was achieved with the added benefit of a simplified correction (i.e., as the correction is performed in image space no off-line reconstruction of the raw data is needed). For this reason the majority of the research that has been conducted following on this has mainly focussed on the RETROICOR method.

However, RETROICOR assumes that each image is acquired instantaneously

so that a unique cardiac or respiratory phase can be assigned to each image. Although this assumption is justified for 2D imaging methods (a single slice is acquired in the order of milliseconds), it may not for 3D imaging methods, where the data that are combined to a single image is often acquired over a period of a few seconds. Nevertheless, several groups have started applying RETROICOR to 3D data [135, 136, 137]. The assumption that is made is that the majority of the contrast is created when the centre of k -space is acquired and thus the cardiac and respiratory phase at that time point are assigned to the entire volume. Although significant improvements using RETROICOR are reported, RETROKCOR surprisingly has not been considered.

We suspect that, although RETROICOR has initially been shown to be superior to RETROKCOR for the correction in 2D acquisitions [71], this might not be the case in 3D imaging. The fact that each k -space segment can be assigned with an individual cardiac and respiratory phase is likely to be beneficial. Assigning a single phase to the entire volume will particularly be problematic for data acquired with cardiac synchronization, because the centre of k -space is not necessarily acquired in consecutive TRs. Depending on the cardiac cycle the segments that lie close to the centre of k -space may be acquired in batches and not in a subsequent order as in a conventional 3D acquisition.

Apart from the question in which domain to perform the regressions, a second important consideration that has to be assessed is which regressors to include in our physiological model. Many papers have been published in which new regressors [76, 138, 77, 139] and optimal regressor sets [8, 95, 80] have been proposed. However, nearly all this work has been focussed on 2D GRE-EPI. It is likely that different sequences, with different contrast mechanisms, will require individually optimized regressor sets.

This chapter therefore discusses a range of experiments that were performed to

optimize retrospective correction methods for 3D bSSFP and SPGR that target the brainstem. Our primary aim is to refine the RETROICOR technique by optimizing the set of physiological regressors, but the performance of RETROICOR on 3D imaging data is also assessed. Throughout this chapter optimizations are also provided for 2D GRE-EPI, the “gold standard” technique for FMRI. Ultimately the temporal stability of all three optimized sequences is compared and their potential for brainstem FMRI is discussed.

7.1.1 Chapter Overview

First, RETROICOR and RETROKCOR correction techniques are compared in simulations. This provided a testbed for the custom software developed for these techniques, and enabled us to study the performance of both techniques over a range of settings (e.g., image SNR). These experiments were intended to provide a better insight into the practical differences between the techniques.

Second, the implementation of the RETROKCOR technique was optimized by altering the way the corrections were performed. It was found that more robust results were obtained when the corrections are performed on the real and imaginary parts of the signal instead of the phase and magnitude (as originally proposed in the literature).

Third, a model selection procedure is introduced, which selects the optimal set of regressors needed to explain as much of the variance as possible without over-fitting the data.

Finally, the optimized models were applied to the 3D data and the achieved temporal stability was assessed and compared to the temporal stability in GRE-EPI.

7.2 Data Acquisition Methods

7.2.1 Data acquisition

The investigations presented in this chapter were performed on a range of data acquired with three different sequences (GRE, bSSFP, and SPGR), and a range of different volume acquisition times (i.e., temporal resolution). Because the various investigations contain some overlap in terms of the data they used (i.e., some data sets were used in multiple investigations, and some investigations used multiple types of data), an overview of the data acquisition is given here first, in order to avoid unnecessary repetition. We adopt a concise naming convention for the different data sets, which we will use throughout the remainder of the chapter. All data sets were acquired on a 3.0 Tesla Siemens TIM Trio system (Siemens Healthcare, Erlangen, Germany) using a 12-channel head receive coil.

Data Set 1: 2D Critically Sampled Data (“2DCRIT-data”)

The single-slice rapidly acquired SPGR data as shown in Section 5.2 served as test data during the initial stages of development. These data enable a comparison of the two correction methods (RETROICOR vs. RETROKCOR) in the case that physiological fluctuations are critically sampled and the single-slice acquisition is effectively instantaneous in both k -space and the image domain. For reference, the scan parameters used for that study were: $\alpha = 30^\circ$, TR/TE = 12/6 ms, FOV = $192 \times 192 \times 2.5$ mm, Matrix = 96×96 , Bandwidth = 1860 Hz/pixel, 8 lines per TR, $T_{\text{vol}} = 152$ ms, 1500 volumes with the shim volume targeted to the brainstem. RF spoiling (quadratic phase increment = 117°) and gradient spoiling (2π phase per voxel) were performed. Physiological monitoring was performed as described below.

Note that, although the acquisition was segmented, the short volume scan time

(152 ms) compared to physiological fluctuations makes it a reasonable approximation to instantaneous (non-segmented) data.

This data set was used to:

- provide realistic time courses for our computer simulations
- compare RETROICOR and RETROKCOR on instantaneous *in vivo* data

Data Set 2: 3D Data with long T_{vol} (“3DLONG-data”)

This data set consists of a series of synchronized and non-synchronized 3D acquisitions acquired along with conventional GRE-EPI data. Resting data were acquired in four healthy volunteers. For each subject five resting scans were acquired in a single scan session: 1) cardiac synchronized bSSFP, 2) non-synchronized bSSFP, 3) cardiac synchronized SPGR, 4) non-synchronized SPGR, and 5) conventional multi-slice GRE-EPI. The scan parameters of the 3D acquisitions and the 2D GRE-EPI are summarised in Table 7.1. For each scan physiological monitoring was performed as described below.

This data set was used to:

- Investigate the effect of cardiac synchronization on the model selection procedure
- Assess the efficacy of RETROKCOR and RETROICOR in 3D data

Data Set 3: 3D Data with short T_{vol} (“3DSHORT-data”)

Additionally, functional data were acquired using synchronized bSSFP and SPGR, and conventional single-shot GRE-EPI. The data were acquired in five healthy subjects. A painless vibrotactile stimulus paradigm was used, which is known to

<i>Parameter</i>	<i>bSSFP</i>	<i>SPGR</i>	<i>GRE-EPI</i>
RF Spoiling	N/A	117° inc.	N/A
Grad. Spoiling	N/A	2 π phase	N/A
Excitation Flip Angle	30°	30°	90°
TR	12 ms	12 ms	3500 ms
TE	6 ms	6 ms	30 ms
Segmentation	8 lines per TR	8 lines per TR	48 lines per TR [†]
Tvol	3500 ms	3500 ms	3500 ms
Imaging Plane	Coronal	Coronal	Coronal
EPI-blip Direction	Right-Left	Right-Left	Right-Left
Field Of View	192×192×48 mm	192×192×48 mm	192×192×48 mm
Matrix Size	96×96×24	96×96×24	96×96×24
Resolution	2×2×2 mm	2×2×2 mm	2×2×2 mm
Parallel Acceleration	None	None	GRAPPA factor 2
Partial Fourier	None	None	None
Readout Bandwidth	1860 Hz/pix	1860 Hz/pix	1860 Hz/pix
Number of Volumes	60	60	60
Total Imaging Time	210 s	210 s	210 s

Table 7.1: Sequence parameters for the “3DLONG-data”. [†]The single-shot GRE-EPI data were acquired with an acceleration factor of 2 in order to reduce the readout length and therefore distortion.

activate the gracile or cuneate nuclei in the inferior part of the medulla [140, 141]. For the purposes of this chapter, however, we will only consider the temporal stability of the acquired time-series. In each subject three functional data sets were collected (synchronized 3D bSSFP, synchronized 3D SPGR, and 2D GRE-EPI data). All data were acquired coronally and localised to the brainstem. In order to prevent bSSFP banding in the area of interest, the shim volume was reduced and carefully placed to cover the brainstem. The scan parameters are listed in table 7.2.

Additional whole brain images were collected for registration purposes using the GRE-EPI, bSSFP, and SPGR sequences (i.e., a single time point). The scan parameters and orientation of the slices were identical to the time-series scans with the only difference being that the number of slices was increased to 80 to allow whole brain coverage.

This data set was used for:

- Optimizing the set of physiological regressors

<i>Parameter</i>	<i>bSSFP</i>	<i>SPGR</i>	<i>GRE-EPI</i>
Synchronization	YES	YES	N/A
RF Spoiling	N/A	117° inc.	N/A
Grad. Spoiling	N/A	2 π phase	N/A
Excitation Flip Angle	30°	30°	77°
TR	12 ms	25 ms	2000 ms
TE	6 ms	12 ms	30 ms
Segmentation	10 lines per TR	20 lines per TR	35 lines per TR [†]
Tvol	1932 ms	2025 ms	2000 ms
Imaging Plane	Coronal	Coronal	Coronal
EPI-blip Direction	Right-Left	Right-Left	Right-Left
Field Of View	180×180×40 mm	180×180×40 mm	180×180×40 mm
Matrix Size	80×80×20	80×80×20	80×80×20
Resolution	2.3×2.3×2 mm	2.3×2.3×2 mm	2.3×2.3×2 mm
Parallel Acceleration	None	None	GRAPPA factor 2
Partial Fourier	None	None	7/8
Readout Bandwidth	1690 Hz/pix	1136 Hz/pix	1358 Hz/pix
Echo Spacing	.78 ms	1.0 ms	.84 ms
Number of Volumes	329	329	329
Total Imaging Time	656 s	664 s	634 s

Table 7.2: Sequence parameters for the “3DSHORT-data”. [†]The single-shot GRE-EPI data were acquired with an acceleration factor of 2 in order to reduce the readout length and therefore distortion.

7.2.2 Physiological Monitoring

The subjects’ cardiac cycles were recorded using two pulse oximeters. The Siemens pulse oximeter was used for cardiac synchronization and a separate pulse oximeter was fed into the Biopac[®] acquisition system to record the cardiac trace. This system was used to simultaneously record the scanner triggers, the cardiac trace, and respiratory waveforms. The pulse-oximeter outputs a trigger each time an R wave is detected. Respiration was measured with a pneumatic belt, which was positioned at the level of the solar plexus. The pneumatic belt measures the positional change that the chest undergoes during each respiratory cycle. The respiratory trace is a continuous waveform, which reflects the relative measure of the depth of respiration. The amount of air inspired with each breath is proportional to the difference between the signal levels at full inspiration and expiration in each cycle. The scanner triggers were recorded simultaneously to accurately time-lock the physiological traces with the time-series data once the data have

been collected. The sequences were programmed to output a single trigger per volume (i.e., coincident with the first RF pulse).

The cardiac phase (φ_c) at any given instance in time was defined as the fraction in the RR interval (the interval between two consecutive R-wave triggers) identical to the definition used for the cardiac synchronization method described in the previous chapter. The respiratory phase (φ_r) was computed using a histogram-equalised transfer function as first introduced by Glover *et al.* [71]. The histogram transfer function takes into account the depth of inhalation and distinguishes between inhale and exhale (for more details the reader is referred to Appendix A).

7.2.3 Definition of the Regressors

The efficacy of correction methods such as RETROICOR and RETROKCOR relies on how well the confound regressors describe the physiological fluctuations in the MR signal. The regressors explored in this study can be grouped into three types: 1) frequency terms, 2) interaction terms, and 3) 'rate' regressors. In this section a brief description of the regressors is given along with their physiological meaning. All regressors were generated using custom in-house code in Matlab (MathWorks, Natick, MA). For a more detailed description on the procedure including mathematical descriptions of the regressors the reader is referred to Appendix A.

The frequency regressors are simple low-order Fourier series that model the independent effects of cardiac and respiratory fluctuations. The frequency regressors are generated by taking the sine and cosine values of the principle frequency and its higher harmonics. Based on findings by Harvey *et al.* [8] 14 frequency regressors and four interaction regressors were generated: for the cardiac signal we generated regressors for the principle frequency and the next two harmonics (i.e., φ_c , $2\varphi_c$, and $3\varphi_c$); the respiratory signal was modelled up to the 3rd harmonic

(i.e., φ_r , $2\varphi_r$, $3\varphi_r$, and $4\varphi_r$). The interaction terms describe the interactions between the cardiac and respiratory (for example the effect of respiratory-induced pressure changes on the cardiac rhythm) and are calculated by $\sin(\varphi_c \pm \varphi_r)$ and $\cos(\varphi_c \pm \varphi_r)$.

In addition to the regressors that were derived from the cardiac and respiratory phase, cardiac and respiratory rate regressors were also included. The respiratory rate has been shown to introduce low frequency variations in the BOLD signal through fluctuations in blood flow and arterial CO_2 [76, 75], which can effectively be modelled with a regressor that describes the respiration volume per time (RVT). Recently, the heart rate regressor was also found to be an effective confound regressor [79, 77, 78]. Although the physiological mechanism behind this regressor is not yet well understood, it has been hypothesized that BOLD signal changes that correlated with the heart rate could be due to arousal related neuronal activity [79]. Based on these findings both regressors were therefore included in our analysis. In order to allow for temporal shifts of the rate-regressors the temporal derivatives were also included.

In total 22 regressors were generated: 14 frequency terms, 4 interaction terms, 2 rate regressors and 2 rate derivatives. The generated regressors are listed in Table 7.3. The numbering and naming convention in this table will be used when the results are discussed.

7.2.4 Performing the Regressions; RETROICOR vs. RETROKCOR

In this section we highlight some of the practical differences between RETROICOR and RETROKCOR, and provide information on the implementation of these techniques. In principle, both techniques employ a similar method to extract physiological signals from the data. That is, both techniques use a linear regression with confound regressors that are generated based on external monitoring of the

<i>No.</i>	<i>Name</i>	<i>Description</i>
1	ev01_cardcos_01	cardiac: First order cosine
2	ev02_cardsin_01	cardiac: First order sine
3	ev03_cardcos_02	cardiac: Second order cosine
4	ev04_cardsin_02	cardiac: Second order sine
5	ev05_cardcos_03	cardiac: Third order cosine
6	ev06_cardsin_03	cardiac: Third order sine
7	ev07_respcos_01	Respiratory: First order cosine
8	ev08_respsin_01	Respiratory: First order sine
9	ev09_respcos_02	Respiratory: Second order cosine
10	ev10_respsin_02	Respiratory: Second order sine
11	ev11_respcos_03	Respiratory: Third order cosine
12	ev12_respsin_03	Respiratory: Third order sine
13	ev13_respcos_04	Respiratory: Fourth order cosine
14	ev14_respsin_04	Respiratory: Fourth order sine
15	ev15_cosadd	Interaction: (card + resp) cosine
16	ev16_cossb	Interaction: (card - resp) cosine
17	ev17_sinadd	Interaction: (card + resp) sine
18	ev18_sinsb	Interaction: (card - resp) sine
19	ev19_cr	Cardiac rate
20	ev20_dcr	Cardiac rate derivative
21	ev21_rvt	Respiratory rate
22	ev22_drvt	Respiratory rate derivative

Table 7.3: The full set of regressors that are tested in this study. The numbering and naming convention shown here will be used throughout the chapter.

respiratory and cardiac cycle. The practical differences between the techniques, however, arise due to the fact that the corrections are applied at different stages in the reconstruction pipeline.

RETROICOR on 2D multi-slice data

In RETROICOR the corrections are performed on the reconstructed (magnitude) images, i.e., at the end of the reconstruction pipeline. The regressors, or explanatory variables (EVs), are generated on a slice-by-slice basis, since each slice is acquired at a different time during the acquisition of a volume. For 2D data, one can reasonably make the approximation that the entire slice was acquired instantaneously (since the 40-50 ms required for excitation and readout is fast relative to physiological effects). This does require knowledge of slice ordering (in our case, A \rightarrow P interleaved), but the ordering will be the same for each collected volume.

This is in many ways the simplest method to implement the correction since most GLM software is written for global regressors.

RETROKCOR on 3D data

In RETROKCOR, corrections are performed in k -space, which we have accomplished through an off-line reconstruction. In our case, each segment can be reasonably considered to have been acquired instantaneously, but the entire k -space volume is acquired over several seconds. The corrections are therefore performed on a segment-by-segment basis, with a different set of regressors created for each k -space segment. We thus need to know the time at which each segment is acquired, and the resulting regressor set will consist of 4D point-wise EVs. Moreover, the order in which the segments are acquired is not fixed for each volume when cardiac synchronization is performed. The information about the ordering of the segments¹ was used to retrieve the time at which each segment was collected in order to generate the segment specific EVs. The regressions were then performed on the data of each individual channel, resulting in a total number of $N_c \times 2$ regressions, where N_c is the number of channels. After the retrospective corrections the GRAPPA algorithm as described in the previous chapter was performed to replace the segments that could not be collected at the correct cardiac phase prior to combining the data using a sum-of-squares combination.

RETROICOR on 3D data

The RETROICOR correction of 3D data is based on the assumption that most of the contrast is generated when the centre of k -space is collected. Thus, a single physiological phase is assigned to the entire volume, which corresponds to the time at which the centre of k -space is acquired. The only difference from 2D

¹The information on the ordering of the segments is stored in the data header. See Section 6.4

RETROICOR, therefore, is that regressors are common to entire volumes, rather than specific regressors being applied to different slices.

7.2.5 ROI Analysis

Time-series images were registered to the MNI template with FLIRT [90] using the whole brain acquisitions and the anatomical scan for each subject as intermediate steps. The transformations obtained from this registration were used to transform the brainstem ROI, which we defined in the MNI space, to the native spaces of each of the acquired scans. This way spatial interpolation of the data, which can bias the temporal variance was avoided. The Harvard-Oxford probabilistic atlas for subcortical structures [88, 130] was used to define the brainstem ROI by setting the threshold of the probability map to 0.95 (i.e., only voxels with a probability greater than 95% were included). A conservative mask was defined in order to restrict the optimization only to the area where the activation is expected. The regions of CSF surrounding the brainstem are known to demonstrate high signal fluctuations and are likely to dominate the optimization if included.

7.3 Method Development and Characterization

RETROICOR and RETROKCOR are mathematically equivalent as long as there is an infinite amount of SNR available. However, in the presence of noise, different physiological effects may be better corrected in one or the other domain, depending on which domain the effects are localized to a high-signal region. Cardiac effects are therefore expected to be better corrected in image space, whereas effects that are more spatially global (e.g., respiratory effects) might be better corrected for in k -space. Further, the fact that each k -space segment can be assigned with an individual cardiac and respiratory phase may play in favour for a k -space

based regression.

This section describes the development steps that were required to develop an optimized retrospective correction method for 3D bSSFP and SPGR. Our main emphasis will lie on RETROKCOR as theoretically this is the more obvious method to correct multi-shot 3D acquisitions. However, RETROICOR is also explored. Alongside the optimization for 3D acquisitions RETROICOR optimization for GRE-EPI will be performed (RETROICOR only).

7.3.1 Initial Characterization of RETROICOR and RETROKCOR

In order to obtain a thorough insight in the differences between the RETROICOR and RETROKCOR, we begin by comparing the performance of the two methods under conditions in which they should be equivalent. The two corrections are identical if both image and k -space can be considered to be instantaneously sampled, which we enforce in phantom simulations and we assume in rapidly-sampled 2D, single-slice data. The simulations allow exploration of various parameters such as SNR and “pCNR”, the contrast-to-noise ratio of the physiological fluctuations. The spatial extent of the physiological fluctuations was also varied. These experiments highlight important considerations for the RETROKCOR method. We then verify that similar results are found in 2DCRIT data. This section serves largely to develop, characterize and validate the methods in preparation for 3D data, where RETROICOR and RETROKCOR are *not* expected to perform equivalently.

Digital Phantom Simulations

Retrospective correction methods were tested on the Shepp-Logan digital phantom in Matlab (MathWorks, Natick, MA). The phantom simulated a single slice of data with matrix size 64×64 and 329 time points. Gaussian random noise was

added in image space to the real and imaginary channels of the data to simulate thermal noise. Noise levels were scaled to create three data sets with SNR values of 10, 25, 50, and 100. Physiological noise was subsequently added to each of the data sets in predefined locations as shown by the mask in Figure 7.1(a). We were particularly interested to test whether the size of regions affected by physiological noise impacted the ability to regress this noise out, since the signal power of effects that are highly-localized in image space are spread throughout k -space. Thus, our mask consisted of one large, contiguous region and a few isolated voxels.

The waveform used to simulate physiological noise was based on the first-order Fourier cosine term estimated from the respiratory trace in one of the 2DCRIT data sets (*ev07_respcos_01*). The amplitude was kept fixed between the data sets in such a way that pCNR values of 1, 2.5, 5, and 10 were obtained for the data sets with SNR values of 10, 25, 50 and 100, respectively. Example magnitude time-courses for a voxel with and without added physiological noise (SNR=100, pCNR=10) are given in Fig. 7.1(b).

For RETROKCOR, the images were first Fourier transformed to k -space, and the magnitude and (temporally unwrapped) phase channels were analysed using the *ev07_respcos_01* EV as the confound regressor. The magnitude and phase residuals were combined to form the corrected complex k -space data, and Fourier transformed back into image space. The RETROICOR correction used the same EV as for RETROKCOR, but was performed on the magnitude data of the image time-series only. The residuals of the regressions represent the corrected data and were used to compare the two methods.

The performance of the RETROICOR and RETROKCOR corrections were assessed on the amount of reduction in the standard deviation over time (tSD), which was calculated per voxel. The percent change in tSD after correction is

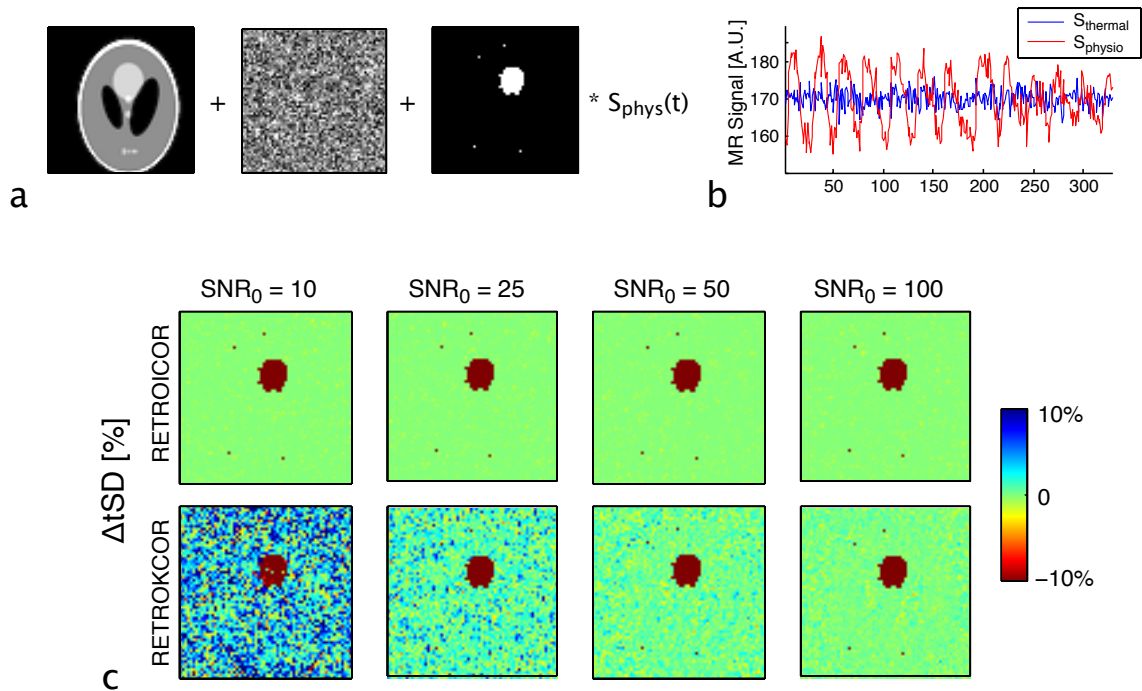


Figure 7.1: The simulation pipeline in (a) showing the addition of thermal noise (globally) and physiological fluctuation (within a given mask). Examples of the simulated time courses are shown in (b). Panel (c) shows the difference in tSD after RETROICOR (top row) and RETROKCOR (bottom row) correction. Red denotes improvement.

shown in Fig. 7.1(c). The RETROICOR and RETROKCOR regressions performed equally well in removing signal fluctuations in areas of simulated physiological noise. However, it was also seen that for RETROKCOR, instabilities were introduced in voxels that contained pure thermal noise. Although the effect was small for the high SNR data set (right column; SNR=100) it is pronounced when a more realistic SNR of 25 is simulated, where the tSD increases by up to 10%. RETROICOR, on the other hand, does not seem to affect areas without simulated physiological noise.

These initial results indicate that the original implementation of RETROKCOR (as proposed by Hu *et al.* [65]) performs worse than its image based counterpart, by introducing instabilities in regions that are not affected by physiological noise. It is observed that this effect is inversely proportional to the image SNR. In the remainder of this section, we will consider several potential improvements

to RETROKCOR that aim to address these problems. Figure 7.2 illustrates several modified RETROKCOR corrections that provide insight into the origin of the discrepancy between the two methods.

One of the differences between RETROICOR and RETROKCOR is the need to correct the phase information appropriately. In RETROICOR this is typically neglected, since most fMRI analyses only consider the image magnitude. In RETROKCOR, however, it is essential to correct the phase information accurately because errors in k -space phase can lead to profound artifacts in image magnitude. Panel (a), shows the change in tSNR when only the magnitude data are corrected in k -space (i.e., combining the corrected magnitude with uncorrected phase data). Several large regions show a reduced tSNR after the correction. In particular a cluster of tSNR decrease in the bottom part of the image is observed, which has a similar shape to the area of simulated physiological noise. Further, although not visible with the windowing used here, the primary cluster of physiological fluctuations had less improvement in tSNR compared to when both magnitude and phase are corrected. These two observations indicate that without phase correction, the physiological signal is mis-localized, which results in a sub-optimal correction in the region of interest and overcompensation in areas that are not affected.

Panel (b) shows the resulting difference in tSNR when both the magnitude and phase are corrected (i.e., the original implementation as shown in Fig. 7.1). Here, the reduced tSNR appears more noise-like (i.e, spatially uncorrelated), which may indicate that the phase inconsistencies are localised to outer k -space. One possibility is that the regression often fails in outer k -space where the low signal magnitude translates into high uncertainty on phase (since phase noise scales non-linearly with signal magnitude). We tested this hypothesis by masking the phase data such that the regressions are performed on all voxels apart from the

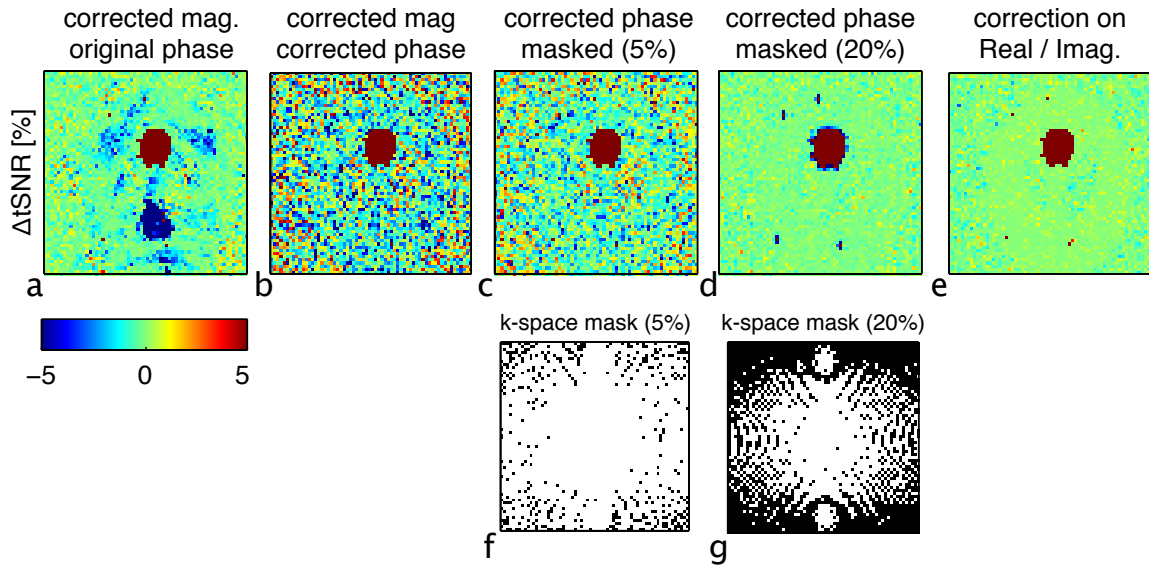


Figure 7.2: Illustration showing the difference in tSNR after RETROKCOR correction (red denotes improvement). Various modifications of the technique are shown in Panels(a)-(e). The k -space masks that were used in method c and d are shown in the panels (f) and (g), respectively. The figure shows that the best results are obtained when the corrections are performed on the real and imaginary channels. The simulation settings were: SNR=25, pCNR=2.5

voxels with the lowest magnitude. Panel (c) and (d) show the results when the bottom 5% and 20% voxels are excluded from the phase regression (the masks are shown in Panels (f) and (g)). The masks that were used are shown below the results panels. Excluding the lower magnitude voxels from the phase correction reduces the “noise” amplification considerably. In Panel (d) most of the voxels that only contain thermal noise no longer show a discrepancy in tSNR before and after correction. However, excluding the outer locations in k -space reduces the spatial resolution of the correction results in reduced tSNR in the voxels adjacent to the regions of true physiological noise (i.e., the blue rings around the red regions indicate reduced tSNR).

Finally we tested the performance when the corrections are performed on the real and imaginary data (R/I) instead of the magnitude and phase (M/P). The original paper by Hu *et al.* [65] argued that the regressions were to be performed on the magnitude and phase of the signal, as respiration was expected to mainly

have an effect on the phase of the signal. However, there is no real reason why the regressions cannot be performed on the real and imaginary parts of the signal instead. This avoids the non-linear scaling of phase noise and also gets around practical issues such as phase unwrapping [142] and the arbitrary nature of k -space masking, which could bias the results. Panel (e) shows the results when the correction is performed on the real and imaginary data. It is seen that the tSNR is no longer reduced in voxels that did not contain physiological noise.

2D *in vivo* Experiments

The phantom simulations suggest that corrections on real and imaginary (R/I) channels appear to have some clear advantages over the magnitude and phase (M/P) corrections. While providing a useful “ground truth” comparison, digital phantoms are not particularly realistic. In order to test if the observations are also present in real data, we performed similar corrections on 2DCRIT data. For this investigation we performed RETROICOR and RETROKCOR corrections on both R/I and M/P channels in a similar way as before. Each slice was considered to be acquired instantaneously, so that a single physiological phase could be assigned to the entire slice. The EVs were generated from the externally recorded cardiac and respiratory traces. Only the first order Fourier terms were considered at this stage, resulting in four EVs (two for cardiac, and two for respiratory).

The results of all four corrections are shown in Fig. 7.3. Panel (a) shows the reduction in temporal SD on the complex k -space data for both RETROKCOR regressions when the corrections are performed on M/P data (left) and R/I data (right). The M/P data shows that the improvement (red) is somewhat confined to the centre of k -space. Some k -space positions towards outer k -space even show an increase tSD after correction (blue). When the regressions are performed on the R/I data the improvement is much wider spread across k -space. The bottom

row shows the change in tSNR of the image magnitude data after RETROKCOR and RETROICOR corrections. This figure shows that for RETROKCOR, a clear improvement is obtained when the corrections are performed on R/I data. In RETROICOR the differences are non-existent. When the corrections are performed on the R/I channels RETROICOR and RETROKCOR are much more comparable, but not equivalent as small differences between the two methods are still observed. In areas where a tSNR increase was obtained RETROKCOR showed a slightly bigger improvement, whereas areas adjacent to this (red arrow) showed a marginal drop in tSNR (up to - 0.5 %).

Our phantom and 2D *in vivo* results confirm previous findings that for 2D corrections, due to the SNR distribution in k -space, RETROKCOR can introduce spatial correlations in the image domain [71]. When an additional variance is introduced to the phase data in outer k -space a thermal noise-like structure is imposed onto the image data (Fig. 7.1). This can be mitigated by confining the corrections to the centre (high SNR) region in k -space at the cost of reducing the spatial resolution of the corrections (Fig. 7.2). Improvements are achieved when the corrections are performed in the R/I channels instead, but small differences are still observed in *in vivo* data (Figs. 7.2 and 7.3).

We will next consider 3D k -space acquisitions, where the comparison between RETROKCOR and RETROICOR may differ. Based on the above results, all corrections in the remainder of this chapter will be performed on the R/I data channels.

7.3.2 Model Optimization; Implementation of the Method

An important factor in optimizing the retrospective correction is determining the specific regressors to include in the regression. In the initial papers by Hu *et al.* (RETROKCOR) [65] and Glover *et al.* (RETROICOR) [71] only the first two order terms (the primary frequency and the first harmonic) were included. Later,

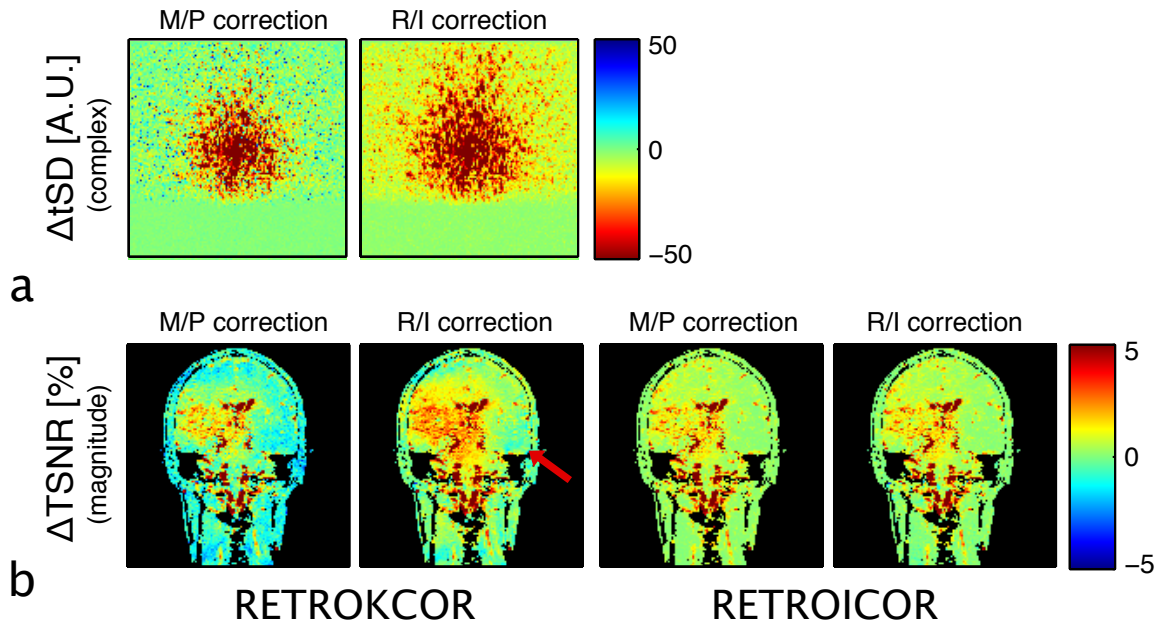


Figure 7.3: 2D *in vivo* results comparing RETROKCOR and RETROICOR. Panel (a) shows the change in complex standard deviation after correction (red is improvement). Panel (b) shows RETROKCOR and RETROICOR results when the regressions are performed on magnitude, and phase and real and imaginary channels. When RETROKCOR is performed on the real and imaginary channels the results are improved compared to the magnitude and phase correction. However, small reductions in tSNR can still be observed as denoted by the red arrow.

Harvey *et al.* [8] showed, that this is not sufficient in areas such as the brainstem. Additionally, several new types of regressors have been introduced since the methods were originally introduced, such as the interaction terms that describe the interaction between respiratory and cardiac [138], and the cardiac and respiratory rate regressors [76, 77, 79]. This has led to a large pool of potential regressors. Simply including all these regressors in the regression, however, may be detrimental to the GLM analysis due to the loss in degrees of freedom (potentially reducing the statistical significance) and the increased chance of correlation between the confound regressors and the regressor of interest (i.e., the stimulus regressor). It is therefore desirable to choose the smallest set of regressors that are able to explain significant physiological variance.

We have previously collaborated on work that employed a Bayesian Informa-

tion Criterion to determine the optimal cut-off for the number of regressors for RETROICOR analysis of GRE-EPI data [143]. In that study the optimal set of regressors was chosen voxelwise. It was shown that the reduced set of regressors was able to explain a similar amount of the variance as full set of regressors without incurring significant reduced degrees-of-freedom. The type of regressors included and the total number was spatially dependent with the largest number of regressors chosen for the brainstem area. We will adopt a similar approach here, albeit with modifications appropriate to k -space corrections and 3D data.

The BIC Selection Procedure

The Bayesian Information Criterion (BIC) [144] is defined as:

$$BIC = N \ln \left(\frac{RSS}{N} \right) + k \ln(N) \quad (7.1)$$

where N is the number of samples (time points), k is the number of regressors, and RSS is the residual sum of squares. The BIC is used to compare different models, particularly models with a different number of parameters (here, regressors), with low BIC indicating a favourable model. BIC selection favours models that explain a large portion of the variance, but penalises a large number of regressors, thereby preventing over-fitting of the data.

Our selection procedure aims to find the optimal set of regressors iteratively by expanding the model to encapsulate one additional regressor in each iteration. Of the set of remaining regressors not yet included in the model, we choose the regressor that explains the greatest amount of variance in the raw data. We then compare the BIC of the model with and without this regressor. If the BIC is smaller with the regressor, we include it in our optimal set; otherwise, we consider

the penultimate model (without the most recently-added regressor) to be optimal. This way the set of regressors is build up step wise for each voxel until the BIC test fails [143].

The voxelwise optimization approach is an elegant method to show the spatial variation in model selection. However, simply copying the same approach in k -space is not compatible with optimizing for a specific spatial location (such as the brainstem). Moreover, running multiple regression models with a different number of EVs across the brain may introduce statistical differences that are driven by the varying DOF instead of the underlying activation when comparing activations across different brain regions.

We therefore propose a modification of the technique, which can be applied to RETROKCOR but still allows optimization for a specific brain region. In short this approach differs from the previous method in that 1) only one model is applied to the entire image (or k -space), and 2) the order in which the regressors are added to the model is determined at the start of the optimization instead of iteratively. The same procedure can be applied to RETROKCOR as well as RETROICOR:

STEP 1: Regressions are run using each of the candidate regressors (22 in total, see Section 7.2.3) as a single EV. Temporal variance maps are calculated for each of the data sets.

STEP 2: The regressors are sorted based on how much variance each regressor explains within the brainstem ROI. This is the order in which they will be added to the model in the selection process. The variance reduction is assessed based on the mean variance across the brainstem.

STEP 3: Using a similar approach to [143] (described above), regressions are performed by expanding the model one regressor at a time, but this time applying the same model throughout the brain. The BIC value is calculated using the mean

RSS within the brainstem ROI. The preferred model is the model with the lowest BIC.

The order of the regressors is determined at the start of the optimization (rather than iteratively) because an iterative scheme would be very computationally expensive for RETROKCOR. In the case of a single voxel RETROICOR model selection, the next candidate regressor is chosen by re-calculating how much of the remaining variance is explained by each of the candidate regressors. For RETROKCOR, each iteration would require a full k -space regression (on all coil channels) with each candidate EV. In our case, when all 22 model expansion steps were to be tested, this would result in a total number of $\binom{23}{2} * N_c * 2 = 6072$ regressions, where N_c is the number of coils. We therefore choose to order the EVs only once at the start of the optimization.

Correlation Between Regressors

The success of this ordering approach is dependent on the amount of correlation between the regressors. When the regressors are highly correlated, the combined effect of two added regressors is much smaller than expected based on the individual regressions. In this case, it might be desirable to instead add a regressor that explains less total variance on its own, but which is orthogonal to the first regressor. In order to test whether the proposed approach is likely to provide an acceptable ordering of the EVs, the correlation between the RETROICOR EVs was calculated and inspected.

Fig. 7.4 shows an example correlation matrix of the RETROICOR regressors generated for the centre slice in a 3DSHORT GRE-EPI data set (in RETROICOR the EVs are determined per slice). It shows that the correlation between the regressors was reasonably small (a maximum correlation of 0.20 was observed off diagonal). These numbers were highly consistent between subjects. Testing the




Figure 7.4: (a) Example correlation matrix between the candidate regressors in one data set. (b) The time courses of the two regressors that demonstrated the highest degree of correlation (only the first 50 time-points are displayed)

correlation between the RETROKCOR EVs is more difficult because the regressors are determined per k -space segment instead of only one per slice. However, as the sampling frequency (T_{vol}) is essentially identical between the GRE-EPI and the 3D acquisitions, we expect similar results for the RETROKCOR EVs. The correlation will be tested in retrospect, once the optimal correction models are selected, by comparing the final reduction in variance with the total reduction in variance predicted by single-EV regressors (which should be equal if the regressors are orthogonal).

7.4 Results

7.4.1 Model Selection Results on 3DSHORT-data

Fig. 7.5 shows the temporal variance of the residuals for single regressions with each of the EVs, averaged over the brainstem ROI. The solid black line is the variance averaged across the subjects with the errorbars representing the standard deviation. For each subject the variances were normalised to the variance of the non-corrected data set.

For the GRE-EPI, the first order cardiac terms, *ev01_cardcos_01* and *ev01_cardsin_01*, explain the greatest amount of variance followed by the cosine of the first order respiratory term (*ev07_respcos_01*) and the cardiac and respiratory rate regressors (*ev01_cr* and *ev21_rvt*). For the remaining EVs it is seen that the amount of variance explained by the interaction terms (ev15 - ev17) is equal to the amount of variance that a randomly constructed regressor would explain (represented by the grey dashed line). On average a randomly constructed regressor would explain $1/N$ of the variance in a data set, where N is the number of time points [143].

For bSSFP it is seen that the largest reduction in variance is achieved by the respiratory and cardiac rate regressors (ev19 and ev21) followed by the first order respiratory terms (ev7 and ev8) and two of the interaction terms (ev15 and ev16). The higher-order respiratory terms (ev9-ev14), however, all seem to introduce an *increase* the variance. This effect is observed in all subjects. This increase is initially confusing, since one would normally assume that regressing out a vector can only reduce or preserve the variance; however, regression in the Fourier domain can re-distribute signal fluctuations, causing local increases in variance (although the total variance in the image domain must by definition be reduced).

RETROKCOR appears to have little effect on the SPGR data in all subjects except one. For one subject a large portion of the regressions caused a strong

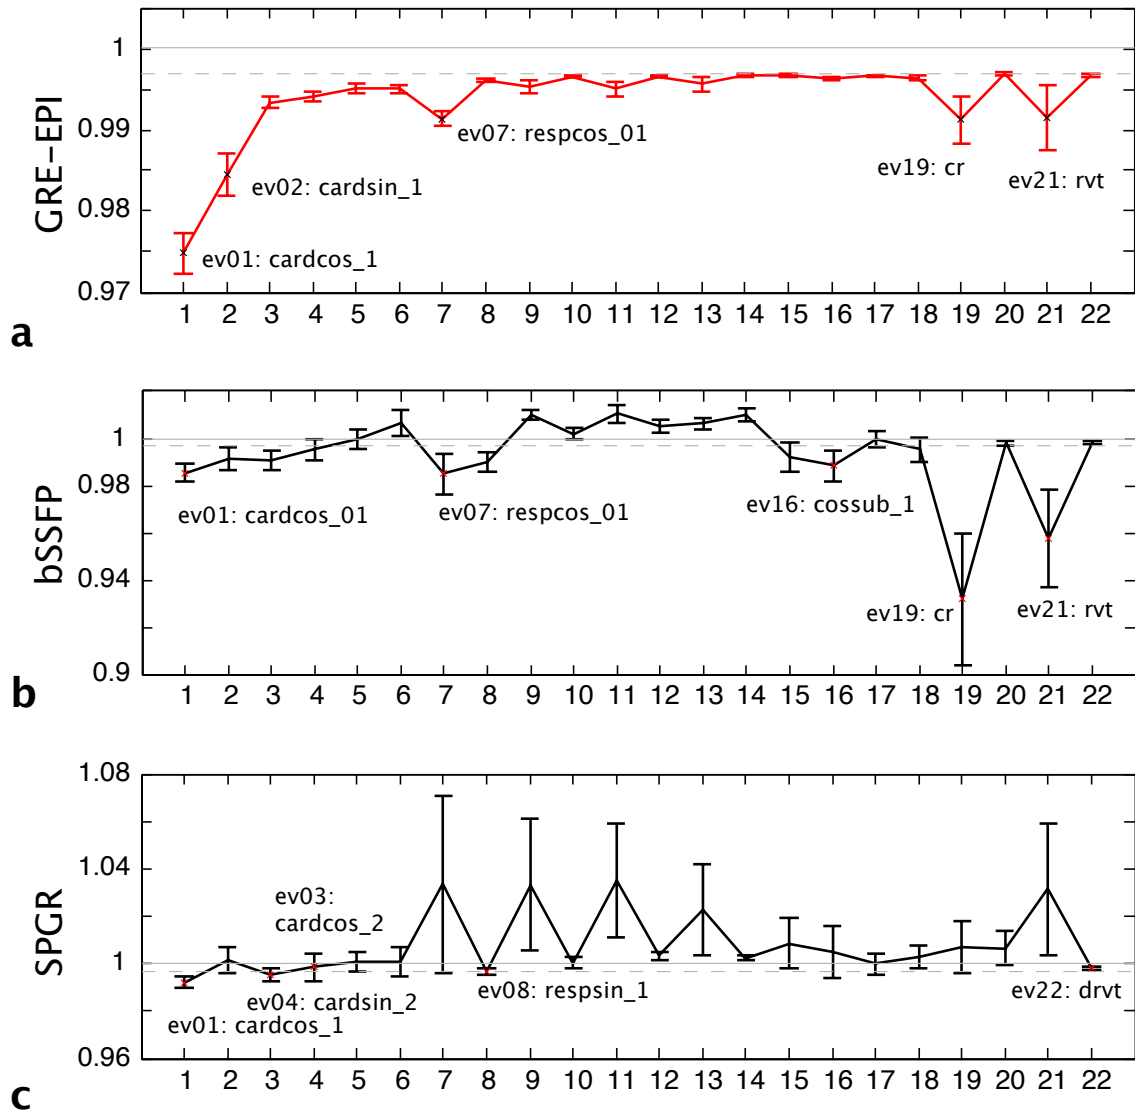


Figure 7.5: The residual variance after single EV-regressions for short- T_{vol} GRE-EPI, short- T_{vol} synchronized bSSFP, and short- T_{vol} synchronized SPGR. The variance is normalised against the total variance in the time-series. For each data set the 5 regressors that show the lowest residual variance (thus explaining a large amount of variance) are marked and listed below the graphs. For GRE-EPI RETROICOR results are shown (red), whereas synchronized bSSFP and SPGR show RETROKCOR results (black). Error bars denote the between subject standard deviation. The dashed grey line represents the expected residual variance obtained by a regression with a randomly constructed regressor

increase in variance. Although the strength of the effect compared to the other subjects suggests that this might have been outlier data, no obvious problems with the data acquisition and reconstruction were found, which is why the subject was not excluded. Inspection of the data showed large fluctuations outside the brainstem, which RETROKCOR corrected for, probably causing over-compensation in the rest of the brain (as was seen in the phantom and 2D *in vivo* data). The total sum of the variance across the entire image was reduced when the correction was performed.

The candidate regressors for the BIC selection were ordered according to the amount of variance each individual regressor explained starting with the regressor that explained the largest amount of variance (i.e., for GRE-EPI this would be *ev01_cardcos_01*). Table 7.4 at the end of this chapter lists the first 10 candidate regressors per data set and correction method. Based on this ordering the correction models were expanded in order to find the optimal model with lowest BIC value. For demonstration purposes, all models were expanded up to the full model that includes all 22 regressors. In practice, however, one could stop the expansion, once the BIC value starts to increase.

The left column in Fig. 7.6 shows the normalised variance of the residuals for each expansion step. The results are averaged across all subjects with error bars denoting the standard deviation. For GRE-EPI the variance is reduced most strongly by the first few EVs, after which the effect flattens to a constant reduction per added regressor. Although the effect of the first few added regressors varied between subjects, the slope to which the curves converged was highly consistent, which is to be expected as the regressors that do not explain actual physiological fluctuations still remove a certain amount of variance equal to that expected by a randomly constructed regressor. For bSSFP a similarly shaped curve as for GRE-EPI is observed. Note that the shape of these curves resembles the shape of the

eigenspectra curve common in PCA analysis, where the “knee” of the plot splits the significant and the presumably unimportant regressors [145]. In PCA the BIC selection is a well-established procedure for the estimation of the dimensionality. For the SPGR data a marginal reduction in variance was observed for the model that contained one regressor only (*ev01_cardcos_01*).

The right column in Fig. 7.6 shows the BIC curves, for each of the sequences, with the minima marked with red arrows. Like the variances, the BIC values are normalised against the non-corrected data set. The figures show that for GRE-EPI and bSSFP the minima of the BIC curves agree well with the location of the “knees” in the variance plots. For GRE-EPI the minimum is obtained when two EVs are included (*ev01_cardcos_01* and *ev02_cardsin_01*). For bSSFP the optimal model contains four regressors (*ev19_cr*, *ev21_rot*, *ev07_respcos_01*, and *ev08_respsin_01*). For SPGR the variance reduction by the regressor *ev01_cardcos_01* was not enough to pass the BIC selection.

7.4.2 Model Selection Results on 3DLONG-data

The BIC model selection approach was applied to the 3DLONG-data to test the consistency between data sets with varying acquisition parameters. We remind the reader that this data differ from the 3DSHORT-data primarily in terms of volume acquisition time ($T_{\text{vol}}=3.5$ s.) and the number of volumes acquired (60 volumes). The 3D data acquired for this data set was acquired with and without cardiac synchronization. RETROKCOR and RETROICOR corrections were performed on both types of acquisition. The effect of cardiac synchronization on the model selection was examined and the efficacy of RETROICOR and RETROKCOR are compared on both types of data.

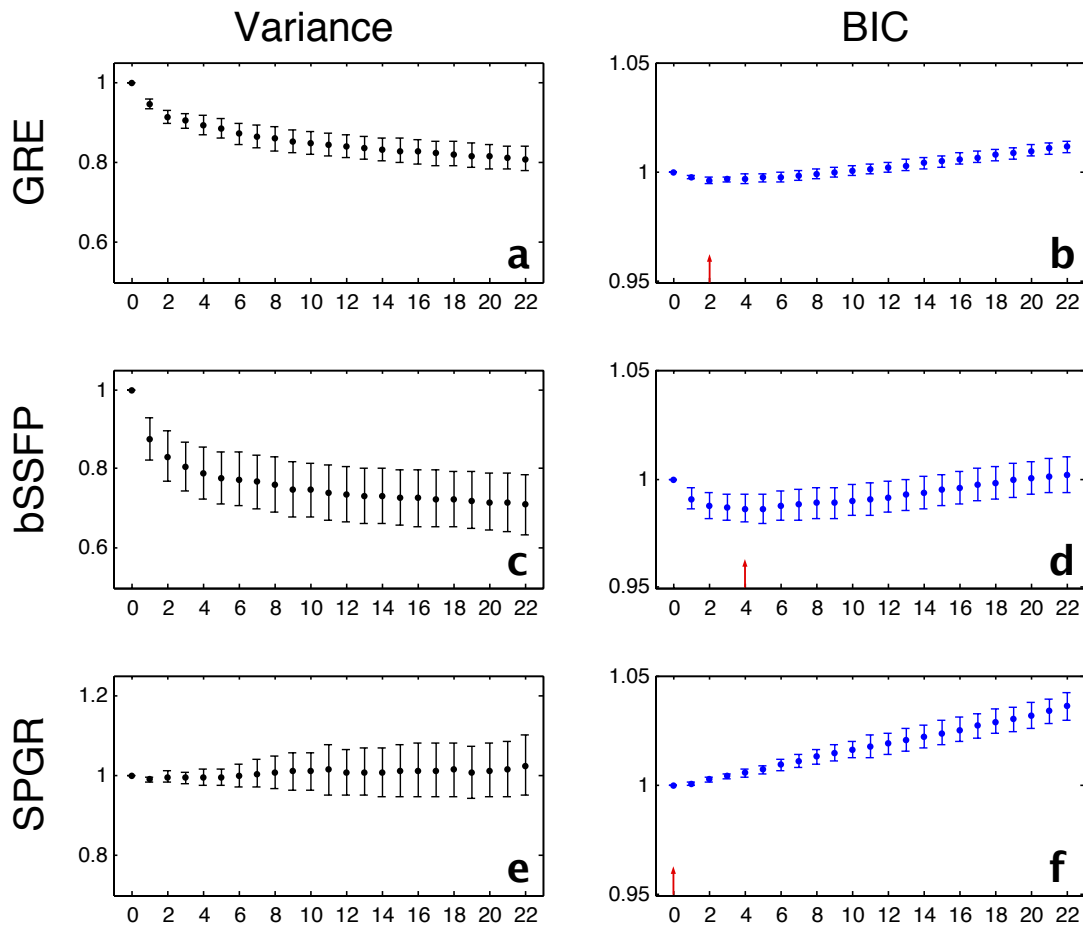


Figure 7.6: Plots showing the normalised variance of short- T_{vol} data after correction after each model expansion step (left column) and the corresponding BIC values (right column). The model with the lowest BIC value is the optimal model. For GRE-EPI the optimal model includes 2 regressors. For synchronized bSSFP four regressors, and for synchronized SPGR no regressors are determined optimal

The Effect of Temporal Sampling Rate and Scan Duration

Fig. 7.7(a) shows the same plot as shown in Fig. 7.5, but for the 3DLONG-data. When comparing the two GRE-EPI data sets, we observe similarly shaped curves, where the first order cardiac terms (ev01 and ev02) explain the greatest amount of variance amongst the EVs tested. Like for the 3DSHORT-data a small dip is observed for the cosine of the first order respiratory term (ev07).

The cardiac and respiratory rate regressors (ev19 and ev21), however, are no longer explaining more variance than a randomly constructed regressor would

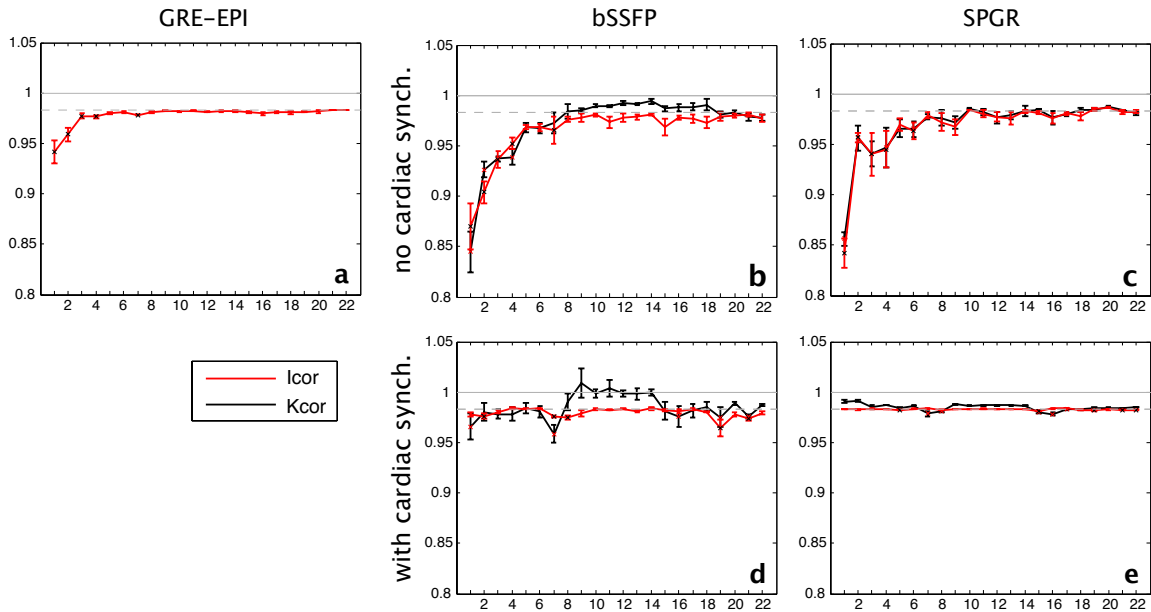


Figure 7.7: The residual variance after single EV-regressions for long- T_{vol} GRE-EPI (a), non-synchronized bSSFP (b), non-synchronized SPGR (c), synchronized bSSFP (d), and synchronized SPGR (e). RETROICOR results are shown in red and RETROKCOR results are shown in black. The mean across subjects is plotted where the error bars represent the standard deviation. The dashed grey line represents the expected residual variance obtained by a regression with a randomly constructed regressor

(grey dashed line). A potential reason for this might be that the total scan time was less in this data set (3.5 min vs. 10.5 min), and therefore less variance is introduced by the (slow varying) respiratory and heart rate. Note, that the amount of variance explained by a random regressor is larger for the 3DLONG-data, causing the variance curve to have a lower upper limit. This is due to the fact that the 3DLONG-data contains less time-points and a randomly constructed regressor therefore explains more of the variance ($1/60^{th}$ instead of $1/329^{th}$).

The Effect of Cardiac Synchronization

Panels (b) and (c) show the individual regression results of the non-synchronized data. Most strikingly is the difference in the amount of variance explained by the cardiac regressors. For both bSSFP and SPGR the cardiac regressors up to the 3rd order (ev01-ev06) show a large reduction in variance. The bSSFP and SPGR data

that were acquired with cardiac synchronization (Panels (d) and (e)) show similar trends as before, where little variance is explained by the cardiac regressors. This is a clear indication that cardiac synchronization is effectively cleaning up the cardiac related fluctuations in the brainstem.

For the other regressors, the differences between non-synchronized and synchronized are smaller. In non-synchronized bSSFP the higher order respiratory terms (ev09-ev14) explain more variance than in the synchronized version of the sequence, although still explaining less than a randomly created regressor. For SPGR all regressors apart from the cardiac regressors lie on the randomly constructed regressor asymptote in both the synchronized and non-synchronized version of the sequence indicating no physiological noise was removed by any of these regressors. Note that the results were very consistent across all subjects compared to the SPGR data presented in Fig. 7.5(c) where one subject showed large increases in variance for a subset of the regressors.

RETROICOR vs. RETROKCOR in 3D Data Sets

Secondary to the investigation of the effect of synchronization we investigated the performance of RETROICOR on both synchronized and non-synchronized data.

A potential pitfall with applying RETROICOR to 3D acquired data is that the cardiac and respiratory cycles are of similar duration to or even a few times shorter than T_{vol} . A 3.5 s. acquisition, for example, will on average contain a full respiratory cycle and approximately 3 cardiac cycles. It was therefore hypothesised that a correction in k -space, where each k -space segment can be assigned with its individual physiological phase and thus regressor, will perform better than RETROICOR with a single phase used for the entire volume.

However, when comparing the two methods we see that the obtained variance reduction of RETROICOR (red line) is comparable or occasionally better than

the results obtained with RETROKCOR. The synchronized bSSFP data show that RETROICOR does not suffer from the increased variance introduced by the higher order respiratory terms (ev09-ev14). For the other regressors the two methods perform equally well. Perhaps most unexpected is the fact that RETROICOR seems to regress out the (faster) cardiac fluctuations with similar success to RETROKCOR in the non-synchronized data when the regressions with individual EVs are compared. When the regressors are combined in our model expansion approach we see small differences between the two methods. For non-synchronized bSSFP the optimal RETROKCOR model contains six regressors (all cardiac, up to the third order; ev01-ev06), whereas for RETROICOR only the first two orders are selected (ev01-ev04). This is due to the fact that for RETROKCOR the total reduction in variance was slightly higher when the regressors were combined, which was just enough to select two extra regressors. For both SPGR data sets four regressors were selected (ev01-ev04).

The fact that the performance of RETROICOR is highly similar to RETROKCOR is welcome, as it vastly simplifies the correction methodology. Being able to use RETROICOR removes the need to collect the RAW k -space data from the scanner and reduces the number of regressions needed considerably, as only one regression on the reconstructed data would suffice (as opposed to separate regressions on each channel).

7.4.3 Optimized Sequence Comparison

In this final section we assess the final tSNR after correction of the data with the optimal model determined by the model selection described in the preceding sections. The regressors used in the correction for each data set are listed in bold type face in Table 7.4. The table shows that for the correction of the 3D bSSFP data in the 3DLONG data set, no regressors were selected by the BIC, in

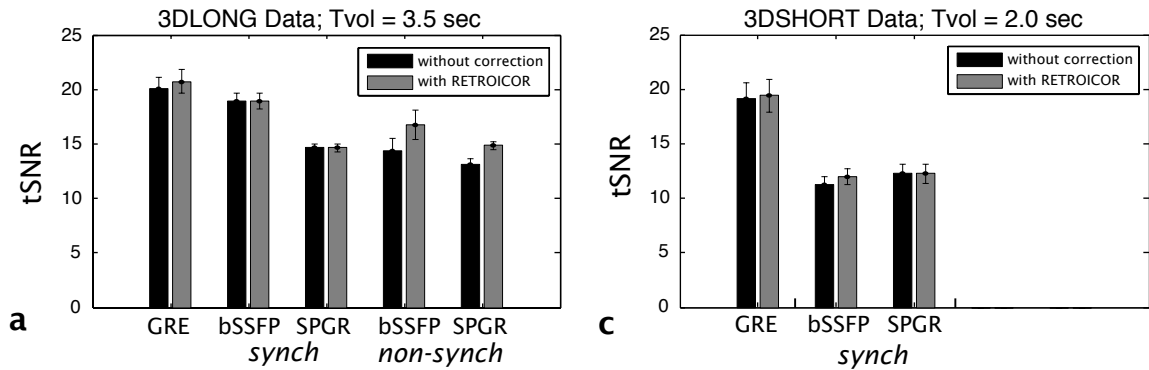


Figure 7.8: Bar plots showing the average tSNR in the brainstem before and after retrospective correction for (a) 3DLONG data and (b) 3DSHORT data. Error bars denote the between subject standard deviation

contrast to the 3DSHORT-data, where four regressors are used. The ordering of the regressors, however, is highly similar between 3DSHORT and 3DLONG bSSFP (for both RETROKCOR and RETROICOR the top four regressors is the same: ev19, ev21, ev07, and ev01), suggesting that short- T_{vol} and long- T_{vol} bSSFP are sensitive to the same physiological processes. The divergence in the BIC selection between the two acquisitions is probably caused by the fact that the 3DLONG time-series contain less time points (60 vs. 329 volumes). When fewer time-points are collected the BIC selection will be more conservative, choosing fewer regressors (Eq. 7.1) in order to prevent over-fitting.

Fig. 7.8 shows the mean tSNR in the brainstem before and after retrospective corrections for all 2D and 3D time series. Because RETROICOR was shown have equal performance to RETROKCOR only RETROICOR results are shown here. This figure shows that RETROICOR has the biggest effect on non-synchronized 3D data (Panel (a)). The increase in tSNR was 18% for bSSFP, and 15% for SPGR. The synchronized versions of the sequences, on the other hand, show identical tSNR before and after, since no regressors were selected by the BIC selection. GRE-EPI shows an improvement of 3.5%. For the 3DSHORT data, shown in Panel (b), the improvements were 2% for GRE-EPI, 6% synchronized bSSFP, and 0% for

synchronized SPGR (as no regressors were selected).

These values indicate that retrospective physiological noise correction only achieves a modest increase in tSNR. In the literature values for tSNR improvement exceeding 25% are often reported [95, 146]. Compared to these studies the numbers found here look very small. We believe the reason for this is two-fold: 1) in this study only very few regressors are used compared to previous literature, and 2) we defined a very conservative brainstem ROI in order to deliberately exclude regions of CSF surrounding the brainstem. Harvey *et al.*, reported a reduction in tSD of up to 15% in the brainstem, although most of the improvement was achieved near the boundaries with CSF. In the centre of the brainstem the reduction was approximately $\sim 5\%$, which is roughly in line with our findings given the reduced number of regressors used in this study. When we perform a regression on the long- T_{vol} GRE-EPI data using the full model of 22 regressors (something which would seriously reduce the DOF in this data set) we observe a tSNR increase of 29%. This is of the same order of magnitude as values previously reported in literature. Hutton *et al.* [95], found a 35% increase in the lateral geniculate nucleus at 7 Tesla when 14 regressors were used (12 cardiac and respiratory frequency regressors and two rate regressors).

The retrospective corrections had the biggest effect on non-synchronized 3D data (Panel (a)). The increase in tSNR was 18% for bSSFP, and 15% for SPGR. For SPGR the improvement in tSNR achieved with RETROICOR equals the improvement by synchronization. For bSSFP, however, we see that the tSNR in the synchronized non-corrected acquisition is higher than the non-synchronized, but retrospectively corrected sequence. Synchronization is therefore preferable, as it has the additional benefit that the DOF is not reduced.

Finally, comparing the 3D sequences to GRE-EPI we see that, when the data are acquired with a long T_{vol} , synchronized bSSFP is the only 3D acquisition

that shows a tSNR that is comparable to the tSNR of GRE-EPI. However, for the short- T_{vol} data the tSNR in both synchronized bSSFP and SPGR has dropped considerably, unlike GRE-EPI, which only shows a marginal reduction in tSNR. Because this limits the applicability of the developed 3D sequences it is important to understand what causes this reduction in tSNR. One possibility is that the reordering is not as effective in short- T_{vol} acquisitions, which might explain why the reduction in tSNR is larger for bSSFP than it is for SPGR. In the short- T_{vol} acquisitions the total number of segments that k -space consists of is less than for the short- T_{vol} because the FOV is reduced in order to bring the T_{vol} down. During the acquisition, the synchronization continuously tries to acquire the optimal segment in k -space based on the current phase in the cardiac cycle. If the pool of segments that can be picked from gets smaller the efficacy of the synchronization will decrease (see Chapter 6). Because short- T_{vol} non-synchronized data were not acquired, a direct assessment of the effect on tSNR was not possible. However, inspection of the cardiac phases at which each of the k -space segments was acquired showed no strong differences in the synchronization efficacy between 3DLONG and 3DSHORT acquisitions. Figure 7.9 demonstrates the tSD of the cardiac phase for each k -space segment. The cardiac phase ranges from 0–1. For a perfectly synchronized acquisition in which each segment could be acquired at the exact same cardiac phase the tSD would be 0, whereas in a non-synchronized sequence the tSD will be close to 1. It is seen that when the data are acquired with a long T_{vol} (Panel (a)) the segments show a smaller tSD compared to the short T_{vol} data (Panel (b)). However, the differences are small compared to the non-synchronized data in (Panel (c)). This indicates that the synchronization efficacy is not explaining the drop in tSNR that is observed.

The second effect that is likely to contribute to this effect is the signal sensitivity that has changed for the 3D sequences, but not for the GRE-data. For 3D




Figure 7.9: Illustration of the efficacy of synchronization in long- T_{vol} synchronized (a), short- T_{vol} synchronized (b), and non-synchronized (c) bSSFP and SPGR. The colormap denotes the tSD of the cardiac phase of the k -space location. A perfect reordering would result in a low variability in cardiac phase over time. It is seen that synchronization is slightly more effective when a longer T_{vol} is used, although differences are small compared to non-synchronized data.

sequences the signal that is acquired is proportional to: $S_{3D} \propto \rho * T_{vol}$, where ρ is the duty cycle that represents the fraction of the time within a TR that is spent receiving the signal. Because in 3D imaging the entire volume is sampled with every TR, the signal is directly proportional to T_{vol} . For 2D multi-slice imaging, on the other hand, a slice is only excited once per T_{vol} . The signal is therefore sampled for approximately 30–50ms (in the case of single-shot EPI), regardless of the T_{vol} . One factor that will influence the signal for 2D multi-slice for different volume acquisition times (=TR) is the longitudinal relaxation. Bloch calculations, however, demonstrate that the longitudinal relaxation for 3DLONG (TR=3500, $\alpha=90^\circ$) and 3DSHORT (TR=2000, $\alpha=90^\circ$) is within a 10% difference, which is in line with the small reduction in tSNR observed in our data. For the 3D data, on the other hand, the reduction from 3.5 s. to 2.0 s. in T_{vol} is predicted to have a much larger effect on the SNR. Unfortunately, however, due to the small FOV in the short- T_{vol} data, no reliable SNR_0 measurement could be made, but this is a topic we plan to look into in the near future.

7.5 General Discussion

The goal of this study was to optimize the retrospective corrections on 3D brain-stem acquisitions in order to further improve the temporal stability such that ultimately the functional contrast to noise ratio (CNR) would be comparable to the “gold-standard” GRE-EPI while affording low distortion images.

7.5.1 Optimization of RETROKCOR

It was shown that the problems associated with the corrections of the higher spatial frequencies in RETROKCOR can be mitigated by performing the correction on the real and imaginary instead of the magnitude and phase channels of the data. A pictorial explanation of this observation is given in Fig. 7.10, which shows the signal S as a vector sum of the real and imaginary channels. Each of the channels has a certain amount of noise ϵ associated to it. Panel (a) depicts the situation, in which $|S| \gg \epsilon$, and Panel (b) the situation where $|S|$ is of the same order as the noise. The black circles denote the standard deviation (σ) The red circles denotes the 95% interval (2σ). When $|S| \gg \epsilon$ the uncertainty of the phase (σ_φ) is proportional to the diameter of the standard deviation circle perpendicular to the magnitude vector $|S|$. The standard deviation of the magnitude and the real and imaginary channels are also depicted. When the signal magnitude decreases (Panel (b)) we see that the angle representing the phase uncertainty (represented by the angle between the green lines) increases in a non-linear fashion. The uncertainty of the real and imaginary signals, on the other hand, is constant throughout the entire range. The standard uncertainty of the phase as a function of signal magnitude is shown in Panel (c) (plot reproduced from [147]).




Figure 7.10: The effect of noise on the phase/magnitude and real/imaginary channels at different SNR. The standard deviation of the noise (ϵ) is represented by the black circle and the 95% interval by the red circle. The uncertainty is proportional to the radius of the σ -circle (the angle between the green lines). As the SNR decreases (Panel (b)) this angle representing the phase uncertainty increases non-linearly. Panel (c) shows the standard uncertainty of the phase as a function of signal magnitude when simulated Gaussian noise is added to the real and imaginary channels with $\sigma=0.01$ (sub-plot reproduced from [147])

7.5.2 Optimization of the Regressor Set

As shown by Hutton *et al.* [95], the inclusion of more confound regressors does not necessarily introduce improvements in BOLD sensitivity. For example, due to the loss in DOF or correlations of the confound regressor with the regressor of interest (i.e., the stimulus). Most of the previous work, however, have not accounted for this and only report reductions in variance [8, 80]. The BIC selection proposed takes into account the loss in DOF by applying a penalty term for the number of regressors used. This weight of penalty term, however, could be varied. Other closely related selection criterion methods exist such as the Akaike information criterion (AIC) [148]. The AIC differs from BIC in that the penalty term for the number of regressors is smaller. Therefore a larger regressor set would be preferred. In order to investigate the optimal weight that should be given to the penalty term, an assessment of the functional sensitivity should be included. In this study, however, such an assessment was not included. Given the disappointing tSNR, the low number of subjects included at this stage, and the small effect size of the stimulus investigated no convincing functional results were obtained.

7.5.3 The Potential for 3D Brainstem FMRI using bSSFP

The results presented in this chapter show that in the most favourable case the tSNR of bSSFP is comparable to that of GRE-EPI. SPGR showed reduced tSNR compared to GRE-EPI. A reasonable predictor of the CNR (e.g., z-statistic) would be the ratio of the signal change to temporal variability. A previous study using visual stimulation and comparable sequences reported a $\sim 2\%$ signal change in bSSFP and a $\sim 4\%$ signal change for GRE [13]. This indicates that the tSNR of bSSFP would have to be at least twice as high compared to GRE-EPI in order to achieve comparable CNR, since the CNR ($\Delta S/\sigma$) is calculated by the product of tSNR (\bar{S}/σ) and the relative signal change ($\Delta S/\bar{S}$).

Our original motivation in pursuing this work was based in part on the higher thermal SNR (i.e., of single images) in bSSFP, but also the early observation of reduced physiological noise in bSSFP in the cortex [13]. The hope was that these effects might combine to overcome the reduced contrast in bSSFP to result in improvements over GRE-EPI. What the present study has established is that the phase interference effects introduced to 3D segmented acquisitions by physiological noise in the brainstem are a significant barrier to this goal, and one that the current methods are insufficient to surmount. However, outside the brainstem there has also been a renewed interest in 3D acquisitions, particularly at 7 Tesla, in order to overcome some of the challenges with GRE-EPI at higher fields. [52, 149]. Given the physiological noise becomes more apparent even in the cortical regions at high fields [86, 150, 95], the methods developed here to correct for physiological noise in 3D acquisitions could find their application in those situations. In addition to noise being more significant, the reduced T_2^* should boost the contrast in short-TE segmented acquisitions. For example, we are in the planning stages of re-visiting the hippocampal FMRI described in chapter 4 at 7 Tesla.

7.6 Conclusion

In this chapter retrospective corrections on 2D, 3D synchronized, and 3D non-synchronized data were characterized and optimized. The efficacy of RETROKCOR and RETROICOR was compared and refinements were made to the existing techniques. It was shown that RETROKCOR can be improved by performing the correction on the real and imaginary channels instead of the phase and magnitude as originally proposed.

Further, a modified version of the BIC selection method by Jenkinson *et al.*, was implemented and used to select the most suitable set of regressors for RETROICOR and RETROKCOR corrections on 2D and 3D data. It was shown that a large portion of the commonly used regressors does not explain more than a randomly constructed regressor and therefore the BIC-selected subset was generally small (the maximum number of regressors chosen was 6).

It was demonstrated that RETROICOR performs equally well as RETROKCOR in highly segmented multi-shot sequences, even though only a single cardiac and respiratory phase can be assigned to each volume in the time-series. This is a welcome result, as RETROICOR simplifies the correction vastly by removing the need for off-line reconstruction of the data. RETROICOR further allows the confound regressors to be included in the GLM together with the stimulus regressor, which statistically is a better approach [74]. For example, if confound regressors are correlated to the stimulus paradigm, a separate regression would always result in an underestimation of the functional activation as the overlapping signal would always be fully contributed to the confound regressor.

Finally it was shown that the tSNR of fully correction 3D acquisitions are only comparable to GRE-EPI when the volume acquisition time is set long enough to achieve sufficient signal sensitivity. For a T_{vol} of 3.5 s. synchronized bSSFP demonstrated comparable tSNR to GRE-EPI, but the tSNR was reduced consider-

ably when the T_{vol} was reduced to 2 s. This shows that particular care should be taken when protocols are set up for 3D FMRI.

Seq.	<i>3DSHORT-Data synch</i>		<i>3DSHORT-Data no synch</i>	
	RETROKCOR	RETROICOR	RETROKCOR	RETROICOR
bSSFP	ev19_cr	ev19_cr		
	ev21_rvt	ev21_rvt		
	ev07_respcos_01	ev07_respcos_01		
	ev01_cardcos_01	ev08_respsin_01		
	ev16_cossb_01c01r	ev11_respcos_03		
	ev08_respsin_01	ev09_respcos_02		
	ev03_cardcos_02	ev02_cardsin_01		
SPGR	ev01_cardcos_01	ev19_cr		
	ev03_cardcos_02	ev07_respcos_01		
	ev08_respsin_01	ev21_rvt		
	ev22_drvt	ev01_cardcos_01		
	ev04_cardsin_02	ev08_respsin_01		
	ev17_sinadd_01c01r	ev02_cardsin_01		
	ev10_respsin_02	ev09_respcos_02		
GRE				ev01_cardcos_01
				ev02_cardsin_01
				ev19_cr
				ev07_respcos_01
				ev21_rvt
			ev03_cardcos_02	
			ev04_cardsin_02	
Seq.	<i>3DLONG-Data synch</i>		<i>3DLONG-Data no synch</i>	
	RETROKCOR	RETROICOR	RETROKCOR	RETROICOR
bSSFP	ev07_respcos_01	ev19_cr	ev01_cardcos_01	ev01_cardcos_01
	ev01_cardcos_01	ev21_rvt	ev02_cardsin_01	ev02_cardsin_01
	ev19_cr	ev08_respsin_01	ev03_cardcos_02	ev03_cardcos_02
	ev21_rvt	ev07_respcos_01	ev04_cardsin_02	ev04_cardsin_02
	ev16_cossb_01c01r	ev02_cardsin_01	ev06_cardsin_03	ev07_respcos_01
	ev04_cardsin_02	ev01_cardcos_01	ev05_cardcos_03	ev06_cardsin_03
	ev03_cardcos_02	ev20_dcr	ev07_respcos_01	ev05_cardcos_03
SPGR	ev16_cossb_01c01r	ev15_cosadd_01c01r	ev01_cardcos_01	ev01_cardcos_01
	ev07_respcos_01	ev19_cr	ev03_cardcos_02	ev03_cardcos_02
	ev15_cosadd_01c01r	ev22_drvt	ev04_cardsin_02	ev04_cardsin_02
	ev08_respsin_01	ev21_rvt	ev02_cardsin_01	ev02_cardsin_01
	ev17_sinadd_01c01r	ev05_cardcos_03	ev06_cardsin_03	ev06_cardsin_03
	ev18_sinsub_01c01r	ev08_respsin_01	ev05_cardcos_03	ev09_respcos_02
	ev21_rvt	ev18_sinsub_01c01r	ev09_respcos_02	ev05_cardcos_03
GRE				ev01_cardcos_01
				ev02_cardsin_01
				ev04_cardsin_02
				ev03_cardcos_02
				ev07_respcos_01
			ev05_cardcos_03	
			ev16_cossb_01c01r	

Table 7.4: This table lists the order of the candidate regressors for the BIC selection. The regressors were ordered based on the amount of variance explained each regressor explained individually. The first 10 regressors are listed. The regressors in bold type face are the regressors that were included by the BIC selection.

Chapter 8

Summary and Future Work

This final chapter summarises the findings in this thesis and includes a discussion on the significance of the developed methods described in the previous chapters. In addition, future research ideas are outlined, including ideas for further refinement of the developed methods and applications for functional imaging at high magnetic field.

8.1 Thesis Summary

Characterization of 3D Sequences for Brainstem FMRI

In this set of initial experiments the potential of 3D bSSFP and RF-spoiled SPGR was investigated. A characterization based on rapidly acquired 2D data suggested a similar advantage in the brainstem for bSSFP over GRE than that was found in the cortex in earlier work [13]. However, when a 3D readout is used the instabilities increased dramatically. The sensitivity to physiological noise of multi-shot 3D readouts is thus a severe obstacle that has to be overcome in order to make 3D techniques viable for brainstem FMRI and take advantage of their unique properties.

Physiological Noise Associated with 3D Acquisitions

In this chapter the influence of multi-shot acquisitions on signal instabilities caused by physiological noise was investigated. It was shown that the signal instabilities mainly originate from regions of CSF and blood and are highly correlated to the cardiac cycle. Several correction methods were explored, including novel ways to spoil CSF signal in bSSFP acquisitions. Using simulations, a potential correction method that employs synchronization of the k -space acquisition with respect to the cardiac cycle was identified as the method to implement *in vivo*.

Cardiac Synchronization with Fixed T_{vol}

A new method which allows real-time cardiac synchronization of the k -space acquisition was developed. A custom parallel imaging reconstruction was also developed to allow for acquisition with a fixed volume frame rate. A fixed volume frame rate has clear benefits for the paradigm design and data analysis in fMRI. The method significantly improved the temporal stability in areas that were most affected by cardiac-related signal fluctuations. In the brainstem the tSNR was found to increase by up to 45% in 3D bSSFP data, and up to 20% in 3D SPGR time-series data.

Optimization of Retrospective Corrections for 2D and 3D Acquisitions

A comprehensive assessment of two currently available correction techniques was provided and their practicalities compared. It was found that the original RETROKCOR method, which performs its corrections on the magnitude and phase channels of the data, is sensitive to imperfections in the correction of outer k -space locations. It was shown that this effect could largely be mitigated by performing corrections on the real and imaginary data instead. Furthermore, a method was implemented that selects the optimal regressor set by taking into account the degrees of freedom in the GLM.

Overall Conclusion

We have shown that the signal instabilities can be significantly reduced in 3D acquisition methods with the correction methods that were developed. The final comparison between optimized 3D bSSFP and GRE-EPI revealed that for brainstem FMRI at 3 Tesla 3D sequences have a comparable, but not superior tSNR. Based on these findings and the fact that a higher functional sensitivity was found for GRE-EPI in the cortex [12] we expect GRE-EPI to have superior CNR in the brainstem. Nevertheless, we are optimistic that the methods developed for this thesis will be valuable for other applications such as 3D FMRI at high field strength. This is further discussed in the Future Research section below. The somewhat unexpected result that RETROICOR appears to be just as capable in removing physiological signal fluctuations as RETROKCOR for 3D FMRI is a welcome finding as it greatly simplifies the correction in these sequences and thus increases their general applicability to FMRI.

8.2 Future Work

As briefly mentioned in the previous chapter there has been a renewed interest in 3D acquisitions for high resolution FMRI at 7 Tesla [149, 9]. In this regime 3D acquisitions may have the greatest benefit, as issues such as limited slice thickness and long volume acquisition times become more pressing for 2D GRE-EPI (i.e., B_0 and B_1 inhomogeneities become more apparent and T_{vol} is proportional to the number of slices in multi-slice imaging). Poser *et al.* [9] have demonstrated that 3D acquisitions allow accelerated parallel imaging and Partial Fourier along two dimensions, affording shorter volumes acquisition times as compared to 2D GRE-EPI. However, the same study also reports increased sensitivity to physiological effects in 3D acquisitions in the cortical regions at 7T. The methods developed in

this thesis could therefore find excellent use in those applications.

High-resolution 3D FMRI at 7 Tesla

We are planning a study with the primary aim to develop methods that allow high resolution, low distortion, FMRI of the visual area and hippocampus at 7 Tesla. As 3D GRE-EPI has already been successfully used for high-resolution FMRI of the visual cortex this will be our starting point. However, like the reduced T_2^* the optimum T_2 is also reduced at 7T, which is a reason for also including the bSSFP as a potential candidate. A T_2 -weighted sequence will have the added benefit that the spatial point spread function (PSF) of the BOLD response is smaller compared to a T_2^* -weighted sequence. At 7T the PSF is 2–3.5 mm for GRE and < 2 mm for SE [48]. If the initial assessment of the sequences indicates that 3D acquisitions are viable for visual and hippocampal FMRI at 7T, further developments will be pursued, including extending the spatial coverage to allow for whole brain coverage acquisitions.

Improving Spatial Coverage

In order to achieve whole brain coverage with a temporal resolution that is sufficient for FMRI we aim to make several changes to the sequence. At the moment, accelerated parallel imaging is not implemented into the sequence. However, as noted above, 3D acquisitions have the advantage to achieve high acceleration factors by subsampling the data in both phase-encode directions. Poser *et al.* achieved compelling functional results using a 3×3 accelerated 3D EPI sequence.

We therefore aim to include parallel acceleration into our cardiac synchronization method. With the protocols tested so far, on average 15% of the segments are identified as being corrupt and replaced using our GRAPPA reconstruction. This indicates that there is enough scope for additional acceleration by regularly

subsampling the k -space lines in two dimensions. With this type of acceleration the hope is to achieve an acceleration factor of 2×2 (i.e., a factor of 2 in each of the phase encode directions).

Another acceleration technique, which has shown promising results is compressed sensing (CS) [151]. Although the use of CS in fMRI is still in its infancy this technique could potentially yield high acceleration factors with minimal loss in image quality. CS requires the sampling pattern to be random such that the aliasing due to the undersampling is non-structured. 3D acquisitions have a benefit over 2D acquisitions in this perspective as they allow a wider range of (semi-) random trajectories. A fairly straightforward trajectory to implement for example would be sample lines along k_x that are randomly placed in the k_y - k_z dimension.

Reducing the Sensitivity to Temporal Frequency Drifts

At 7 Tesla temporal frequency drifts are expected to have a bigger effect on bSSFP data. It is crucial for bSSFP fMRI that the region of interest is placed in the pass-band region of the bSSFP signal profile. Temporal B_0 drifts caused by gradient heating or oscillatory B_0 fluctuations induced by the respiratory cycle could shift the bands into the tissue of interest over the course of the experiment. If this proves to be a limiting factor at 7 Tesla we intend to implement a real-time feedback method, which measures the B_0 offset and updates the shims dynamically using a simple navigator or FID similar to the method proposed in [58].

Optimization of the BIC approach

In the previous chapter the BIC criterion was used to place a certain penalty on the inclusion of new confound regressors in RETROICOR or RETRKCOR corrections. However, other selection criterions exist such as the Akaike information criterion (AIC), which put different weights on the inclusion of additional regressors. In order to fully test, which of the criterions is best to use, a functional study

would be needed in which the functional sensitivity could be compared between data sets that are corrected using each of the optimization schemes. This type of assessment could be included in the hippocampus study proposed above.

References

- [1] A. C. Guyton and J. E. Hall, *Textbook of Medical Physiology*. W. B. Saunders Company, tenth ed., 2000.
- [2] H. Breivik, B. Collett, V. Ventafridda, R. Cohen, and D. Gallacher, "Survey of chronic pain in europe: prevalence, impact on daily life, and treatment.," *Eur J Pain*, vol. 10, no. 4, pp. 287–333, 2006.
- [3] L. Zambreanu, R. G. Wise, J. C. W. Brooks, G. D. Iannetti, and I. Tracey, "A role for the brainstem in central sensitisation in humans. Evidence from functional magnetic resonance imaging," *Pain*, vol. 114, no. 3, pp. 397–407, 2005.
- [4] M. Fairhurst, K. Wiech, P. Dunckley, and I. Tracey, "Anticipatory brainstem activity predicts neural processing of pain in humans," *Pain*, vol. 128, no. 1-2, pp. 101–110, 2007.
- [5] F. Beissner, R. Deichmann, and S. Baudrexel, "fMRI of the brainstem using dual-echo EPI," *Neuroimage*, Jan 2011.
- [6] C. M. Cahill and P. W. Stroman, "Mapping of neural activity produced by thermal pain in the healthy human spinal cord and brain stem: a functional magnetic resonance imaging study," *Magn Reson Imaging*, vol. 29, pp. 342–52, Apr 2011.
- [7] S. V. Astafiev, A. Z. Snyder, G. L. Shulman, and M. Corbetta, "Comment on "modafinil shifts human locus coeruleus to low-tonic, high-phasic activity during functional MRI" and "homeostatic sleep pressure and responses to sustained attention in the suprachiasmatic area"," *Science*, vol. 328, p. 309; author reply 309, Apr 2010.
- [8] A. K. Harvey, K. T. S. Pattinson, J. C. W. Brooks, S. D. Mayhew, M. Jenkinson, and R. G. Wise, "Brainstem functional magnetic resonance imaging: disentangling signal from physiological noise.," *J Magn Reson Imaging*, vol. 28, no. 6, pp. 1337–1344, 2008.
- [9] B. Poser, P. Koopmans, L. Wald, and M. Barth, "Investigation into the benefits of 3D-EPI for high-resolution fMRI at 7T," in *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*, p. 3657, 2009.

- [10] K. L. Miller, S. M. Smith, P. Jezzard, and J. M. Pauly, "High-resolution fMRI at 1.5T using balanced SSFP," *Magn Reson Med*, vol. 55, no. 1, pp. 161–170, 2006.
- [11] J. H. Lee, S. O. Dumoulin, E. U. Saritas, G. H. Glover, B. A. Wandell, D. G. Nishimura, and J. M. Pauly, "Full-brain coverage and high-resolution imaging capabilities of passband b-SSFP fMRI at 3T.," *Magn Reson Med*, vol. 59, no. 5, pp. 1099–1110, 2008.
- [12] K. Zhong, J. Leupold, J. Hennig, and O. Speck, "Systematic investigation of balanced steady-state free precession for functional MRI in the human visual cortex at 3 Tesla," *Magnetic Resonance in Medicine*, vol. 57, no. 1, pp. 67–73, 2007.
- [13] K. L. Miller, S. M. Smith, P. Jezzard, G. C. Wiggins, and C. J. Wiggins, "Signal and noise characteristics of SSFP fMRI: A comparison with GRE at multiple field strengths," *NeuroImage*, vol. 37, no. 4, pp. 1227–1236, 2007.
- [14] L. G. Hanson, "Is quantum mechanics necessary for understanding magnetic resonance," *Concepts in Magnetic Resonance*, vol. 32A, no. 5, pp. 329–240, 2008.
- [15] M. Haacke, R. Brown, M. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley-Liss, June 1999.
- [16] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI pulse sequences*. Elsevier Academic Press, 1st ed., 2004.
- [17] D. G. Nishimura, *Principles of Magnetic Resonance Imaging*. Lulu, 2010.
- [18] S. A. Huettel, A. W. Song, and G. McCarthy, *Functional Magnetic Resonance Imaging*. Sunderland, MA: Sinauer Associates, Inc., 2004.
- [19] R. B. Buxton, *Introduction to Functional Magnetic Resonance Imaging*. Cambridge: Cambridge University Press, 1st ed., 2002.
- [20] E. L. Hahn, "Spin echoes," *Phys Rev*, vol. 80, no. 4, pp. 580–594, 1950.
- [21] J. P. Wansapura, S. K. Holland, R. S. Dunn, and W. S. B. Jr., "NMR relaxation times in the human brain at 3.0 Tesla," *Journal of Magnetic Resonance Imaging*, vol. 9, no. 4, pp. 531–538, 1999.
- [22] M. J. Donahue, H. Lu, C. K. Jones, R. A. E. Edden, J. J. Pekar, and P. C. M. van Zijl, "Theoretical and experimental investigation of the VASO contrast mechanism," *Magn Reson Med*, vol. 56, pp. 1261–73, Dec 2006.
- [23] S. K. Piechnik, J. Evans, L. H. Bary, R. G. Wise, and P. Jezzard, "Functional changes in CSF volume estimated using measurement of water T2 relaxation," *Magn Reson Med*, vol. 61, pp. 579–86, Mar 2009.

- [24] F. Bloch, "Nuclear induction," *Phys.Rev.*, vol. 70, no. 7-8, pp. 460–474, 1946.
- [25] P. Mansfield and I. Pykett, "Biological and medical imaging by NMR," *Journal of Magnetic Resonance*, vol. 29, pp. 255–373, 1978.
- [26] J. Frahm, A. Haase, and D. Matthaei, "Rapid NMR imaging of dynamic processes using the FLASH technique," *Magn Reson Med*, vol. 3, pp. 321–7, Apr 1986.
- [27] A. P. Crawley, M. L. Wood, and R. M. Henkelman, "Elimination of transverse coherences in FLASH MRI," *Magn Reson Imaging*, vol. 8, no. 3, pp. 248–260, 1988.
- [28] A. Haase, J. Frahm, D. Matthaei, W. Hanicke, and M. K., "FLASH imaging. rapid NMR imaging using low flip-angle pulses," *J Magn Reson*, vol. 67, pp. 258–266, 1986.
- [29] K. L. Miller, R. H. Tijssen, and T. W. Stikov N. Okell, "Steady-state MRI: Methods for neuroimaging," *Imaging in Medicine*, vol. 3, pp. 93–105, 2011.
- [30] Y. Zur, M. L. Wood, and L. J. Neuringer, "Spoiling of transverse magnetization in steady-state sequences," *Magn Reson Med*, vol. 21, no. 2, pp. 251–263, 1991.
- [31] J. H. Duyn, "Steady state effects in fast gradient echo magnetic resonance imaging," *Magn Reson Med*, vol. 37, pp. 559–68, Apr 1997.
- [32] K. Scheffler and S. Lehnhardt, "Principles and applications of balanced SSFP techniques," *European Radiology*, vol. 13, no. 11, pp. 2409–2418, 2003.
- [33] H. Y. Carr, "Steady-state free precession in nuclear magnetis resonance," *Phys.Rev.*, vol. 112, pp. 1693–1701, 1958.
- [34] R. R. Ernst and W. A. Anderson, "Application of Fourier transform spectroscopy to magnetic resonance," *Rev.Sci.Instrum.*, vol. 37, pp. 93–102, 1966.
- [35] R. Freeman and H. D. W. Hill, "Phase and intensity anomalies in Fourier transform NMR," *Journal of Magnetic Resonance*, vol. 4, no. 3, pp. 366–383, 1971.
- [36] E. Jaynes, "Matrix treatment of nuclear induction," *Phys.Rev.*, vol. 98, pp. 1099–1105, 1955.
- [37] B. Aldefeld and P. Börnert, "Effects of gradient anisotropy in MRI," *Magn Reson Med*, vol. 39, pp. 606–14, Apr 1998.
- [38] A. Zakhor, R. M. Weisskoff, and R. Rzedzian, "Optimal sampling and reconstruction of MRI signals resulting from sinusoidal gradients," *IEEE Trans Signal Proc*, vol. 39, pp. 2056–2065, 1991.

- [39] A. Jesmanowicz, E. E. Wong, and J. S. Hyde, "Phase correction for EPI using internal reference lines," in *Proceedings of the 12th SMRM*, 1993.
- [40] J. M. Pauly, "Non-cartesian reconstruction; partial k-space reconstruction," Stanford University, 2004.
- [41] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, "The NMR phased array," *Magn Reson Med*, vol. 16, no. 2, pp. 192–225, 1990.
- [42] J. W. Carlson and T. Minemura, "Imaging time reduction through multiple receiver coil data acquisition and image reconstruction.," *Magn Reson Med*, vol. 29, no. 5, pp. 681–687, 1993.
- [43] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI.," *Magn Reson Med*, vol. 42, no. 5, pp. 952–962, 1999.
- [44] 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)," *Magn Reson Imaging*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [45] T. Liu, B. Kressler, K. Wang, and Y. Wang, "Block circulant quasi-band matrix property for the SENSE unfolding in k-space and justification for GRAPPA," *Conf Proc IEEE Eng Med Biol Soc*, vol. 2008, pp. 1659–62, 2008.
- [46] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magn Reson Med*, vol. 43, pp. 682–90, May 2000.
- [47] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation," *Proc Natl Acad Sci U S A*, vol. 87, no. 24, pp. 9868–9872, 1990.
- [48] D. G. Norris, "Principles of magnetic resonance assessment of brain function," *J Magn Reson Imaging*, vol. 23, pp. 794–807, Jun 2006.
- [49] K. L. Miller, B. A. Hargreaves, J. Lee, D. Ress, R. C. de Charms, and J. M. Pauly, "Functional MRI using a blood oxygenation sensitive steady state," vol. 50, no. 4, pp. 675–683, 2003.
- [50] P. Jezzard and S. Clare, "Sources of distortion in functional MRI data.," *Hum Brain Mapp*, vol. 8, no. 2-3, pp. 80–85, 1999.
- [51] F. Hennel, "Multiple-Shot Echo-Planar Imaging," *concepts in magnetic resonance*, vol. 9, no. 1, pp. 43–58, 1997.

- [52] B. A. Poser, P. J. Koopmans, T. Witzel, L. L. Wald, and M. Barth, "Three dimensional echo-planar imaging at 7 Tesla," *Neuroimage*, vol. 51, pp. 261–6, May 2010.
- [53] Y. Hu and G. H. Glover, "Three-dimensional spiral technique for high-resolution functional MRI.," *Magn Reson Med*, vol. 58, no. 5, pp. 947–951, 2007.
- [54] M. Barth and D. G. Norris, "Very high-resolution three-dimensional functional MRI of the human visual cortex with elimination of large venous vessels," *NMR Biomed*, vol. 20, pp. 477–84, Aug 2007.
- [55] K. Scheffler, E. Seifritz, D. Bilecen, R. Venkatesan, J. Hennig, M. Deimling, and E. M. Haacke, "Detection of BOLD changes by means of a frequency-sensitive trueFISP technique: Preliminary results," *NMR Biomed*, vol. 14, no. 7-8, pp. 490–496, 2001.
- [56] C. V. Bowen, R. S. Menon, and J. S. Gati, "High field balanced-SSFP fMRI: A BOLD technique with excellent tissue sensitivity and superior large vessel suppression," in *Proceedings of the 13th Annual Meeting of ISMRM*, (Miami), p. 119, 2005.
- [57] C. V. Bowen, J. Mason, R. S. Menon, and J. S. Gati, "High field balanced-SSFP fMRI: Examining a diffusion contrast mechanism using varied flip angles," in *Proceedings of the 14th Annual Meeting of ISMRM*, (Seattle), p. 665, 2006.
- [58] J. Lee, J. M. Santos, S. M. Conolly, K. L. Miller, B. A. Hargreaves, and J. M. Pauly, "Respiration-induced B₀ field fluctuation compensation in balanced SSFP: real-time approach for transition-band SSFP fMRI.," *Magn Reson Med*, vol. 55, no. 5, pp. 1197–1201, 2006.
- [59] E. N. Marieb, *Human Anatomy and Physiology*. Benjamin Cummings, fifth ed., 2001.
- [60] K. T. S. Pattinson, G. D. Mitsis, A. K. Harvey, S. Jbabdi, S. Dirckx, S. D. Mayhew, R. Rogers, I. Tracey, and R. G. Wise, "Determination of the human brainstem respiratory control network and its cortical connections in vivo using functional and structural imaging," *Neuroimage*, vol. 44, pp. 295–305, Jan 2009.
- [61] L. M. Slabu, "The effect of slice orientation on auditory fMRI at the level of the brainstem," *Brain Topogr*, vol. 23, pp. 301–10, Sep 2010.
- [62] I. Tracey and G. D. Iannetti, "Brainstem functional imaging in humans.," *Supplements to Clinical neurophysiology.*, vol. 58, pp. 52–67, 2006.

- [63] W.-T. Zhang, C. Mainero, A. Kumar, C. J. Wiggins, T. Benner, P. L. Purdon, D. S. Bolar, K. K. Kwong, and A. G. Sorensen, "Strategies for improving the detection of fMRI activation in trigeminal pathways with cardiac gating," *Neuroimage*, vol. 31, pp. 1506–12, Jul 2006.
- [64] K. D'Ardenne, S. M. McClure, L. E. Nystrom, and J. D. Cohen, "BOLD responses reflecting dopaminergic signals in the human ventral tegmental area," *Science*, vol. 319, pp. 1264–7, Feb 2008.
- [65] X. Hu, T. H. Le, T. Parrish, and P. Erhard, "Retrospective estimation and correction of physiological fluctuation in functional MRI," *Magn Reson Imaging*, vol. 34, no. 2, pp. 201–212, 1995.
- [66] P. Jezzard, "Physiological noise: Strategies for correction," in *Functional MRI* (C. T. W. Moonen and P. A. Bandettini, eds.), pp. 173–182, Springer, 1 ed., 1999.
- [67] B. Wowk, M. C. McIntyre, and J. K. Saunders, "k-Space detection and correction of physiological artifacts in fMRI," *Magn Reson in Med*, vol. 38, no. 6, pp. 1029–1034, 1997.
- [68] D. R. Enzmann and N. J. Pelc, "Brain motion: Measurement with phase-contrast MR imaging," *Radiology*, vol. 185, no. 3, pp. 653–660, 1992.
- [69] B. Biswal, A. E. DeYoe, and J. S. Hyde, "Reduction of physiological fluctuations in fMRI using digital filters.," *Magn Reson Med*, vol. 35, no. 1, pp. 107–113, 1996.
- [70] B. G. Goodyear, H. Zhu, R. A. Brown, and J. R. Mitchell, "Removal of phase artifacts from fMRI data using a Stockwell transform filter improves brain activity detection," *Magn Reson Med*, vol. 51, pp. 16–21, Jan 2004.
- [71] G. H. Glover, T. Q. Li, and D. Ress, "Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR," *Magn Reson in Med*, vol. 44, no. 1, pp. 162–167, 2000.
- [72] E. B. Beall and M. J. Lowe, "Isolating physiologic noise sources with independently determined spatial measures.," *Neuroimage*, vol. 37, no. 4, pp. 1286–1300, 2007.
- [73] R. H. R. Deckers, P. van Gelderen, M. Ries, O. Barret, J. H. Duyn, V. N. Ikonomidou, M. Fukunaga, G. H. Glover, and J. A. de Zwart, "An adaptive filter for suppression of cardiac and respiratory noise in MRI time series data," *Neuroimage*, vol. 33, pp. 1072–81, Dec 2006.
- [74] T. E. Lund, K. H. Madsen, K. Sidaros, W.-L. Luo, and T. E. Nichols, "Non-white noise in fMRI: does modelling have an impact?," *Neuroimage*, vol. 29, no. 1, pp. 54–66, 2006.

- [75] 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*, vol. 21, pp. 1652–64, Apr 2004.
- [76] 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*, vol. 31, no. 4, pp. 1536–1548, 2006.
- [77] 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*, vol. 38, pp. 306–20, Nov 2007.
- [78] M. Bianciardi, M. Fukunaga, P. van Gelderen, S. G. Horovitz, J. A. de Zwart, K. Shmueli, and J. H. Duyn, "Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study," *Magn Reson Imaging*, vol. 27, pp. 1019–29, Oct 2009.
- [79] C. Chang, J. P. Cunningham, and G. H. Glover, "Influence of heart rate on the BOLD signal: the cardiac response function," *Neuroimage*, vol. 44, pp. 857–69, Feb 2009.
- [80] T. B. Jones, P. A. Bandettini, and R. M. Birn, "Integration of motion correction and physiological noise regression in fMRI," *Neuroimage*, vol. 42, pp. 582–90, Aug 2008.
- [81] K. Scheffler and J. Hennig, "Is TrueFISP a gradient-echo or a spin-echo sequence?," *Magn Reson Imaging*, vol. 49, no. 2, pp. 395–397, 2003.
- [82] E. Yacoub, P.-F. Van De Moortele, A. Shmuel, and K. Ugurbil, "Signal and noise characteristics of Hahn SE and GE BOLD fMRI at 7 T in humans," *Neuroimage*, vol. 24, pp. 738–50, Feb 2005.
- [83] G. Krüger and G. H. Glover, "Physiological noise in oxygenation-sensitive magnetic resonance imaging," *Magn Reson Imaging*, vol. 46, no. 4, pp. 631–637, 2001.
- [84] C. Triantafyllou, J. Polimeni, and L. Wald, "Physiological noise in gradient echo and spin echo EPI using multi-channel array coils," in *Proceedings 16th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Toronto*, p. 2465, 2008.
- [85] 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," *Magn Reson Med*, Jun 2011.
- [86] 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*, vol. 26, no. 1, pp. 243–250, 2005.
- [87] C. Triantafyllou, J. R. Polimeni, and L. L. Wald, "Physiological noise and signal-to-noise ratio in fmri with multi-channel array coils," *Neuroimage*, vol. 55, pp. 597–606, Mar 2011.
- [88] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. J. Behrens, H. Johansen-Berg, P. R. Bannister, M. D. Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. D. Stefano, J. M. Brady, and P. M. Matthews, "Advances in functional and structural MR image analysis and implementation as FSL," *NeuroImage*, vol. 23, no. SUPPL. 1, 2004.
- [89] S. M. Smith, "Fast robust automated brain extraction," *Hum Brain Mapp*, vol. 17, pp. 143–55, Nov 2002.
- [90] 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*, vol. 17, pp. 825–41, Oct 2002.
- [91] Y. Zhang, M. Brady, and S. Smith, "Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm," *IEEE Trans Med Imaging*, vol. 20, pp. 45–57, Jan 2001.
- [92] B. Patenaude, S. M. Smith, D. N. Kennedy, and M. Jenkinson, "A Bayesian model of shape and appearance for subcortical brain segmentation," *Neuroimage*, vol. 56, pp. 907–22, Jun 2011.
- [93] G. S. Pell, R. S. Briellmann, A. B. Waites, D. F. Abbott, D. P. Lewis, and G. D. Jackson, "Optimized clinical T2 relaxometry with a standard CPMG sequence," *J Magn Reson Imaging*, vol. 23, pp. 248–52, Feb 2006.
- [94] C. Triantafyllou, J. Polimeni, M. Elschot, and L. Wald, "Physiological noise in gradient echo and spin echo EPI at 3T and 7T," in *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*, p. 122, 2009.
- [95] 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*, vol. 57, pp. 101–12, Jul 2011.
- [96] P. Kellman and E. R. McVeigh, "Image reconstruction in SNR units: a general method for SNR measurement.," *Magn Reson Med*, vol. 54, no. 6, pp. 1439–1447, 2005.
- [97] G. Schroth and U. Klose, "Cerebrospinal fluid flow. I. physiology of cardiac-related pulsation.," *Neuroradiology*, vol. 35, no. 1, pp. 1–9, 1992.

-
- [98] S. Peltier and D. C. Noll, "Physiological noise in multi-shot functional MRI," 2002.
- [99] S. S. Vasanaawala, J. M. Pauly, and D. G. Nishimura, "Linear combination steady-state free precession MRI," *Magn Reson Med*, vol. 43, pp. 82–90, Jan 2000.
- [100] N. K. Bangerter, B. A. Hargreaves, S. S. Vasanaawala, J. M. Pauly, G. E. Gold, and D. G. Nishimura, "Analysis of multiple-acquisition SSFP," *Magn Reson Med*, vol. 51, pp. 1038–47, May 2004.
- [101] K. Fliessbach, P. Trautner, C. M. Quesada, C. E. Elger, and B. Weber, "Cerebellar contributions to episodic memory encoding as revealed by fMRI.," *Neuroimage*, vol. 35, no. 3, pp. 1330–1337, 2007.
- [102] R. Weisskoff, J. Baker, J. Belliveau, T. Davis, K. Kwong, M. Cohen, and B. Rosen, "Power spectrum analysis of functionally-weighted MR data: what's in the noise?," in *Proceedings of the 12th SMRM*, (New York), p. 7, 1993.
- [103] U. Klose, S. Friese, M. Erb, and W. Grodd, "Physiological MR signal variations within the brain at 3 T.," *Biomed Tech (Berl)*, vol. 52, no. 1, pp. 126–129, 2007.
- [104] D. A. Feinberg and A. S. Mark, "Human brain motion and cerebrospinal fluid circulation demonstrated with MR velocity imaging," *Radiology*, vol. 163, pp. 793–9, Jun 1987.
- [105] R. L. Barry, L. M. Klassen, J. M. Williams, and R. S. Menon, "Hybrid two-dimensional navigator correction: a new technique to suppress respiratory-induced physiological noise in multi-shot echo-planar functional MRI.," *Neuroimage*, vol. 39, pp. 1142–1150, Feb 2008.
- [106] R. S. Menon, C. G. Thomas, and J. S. Gati, "Investigation of BOLD contrast in fMRI using multi-shot EPI," *NMR Biomed*, vol. 10, no. 4-5, pp. 179–82, 1997.
- [107] A. R. Guimaraes, J. R. Melcher, T. M. Talavage, J. R. Baker, P. Ledden, B. R. Rosen, N. Y. Kiang, B. C. Fullerton, and R. M. Weisskoff, "Imaging subcortical auditory activity in humans," *Hum Brain Mapp*, vol. 6, no. 1, pp. 33–41, 1998.
- [108] M. H. Cho, W. S. Kim, and Z. H. Cho, "CSF flow artifact reduction using cardiac cycle ordered phase-encoding method," *Magn Reson Imaging*, vol. 8, no. 4, pp. 395–405, 1990.

- [109] D. R. Bailes, D. J. Gilderdale, G. M. Bydder, A. G. Collins, and D. N. Firmin, "Respiratory ordered phase encoding (ROPE): a method for reducing respiratory motion artefacts in MR imaging," *J Comput Assist Tomogr*, vol. 9, no. 4, pp. 835–8, 1985.
- [110] V. A. Stenger, S. Peltier, F. E. Boada, and D. C. Noll, "3D spiral cardiac/respiratory ordered fMRI data acquisition at 3 Tesla," *Magn Reson Med*, vol. 41, pp. 983–91, May 1999.
- [111] P. Storey, W. Li, Q. Chen, and R. R. Edelman, "Flow artifacts in steady-state free precession cine imaging.," *Magn Reson Med*, vol. 51, no. 1, pp. 115–122, 2004.
- [112] M. Markl, M. T. Alley, C. J. Elkins, and N. J. Pelc, "Flow effects in balanced steady state free precession imaging.," *Magn Reson Med*, vol. 50, no. 5, pp. 892–903, 2003.
- [113] O. Bieri and K. Scheffler, "Flow compensation in balanced SSFP sequences.," *Magn Reson Med*, vol. 54, no. 4, pp. 901–907, 2005.
- [114] D. L. Foxall, "Frequency-modulated steady-state free precession imaging," *Magn Reson Med*, vol. 48, pp. 502–8, Sep 2002.
- [115] L. Howden, D. Giddings, H. Power, A. Aroussi, M. Vloeberghs, M. Garnett, and D. Walker, "Three-dimensional cerebrospinal fluid flow within the human ventricular system," *Comput Methods Biomech Biomed Engin*, vol. 11, pp. 123–33, Apr 2008.
- [116] A. A. Linninger, B. Sweetman, and R. Penn, "Normal and hydrocephalic brain dynamics: the role of reduced cerebrospinal fluid reabsorption in ventricular enlargement," *Ann Biomed Eng*, vol. 37, pp. 1434–47, Jul 2009.
- [117] B. A. Hargreaves, S. S. Vasanaawala, J. M. Pauly, and D. G. Nishimura, "Characterization and reduction of the transient response in steady-state MR imaging," *Magn Reson Med*, vol. 46, pp. 149–58, Jul 2001.
- [118] K. Scheffler, S. Maderwald, M. E. Ladd, and O. Bieri, "Oscillating steady states," *Magn Reson Med*, vol. 55, pp. 598–603, Mar 2006.
- [119] W. R. Overall, S. M. Conolly, D. G. Nishimura, and B. S. Hu, "Oscillating dual-equilibrium steady-state angiography," *Magn Reson Med*, vol. 47, pp. 513–22, Mar 2002.
- [120] H. Lu, L. M. Nagae-Poetscher, X. Golay, D. Lin, M. Pomper, and P. C. M. van Zijl, "Routine clinical brain MRI sequences for use at 3.0 Tesla," *J Magn Reson Imaging*, vol. 22, pp. 13–22, Jul 2005.

- [121] S. K. Piechnik, P. A. Chiarelli, and P. Jezzard, "Modelling vascular reactivity to investigate the basis of the relationship between cerebral blood volume and flow under CO₂ manipulation," *Neuroimage*, vol. 39, pp. 107–18, Jan 2008.
- [122] R. R. Ingle and D. G. Nishimura, "The "central signal singularity" phenomenon in balanced SSFP," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, 2011.
- [123] C. Ganter, "Steady state of gradient echo sequences with radiofrequency phase cycling: analytical solution, contrast enhancement with partial spoiling," *Magn Reson Med*, vol. 55, pp. 98–107, Jan 2006.
- [124] S. Gupta, M. Soellinger, P. Boesiger, D. Poulikakos, and V. Kurtcuoglu, "Three-dimensional computational modeling of subject-specific cerebrospinal fluid flow in the subarachnoid space," *J Biomech Eng*, vol. 131, p. 021010, Feb 2009.
- [125] M. Blaimer, F. A. Breuer, M. Mueller, N. Seiberlich, D. Ebel, R. M. Heidemann, M. A. Griswold, and P. M. Jakob, "2D-GRAPPA-operator for faster 3D parallel MRI," *Magn Reson Imaging*, vol. 56, no. 6, pp. 1359–1364, 2006.
- [126] P. Jakob, M. Griswold, F. Breuer, M. Blaimer, and N. Seiberlich, "a 3D GRAPPA algorithm for volumetric parallel imaging," in *Proceedings of the 14th Annual Meeting of ISMRM*, (Seattle), p. 286, 2006.
- [127] R. H. Britt and G. T. Rossi, "Quantitative analysis of methods for reducing physiological brain pulsations.," *J Neurosci Methods*, vol. 6, no. 3, pp. 219–229, 1982.
- [128] N. A. Shkumat, J. H. Siewerdsen, A. C. Dhanantwari, D. B. Williams, N. S. Paul, J. Yorkston, and R. Van Metter, "Cardiac gating with a pulse oximeter for dual-energy imaging," *Phys Med Biol*, vol. 53, pp. 6097–112, Nov 2008.
- [129] E. M. Haacke, E. D. Lindskog, and W. Lin, "A fast, iterative, partial-fourier technique capable of local phase recovery," *Journal of Magnetic Resonance*, vol. 92, no. 1, 1991.
- [130] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson, D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman, M. S. Albert, and R. J. Killiany, "An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest," *Neuroimage*, vol. 31, pp. 968–80, Jul 2006.
- [131] U. Goerke, H. E. Moller, D. G. Norris, and C. Schwarzbauer, "A comparison of signal instability in 2D and 3D EPI resting-state fMRI.," *NMR Biomed*, vol. 18, no. 8, pp. 534–542, 2005.

- [132] Y. Hu and G. H. Glover, "Partial-k-space acquisition method for improved SNR efficiency and temporal resolution in 3D fMRI," *Magn Reson Med*, vol. 55, pp. 1106–13, May 2006.
- [133] D. C. Noll and W. Schneider, "Simulation and compensation of physiological motion artifacts in functional MRI," in *IEEE International Conference on Image Processing, Austin*, pp. 40–44, 1994.
- [134] O. Bieri, M. Markl, and K. Scheffler, "Analysis and compensation of Eddy currents in balanced SSFP.," *Magn Reson Med*, vol. 54, no. 1, pp. 129–137, 2005.
- [135] J. Jorge, P. Figueirido, W. van der Zwaag, M. Narsude, and J. Marques, "Sources of signal fluctuations in single-shot 2d EPI and segmented 3D EVI acquisitions for fMRI at 7T," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, 2011.
- [136] A. Lutti, O. Josephs, D. Thomas, R. Lawson, J. Roiser, C. Hutton, and N. Weiskopf, "Optimized physiological noise correction for 3D EPI time series," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 3635, 2011.
- [137] R. L. Barry, S. C. Strother, and J. C. Gore, "Complex and magnitude-only preprocessing of 2D and 3D bold fMRI data at 7 T," *Magn Reson Med*, vol. (Epub ahead of print), Jul 2011.
- [138] 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*, vol. 39, no. 2, pp. 680–692, 2008.
- [139] S. Friese, U. Hamhaber, M. Erb, and U. Klose, "B-waves in cerebral and spinal cerebrospinal fluid pulsation measurement by magnetic resonance imaging," *J Comput Assist Tomogr*, vol. 28, no. 2, pp. 255–62, 2004.
- [140] V. Ibañez, M. P. Deiber, and F. Mauguière, "Interference of vibrations with input transmission in dorsal horn and cuneate nucleus in man: a study of somatosensory evoked potentials (seps) to electrical stimulation of median nerve and fingers," *Exp Brain Res*, vol. 75, no. 3, pp. 599–610, 1989.
- [141] J. W. Morley and M. J. Rowe, "Perceived pitch of vibrotactile stimuli: effects of vibration amplitude, and implications for vibration frequency coding," *J Physiol*, vol. 431, pp. 403–16, Dec 1990.
- [142] G. E. Hagberg, M. Bianciardi, V. Brainovich, A. M. Cassara, and B. Maraviglia, "The effect of physiological noise in phase functional magnetic resonance imaging: from blood oxygen level-dependent effects to direct detec-

- tion of neuronal currents.," *Magn Reson Imaging*, vol. 26, no. 7, pp. 1026–1040, 2008.
- [143] M. Jenkinson, R. Tijssen, J. Brooks, and K. Miller, "Voxel-wise regressor selection for physiological noise correction with RETROICOR," in *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*, p. 1582, 2009.
- [144] R. Kass and A. Raftery, "Bayes factors and model uncertainty," Tech. Rep. 254, University of Washington, 1993.
- [145] C. F. Beckmann and S. M. Smith, "Probabilistic independent component analysis for functional magnetic resonance imaging," *IEEE Trans Med Imaging*, vol. 23, pp. 137–52, Feb 2004.
- [146] R. M. Birn, K. Murphy, P. Bandettini, and J. Bodurka, "The use of multiple physiologic parameter regression increases gray matter temporal signal to noise by up to 50%," in *Proceedings 14th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Seattle*, 2006.
- [147] C. M. Williams, D. F. ans Wang and U. Arz, "In-phase/quadrature covariance-matrix representation of the uncertainty of vectors and complex numbers," in *68th ARTFG Conf. Dig.*, pp. 62–65, 2006.
- [148] H. Akaike, "A new look at the statistical model identification," *IEEE Trans Autom Control*, 1974.
- [149] M. Barth, H. Meyer, S. A. R. Kannengiesser, J. R. Polimeni, L. L. Wald, and D. G. Norris, "T2-weighted 3D fMRI using S2-SSFP at 7 Tesla," *Magn Reson Med*, vol. 63, pp. 1015–20, Apr 2010.
- [150] P. Van de Moortele, J. Pfeuffer, G. H. Glover, K. Ugurbil, and X. Hu, "Respiration-induced B0 fluctuations and their spatial distribution in the human brain at 7 Tesla.," *Magn Reson Med*, vol. 47, no. 5, pp. 888–895, 2002.
- [151] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn Reson Med*, vol. 58, pp. 1182–95, Dec 2007.

Appendix A

Definition of the Confound Regressors

This appendix provides a more detailed description on how the confound regressors used in Chapter 7 were calculated. Where appropriate, mathematical descriptions are given.

Cardiac Phase

The cardiac waveform that is a trigger waveform that records a trigger each time an RR wave in the ECG signal is detected. From the trigger waveform the phase is calculated by:

$$\varphi_c[t] = t_{trig}[t]/T_{RR} \quad (\text{A.1})$$

where t_{trig} is the time that since the last trigger and T_{RR} the duration of the current RR interval. This definition is identical to the definition of the cardiac phase used for cardiac synchronization. The difference, however, is that the interval T_{RR} of each cardiac cycle can be determined by calculating the time difference between the preceding and following trigger. In the case of cardiac synchronization, which is performed in real-time, T_{RR} had to be estimated from the average heart rate.

Respiratory Phase

In order to define the respiratory phase, the histogram-equalised transfer function as described in Glover *et al.* [71] was used. This approach takes into account the depth of breathing and distinguishes between inhale and exhale. The continuous respiratory waveform, $R(t)$, is normalised to the maximum amplitude of inhalation, R_{max} such that the range is $(0, R_{max})$. From this normalised waveform a histogram, $H(b)$, is calculated. For a given value $R(t)$ the respiratory phase is defined by the area under the histogram left of $R(t)$ divided by the area under the entire histogram, and multiplied by the sign of dR/dt :

$$\varphi_r[t] = \pi \frac{\sum_{b=1}^{\lceil R(t)/R_{max} \rceil} H[b]}{\sum_{b=1}^{\lceil R_{max} \rceil} H[b]} \text{sgn}\left(\frac{dR}{dt}\right) \quad (\text{A.2})$$

Using this approach, the end-expiration is assigned a phase of 0π , whereas peak inspiration is assigned a value of $\pm\pi$ depending on whether we are at the end of inspiration or at the start of expiration.

Frequency Regressors

The frequency regressors are defined by the low order Fourier terms. The first order is the primary frequency, whereas the higher orders are also called the harmonics. Per order two regressors are generated, which are calculated by taking the sine and the cosine value for the current cardiac or respiratory phase. The signal that describes the cardiac ($C[t]$) and respiratory ($R[t]$) fluctuations can be written as the sum of the regressors:

$$\begin{aligned} C[t] &= \sum_{m=1}^M [\beta_{c1}^m \cos(m\varphi_c[t]) + \beta_{c2}^m \sin(m\varphi_c[t])] \\ R[t] &= \sum_{n=1}^N [\beta_{r1}^n \cos(n\varphi_r[t]) + \beta_{r2}^n \sin(n\varphi_r[t])] \end{aligned} \quad (\text{A.3})$$

Here, m and n are the order for the cardiac and respiratory waveforms, respectively. For the study in Chapter 7 we used: $M \leq 3$ and $N \leq 4$. The β 's are the individual amplitude of each of the regressors that are estimated by the linear regression.

Interaction Regressors

The interaction terms describe the interactions between the cardiac and respiratory (for example the effect of respiratory-induced pressure changes on the cardiac rhythm). Per order four regressors are created. The regressors are given by:

$$\begin{aligned} X[t] &= \sum_{m=1}^M \sum_{n=1}^N [\beta_{x1}^{m,n} \sin(m\varphi_c[t] + n\varphi_r[t]) + \beta_{x2}^{m,n} \cos(m\varphi_c[t] + n\varphi_r[t]) + \\ &\quad \beta_{x3}^{m,n} \sin(m\varphi_c[t] - n\varphi_r[t]) + \beta_{x4}^{m,n} \cos(m\varphi_c[t] - n\varphi_r[t])] \end{aligned} \quad (\text{A.4})$$

Again, β denotes the amplitude for each regressor that is estimated by the linear regression, and M and N represent the Fourier orders that are generated. In Chapter 7 only first order regressors ($M = N = 1$) were used.

Cardiac and Respiratory Rate

Using the approach described in Shmueli *et al.* [77] the cardiac rate was calculated by taking the inverse of the period of each RR interval. This waveform was then smoothed with a moving average filter that spanned 10 s. The resulting signal was re-sampled at the acquisition rate of each k -space segment (in the case of RETROKCOR) or each slice (in the case of RETROICOR).

For respiratory, a regressor that takes into account the respiratory volume per time (RVT) was generated. The same procedure as described in [76] was followed. First, the amount of air inspired by each breath was calculated by taking the difference between the minimum and maximum peaks of the respiratory waveform. This difference is then divided by the duration of the respiratory cycle (i.e., the time between the peaks). The resulting waveform was smoothed with a moving average filter that spanned 10 s.

For both these regressors the derivative was included to allow for temporal shifts of the regressors.