Diffusion-Weighted Magnetic Resonance Imaging with Readout-Segmented Echo-Planar Imaging

Robert Frost

FMRIB Centre
and Magdalen College

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

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

Trinity Term 2012

Diffusion-Weighted Magnetic Resonance Imaging with Readout-Segmented Echo-Planar Imaging

Robert Frost

Magdalen College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Trinity Term 2012

Abstract

Diffusion-weighted (DW) magnetic resonance imaging is an important neuroimaging technique that has successful applications in diagnosis of ischemic stroke and methods based on diffusion tensor imaging (DTI). Tensor measures have been used for detecting changes in tissue microstructure and for non-invasively tracing white matter connections *in vivo*. The most common image acquistion strategy is to use a DW single-shot echo-planar imaging (ss-EPI) pulse sequence, which is attractive due to its robustness to motion artefacts and high imaging speed. However, this sequence has limited achievable spatial resolution and suffers from geometric distortion and blurring artefacts.

Readout-segmented echo-planar imaging (rs-EPI) is a DW sequence that is capable of acquiring high-resolution images by segmenting the acquisition of k-space into multiple shots. The fast, short readouts reduce distortion and blurring and the problem of artefacts due to motion-induced phase changes between shots can be overcome with navigator techniques.

The rs-EPI sequence has two main shortcomings. (i) The method is slow to produce image volumes, which is limiting for clinical scans due to patient welfare and prevents us from acquiring very many directions in DTI. (ii) The sequence (like other diffusion techniques) is far from the optimum repetition time (TR) for acquiring data with the highest possible signal-to-noise ratio (SNR) in a given time. The work in this thesis seeks to address both of these important issues using a range of approaches.

In Chapter 4 a partial Fourier extension is presented, which addresses point (i) by reducing the number of readout segments acquired and estimating the missing data. This allows reductions in scan time by approximately 40% and the reliability of the images is demonstrated in comparisons with the original images.

The application of a simultaneous multi-slice scheme to rs-EPI, to address points (i) and (ii), is described in Chapter 5. Using the slice-accelerated rs-EPI sequence, tractography data were compared to ss-EPI data and high-resolution trace-weighted data were acquired in clinically relevant scan times.

Finally, a 3D multi-slab extension that addresses point (i) is presented in Chapter 6. A 3D sequence could also allow higher resolution in the slice direction than 2D multi-slice methods, which are limited by the difficulties in exciting thin, accurate slices. A 3D version of rs-EPI was simulated and implemented and a k-space acquisition synchronised to the cardiac cycle showed substantial improvements in image artefacts compared to a conventional k-space acquisition.

Acknowledgements

I would like to thank all the people who have contributed to this thesis in various ways. It has been fascinating to work in the same building as so many talented people who are always willing to share their expertise. The friendly atmosphere in FMRIB has also made it a very enjoyable time.

I am extremely grateful to my supervisors Prof. Peter Jezzard and Dr. Karla Miller. It has been a pleasure to work with and learn from them and I thank them for their unfailingly patient and honest advice. Peter's extensive knowledge of MRI and attention to detail is amazing and his guidance and support throughout my DPhil has been invaluable. Karla has always been available to explain concepts and discuss practicalities with me at length and whose comments on drafts of this thesis and on recent work have been incredibly helpful and deserve special mention. Dr. David Porter has also been very helpful and supportive as an external supervisor and I am very grateful for the time and advice (and of course the sequence!) he has given us during this project. I consider myself very lucky to have had the opportunity to work closely with such talented, knowledgeable researchers whose work is always of the highest standard.

Many physics group students and post-docs also deserve thanks, especially James Meakin, who has painstakingly explained concepts in MRI, programming etc. on numerous occasions. His knowledge of MRI and many other topics (especially cycling!) is incredible! Tom Okell and Rob Tijssen have also been a great help by sharing their expertise and code with me. Linqing Li and James have been valued partners at late-night and weekend scanning sessions for which I'm very grateful, especially because they were well aware that my "30 minute" scans are in practice often considerably longer! Thanks to them and the rest of the physics group for their company and friendship. Thanks to Wilfred Lam, Way Cherng Chen, Mike Germuska, Clark Lemke, Alex Gardener, Jennifer McNab, Molly Bright, Stuart Clare, Dan Gallichan, Jingyi Xie, Mike Kelly, Jamie Near, Dan Bulte, Sean Foxley and Nic Blockley for help, laughs, football, food, pub trips, ice hockey chats, music, conferences, cakes and conversation - it has been great fun working here.

Many other people around FMRIB have also given me valuable advice on pulse sequences, programming, stats, FSL, brain anatomy etc. Thank you to Chris Rodgers, Matt Robson, Jesper Andersson, Adrian Groves, Saad Jbabdi, Alek Petrovic, Reza Khorshidi, Stam Sotoripoulos, Gwenaelle Douad, Natalie Voets and Holly Bridge in particular. Thank you Marilyn Goulding and Sue Field for your

cheerfulness and help and for telling me when it's time to go home! Thank you Duncan Mortimer, Dave Flitney and Iain Wilson for repeatedly saving my laptop and data with ingenious fixes and workarounds!

Thank you to friends around Oxford and at home and everyone who has been to visit. Thanks to Oxford University, Magdalen, Jesus and Mansfield Road football clubs for letting me play and to the friends I've made there. Thank you to the Medical Research Council, Magdalen College and the Institute of Physics for funding I have had during my time here.

Most importantly, thank you to my parents, Stephen and Anna-Lena, and my brothers, William and James, for giving me every opportunity and all their support.

Contents

Preface

Declaration of the Contributions of Others

I would like to acknowledge the conribution of Dr. David Porter (Siemens Healthcare, Erlangen, Germany) who wrote and developed the original version of the source code for the rs-EPI pulse sequence, which formed the basis for the work in this thesis. Modifications were then made by myself to extend the scope of the original sequence, as described in this thesis. Some other sequence code, used in the implementation of the on-line cardiac reordering of shots in the 3D rs-EPI sequence, was derived from code written by Dr. Rob Tijssen (previously of FMRIB Centre, University of Oxford).

List of Publications Arising from this Thesis

Journal Articles

1. **R Frost**, DA Porter, KL Miller, and P Jezzard. Implementation and assessment of diffusion-weighted partial Fourier readout-segmented echo-planar imaging. *Magn Reson Med*, 68(2):441–451, Aug 2012.

2. **R Frost**, KL Miller, DA Porter, RH Tijssen, and P Jezzard. 3D multi-slab diffusion-weighted readout-segmented echo-planar imaging with real time cardiac-reordered k-space acquisition. Submitted to *Magn Reson Med*.

3. **R Frost**, P Jezzard, DA Porter, and KL Miller. Diffusion-weighted readout-segmented echo-planar imaging with a blipped-CAIPI simultaneous multi-slice modification. In preparation.

Conference Proceedings

1. **R Frost**, DA Porter, T Feiweier, and P Jezzard. Homodyne reconstruction of partial Fourier readout-segmented EPI for diffusion imaging. In *Proceedings of the 18th Annual Meeting of ISMRM*, Stockholm, Sweden, 2010 (abstract 1625).

2. RH Tijssen, SM Smith, P Jezzard, **R Frost**, M Jenkinson, KL Miller. Characterization and Correction of Physiological Instabilities in 3D FMRI. In *Proceedings of the 18th Annual Meeting of ISMRM*, Stockholm, Sweden, 2010 (abstract 1163).

3. **R Frost**, DA Porter, G Douaud, P Jezzard, and KL Miller. Reduction of diffusion-weighted readout-segmented EPI scan time using a blipped-CAIPI modification. In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 116).

4. **Frost R**, Miller KL, Porter DA, Tijssen RH, Jezzard P (2013). 3D multi-slab diffusion-weighted readout-segmented echo-planar imaging with real-time cardiac-reordered k-space acquisition. In *Proceedings of the 21st Annual Meeting of ISMRM*, Salt Lake City, USA, 2013 (abstract 3962).

5. **Frost R**, Jezzard P, Porter DA, Miller KL (2013). Simultaneous multi-slab acquisition in 3D multi-slab diffusion-weighted readout-segmented echo-planar imaging. In *Proceedings of the 21st Annual Meeting of ISMRM*, Salt Lake City, USA, 2013 (abstract 3907).

Patents

1. **R Frost**, P Jezzard, KL Miller and DA Porter. Diffusion-weighted magnetic resonance imaging using 3D mosaic segmentation and 3D navigator phase correction. US Patent 2012/0286777. China, Japan and Korea Patent, Pending.

Chapter 1

Introduction

This DPhil thesis concerns the development of a high-resolution diffusion-weighted pulse sequence for magnetic resonance imaging (MRI) of the brain. In neuroimaging, diffusion-weighted MRI is the only widely available clinical tool to diagnose the early stages of ischemic stroke. It is also the only way to non-invasively trace white matter connections in the brain. This chapter describes the motivation for the research and an outline of the thesis.

1.1 Motivation

Diffusion-weighted MRI generates image contrast based on the self-diffusion of water molecules. Each image is sensitised to diffusion in a chosen direction so the contrast is sensitive to the average displacement during the echo time of water molecules along that particular direction. The further molecules diffuse in the sensitisation direction, the lower the signal in the diffusion-weighted image. For example, in the brain, white and grey matter will generally have higher signal than cerebrospinal fluid (CSF), since there are more obstructions to movement of water molecules. This signal has been shown to change as the result of a stroke, as soon as minutes after the onset of the infarct. Ischemic stroke is the most common

neurological cause of severe disability and death [1] and it is estimated that in a typical patient 1.9 million neurons are lost every minute the stroke is untreated [2]. Using diffusion-weighted MRI it is possible to diagnose the location and the severity of the stroke almost immediately afterwards. Smaller strokes with more benign symptoms or transient ischemic attacks that do not cause lasting damage may be predictive of more damging strokes in the future. Often these small infarcts can not be seen at the resolution of standard diffusion-weighted images so a higher resolution image acquisition method would be valuable.

The directional dependence of the diffusion-weighted signal can also inform on the architecture of the brain. In white matter, there is less diffusion across the direction of the ordered fibre tracts than along them, so the diffusion is said to be anisotropic, whereas in CSF, the diffusion is approximately isotropic. Differences in signal can be observed by acquiring multiple images with diffusion encoding in different directions and the degree of anisotropy can be mapped with a technique termed diffusion tensor imaging [3]. Voxel-wise estimates which quantify the anisotropy can be used to detect changes in tissue microstructure due to pathology or normal development. By assuming that white matter tracts run in the direction of fastest diffusion, fibre pathways can be traced through the brain non-invasively. These white matter tracts transmit information between localised regions of specific function in grey matter cortex. Although the fibre tracing methods have limitations, knowledge of how these connections are wired is important for understanding of neurological function and can also have clinical applications [4].

The applications of diffusion-weighted imaging would benefit from improvements to image acquisition methods. The conventional diffusion imaging sequence encodes the image with a single-shot echo-planar imaging (ss-EPI) readout [5]. The resolution is limited to around $2\times2\times2$ mm voxels because the image

must be encoded while there is sufficient signal following a single excitation of the spins. Also, the long k-space readout time degrades the quality of the images, specifically through spatial distortions and blurring. Resolution and image quality can be improved by using multiple short readouts to encode the image. However, because diffusion imaging sequences are sensitive to motion on a micron scale, motion during the diffusion-weighting section will cause encoding inconsistencies between readouts that ruin the images if they are not properly accounted for. The ss-EPI sequence is not affected by this problem to the same degree because the whole image is acquired in one readout. So-called multi-shot sequences can be corrected using "navigator" techniques that effectively measure the motion at each shot, so that it can be accounted for in reconstruction. One sequence that has addressed this motion problem is readout-segmented echo-planar imaging (rs-EPI) [6], which has been implemented by Porter and co-workers and is the subject of the research in this thesis. This method enables robust high-resolution diffusion-weighted imaging with less distortion and blurring by virtue of the shorter readouts. However, apart from motion insensitivity, another factor in the continued success of ss-EPI is its rapid imaging speed, which allows acquisition of all slices of a whole brain volume in roughly 10 seconds. This is particularly important in DTI and related methods where 60 or more images with diffusion-weighting in different directions are typically acquired. Segmenting the image acquisition into different shots increases the scan time per volume (roughly by the number of shots) and therefore DTI studies can become prohibitively long. Before high-resolution diffusion imaging can be used in this context, methods for reducing the scan time are required.

Another limitation of the current diffusion imaging, is the reliance on 2D multi-slice sequences, which "excite" and then encode signal from 1.5-5 mm thick slices of the brain. These slices are shifted in separate excitations so that the

whole brain is covered. It is difficult to accurately excite slices much thinner than 1.5 mm thick and because signal-to-noise ratio scales with voxel volume, signal decreases with thin slices. 3D sequences, which excite thicker slabs of signal and then encode slices in the through-slab direction, can be used to image high-resolution slices. However, despite a large volume of previous research 3D diffusion-weighted sequences have not been adopted for standard use. The difficulties arise in accounting for motion in the slice direction when exciting a thicker slab because measuring the motion in all three dimensions is challenging in the short time available for signal encoding. However, recent research indicates that 3D diffusion-weighted imaging could be possible by limiting the thickness of the slabs.

1.2 Thesis Outline

This thesis describes the implementation of image acquisition and reconstruction methods to develop the rs-EPI sequence [6], which was shared with the University of Oxford as part of a research collaboration with Siemens Healthcare. Chapter 2 gives a brief background in the basics of MRI, fast imaging methods which are relevant to diffusion imaging, the origin of the diffusion-weighted signal and the image analysis methods used in stroke diagnosis and diffusion tensor imaging.

Chapter 3 describes the origin of artefacts in diffusion-weighted imaging and some correction strategies and assesses current diffusion imaging sequences and the rs-EPI sequence in particular.

Chapter 4 is the first research chapter on the reduction of the scan time in rs-EPI by reducing the number of acquired shots and reconstructing the missing data with a partial Fourier algorithm. The homodyne and projection onto convex sets (POCS) algorithms are compared relative to the original (full k-space) acquisition

in terms of image quality and in region of interest (ROI) analyses of estimates of diffusion parameters.

Chapter 5 presents another approach to reducing scan time where the number of excitations required to encode a full image is reduced. A recently published method is adapted for rs-EPI to simultaneously excite and readout multiple slices together. The images from the different slices are collapsed on top of each other in the reconstructed images but a reconstruction that takes advantage of multiple receive channels can be used to separate the data into separate slices. The modified rs-EPI sequence is used to acquire tractography data for comparison with ss-EPI data and trace-weighted images which are used for clinical stroke diagnosis.

In Chapter 6 the 2D multi-slice rs-EPI sequence is adapted for 3D multi-slab imaging. A model for the pulsatile motion of the brain during the cardiac cycle is used to simulate the effects of motion on image encoding. Using this simulation framework, different approaches to correction of motion artefacts in 3D imaging are tested. The simulations suggest that a correction can be applied using 2D navigator data and that an encoding scheme that is synchronised to the cardiac cycle reduces artefacts. The simulation results are confirmed in experiment and a simultaneous multi-slab acquisition is presented, which builds upon the work in Chapter 5.

Chapter 7 summarises the research in the thesis and discusses possible future directions of research.

Chapter 2

Background

2.1 Principles of MRI

This section gives a brief introduction to magnetic resonance imaging (MRI) that is based on descriptions in standard texts [7, 8, 9, 10]. The magnetic resonance (MR) phenomenon is explained using classical physics with the postulate that sub-atomic particles have an intrinsic quantised angular momentum, or "spin". A rigorous derivation requires detailed quantum mechanics but a classical overview is sufficient for understanding of the macroscopic MR physics.

Observation of the electron spin (in 1924 by Stern and Gerlach) [11] and measurement of nuclear magnetic moment (in 1938 by Rabi and co-workers) [12] led to the independent discoveries of nuclear magnetic resonance (NMR) by Bloch and Purcell [13, 14] in 1946, for which they received the 1952 Nobel Prize in Physics. Nuclei with an odd number of protons and/or neutrons have a net nuclear spin angular momentum and this, coupled with their charge, creates a nuclear magnetic moment. This magnetic moment is the origin of the NMR signal and the nuclei being studied are often referred to as spins. Hydrogen-1 (with a nucleus consisting of a single proton) is the most commonly used nucleus in NMR and

MRI experiments because it has a strong nuclear magnetic moment (strong signal) and is the most abundant in biological tissue (which is mostly water). NMR has allowed the study of the biochemistry of nuclei such as hydrogen-1, phosphorous-31 and carbon-13 *in vitro* and *in vivo*.

MRI was first described by Lauterbur and Mansfield in 1973 [15, 16]. Their innovation was to apply spatial localisation to NMR using linear magnetic field gradients. MRI is very powerful due to its flexible image contrast and experimental set up and, importantly, is non-invasive and safe. Despite the drawbacks of the expensive hardware and lengthy scan times in some situations, use of MRI is widespread and in 2003 Lauterbur and Mansfield shared the Nobel Prize in Physiology and Medicine.

In MRI, manipulation of the hydrogen-1 spins using the main magnetic field ($\mathbf{B_0}$), radiofrequency magnetic fields ($\mathbf{B_1}$) and linear gradient fields ($\mathbf{G}$) enables polarisation, excitation and spatial localisation, respectively. These concepts and the k-space formalism for data acquisition are discussed in the remainder of this section.

2.1.1 The MR Signal

When a hydrogen nucleus is placed in a magnetic field $\mathbf{B_0}$, which by convention is defined to be along the z direction, two energy levels are created by the quantum Zeeman effect. These energy states can be thought of as two possible orientations of the z component of the nuclear magnetic moment. The lower energy level corresponds to the z component of the spin being aligned with $\mathbf{B_0}$ and in the higher level the z component is anti-aligned with $\mathbf{B_0}$. These energy states are separated by ΔE:

$$\Delta E = \gamma \hbar B_0 \tag{2.1}$$

where γ is the gyromagnetic ratio of hydrogen (equal to $2\pi \times 42.577 \times 10^6$ rad/s/T), $\hbar$ is Planck's constant divided by 2π, and B_0 is the magnitude of $\mathbf{B_0}$. The ratio of the anti-parallel population ($n_\downarrow$) to the parallel population ($n_\uparrow$) is given by the Boltzmann distribution:

$$\frac{n_\downarrow}{n_\uparrow} = e^{-\Delta E/k_B T} \tag{2.2}$$

where k_B is Boltzmann's constant and T is the absolute temperature. The lower energy parallel state is favoured but at 3 T and physiological temperatures $k_B T \gg \gamma \hbar B_0$ so there is only an excess of ~10 spins per million in the parallel state. The net magnetisation $\mathbf{M}$ of a sample in a magnetic field is the vector sum of all the spins' magnetic moments $\boldsymbol{\mu}$ ($\mathbf{M} = \sum \boldsymbol{\mu}$) so most of the magnetic moments cancel out but the imbalance in the spin populations causes a small polarisation. The magnitude of $\mathbf{M}$ is approximately:

$$M_0 \approx \rho_0 \Big(\frac{\gamma \hbar}{2}\Big)^2 \frac{B_0}{k_B T} \tag{2.3}$$

where ρ_0 is the number of hydrogen nuclei per unit volume. This net magnetisation is detectable as signal due to the high number density of hydrogen in water-based samples such as biological tissue.

At equilibrium $\mathbf{M}$ and $\mathbf{B_0}$ are aligned but if $\mathbf{M}$ is perturbed to point in a different direction it will precess around $\mathbf{B_0}$. This can be seen from the classical definition of torque (rate of change of angular momentum, $\mathbf{S}$) on the individual magnetic moments $\boldsymbol{\mu} = \gamma \mathbf{S}$ in the $\mathbf{B_0}$ field:

$$\frac{\mathrm{d}\mathbf{S}}{\mathrm{d}t} = \boldsymbol{\mu} \times \mathbf{B_0} \tag{2.4}$$

Multiplying by γ and summing all the magnetic moments gives:

$$\frac{\mathrm{d}\boldsymbol{\mu}}{\mathrm{d}t} = \boldsymbol{\mu} \times \gamma \mathbf{B_0} \tag{2.5}$$

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \mathbf{M} \times \gamma \mathbf{B_0} \tag{2.6}$$

The solutions to Equations 2.6 and 2.5 are for $\boldsymbol{\mu}$ and $\mathbf{M}$ to precess around $\mathbf{B_0}$ at the Larmor frequency:

$$\nu_0 = \frac{\gamma}{2\pi} B_0 \tag{2.7}$$

where ν_0 ($= \omega_0/2\pi$) is the frequency in Hz, $\gamma/2\pi$ is in Hz/T and B_0 is in T. For example, at 3 T, individual hydrogen-1 spins and their resultant magnetisation will precess at ~128 MHz. This is also the frequency associated with the energy level separation, ΔE, in Eq. 2.1. If the magnetisation is excited with a radiofrequency (RF) magnetic field, $\mathbf{B_1}$, at the Larmor frequency a transition can be induced between the energy states. From a classical perspective, the $\mathbf{B_1}$ field exerts a torque on $\mathbf{M}$ so that it rotates away from the equilibrium z direction whilst precessing. The magnetisation will be perturbed from the z direction and precess in the x-y plane. The angle of perturbation or the flip angle, θ, depends on the magnitude, B_1, and duration, τ, of the RF pulse:

$$\theta = \gamma \int_0^\tau B_1(t)\,\mathrm{d}t \tag{2.8}$$

The z axis and the x-y plane are often referred to as the longitudinal axis and the transverse plane, respectively. After $\mathbf{B_1}$ is switched off, the transverse component of $\mathbf{M}$, M_{xy}, can be detected as a small electromotive force in a conducting coil placed around the sample via Faraday's law of induction. The same RF coil used to transmit the $\mathbf{B_1}$ field can be switched to receive mode to detect the signal. The maximum signal will be seen after a 90° RF pulse and no signal is generated

by a 180° pulse because there is no component of magnetisation in the transverse plane. The longitudinal component of $\mathbf{M}$, M_z, subsequently regrows with time constant, T_1. The precession of any transverse component will generate an oscillating signal called a free induction decay (FID) which decays under an exponential envelope. This decay is caused by small fluctuations in the magnetic field which means that the spins precess at frequencies that are slightly different from the intended Larmor frequency. Therefore some spins will be ahead of the resonance frequency and some spins will be behind, so the transverse magnetisation vector loses phase coherence or dephases. Following an excitation pulse, the FID will decay according to the apparent transverse relaxation time constant, T_2^*, due to variations in field caused by imperfections in the main magnetic field as well as random fluctuations caused by the motion of the molecules. The dephasing caused by inhomogeneities in the main magnetic field can be refocused with a 180° pulse, in which case the signal will depend on the transverse relaxation time constant, T_2. The phenomenological relaxation terms, T_1 and T_2, and the evolution of $\mathbf{M}$ following RF excitation are summarised by the Bloch equation:

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \mathbf{M} \times \gamma\mathbf{B_0} - \frac{M_x\hat{\mathbf{i}} + M_y\hat{\mathbf{j}}}{T_2} - \frac{(M_z - M_0)\hat{\mathbf{k}}}{T_1} \tag{2.9}$$

where $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$ and $\hat{\mathbf{k}}$ are unit vectors in the x, y and z directions, respectively.

2.1.2 Image Formation

In MRI, the NMR signal is spatially resolved by applying linear magnetic field gradients, $\mathbf{G}$, on top of the main magnetic field B_0. Gradients are used to selectively excite slices of magnetisation and for the subsequent frequency- and phase-encoding of the signal required in Fourier imaging. These concepts are explained below, after a brief description of how gradients modify the precessional

frequency of spins in order to determine their position.

Under the influence of a gradient, the main magnetic field becomes a function of position, **r**:

$$\mathbf{B} = (B_0 + \mathbf{G} \cdot \mathbf{r})\,\hat{\mathbf{k}} \tag{2.10}$$

where

$$\mathbf{G} = \begin{pmatrix} G_x \\ G_y \\ G_z \end{pmatrix} = \begin{pmatrix} \mathrm{d}B_z/\mathrm{d}x \\ \mathrm{d}B_z/\mathrm{d}y \\ \mathrm{d}B_z/\mathrm{d}z \end{pmatrix} \tag{2.11}$$

It is important to note that the magnetic field still points in the z direction but the field strength varies with position. Hence the Larmor frequency of precession becomes a function of position. For example, in the presence of a gradient in the x direction, G_x the Larmor frequency is:

$$\nu(x) = \frac{\gamma}{2\pi}(B_0 + G_x x) = \nu_0 + \frac{\gamma}{2\pi} G_x x \tag{2.12}$$

The time domain signal under the influence of a gradient can be Fourier transformed [17] to obtain the signal contributions as a function of frequency. The signal can then be localised by using Eq. 2.12 to map frequency to position. Typically the maximum gradient strength on a clinical scanner is 40 mT/m, which creates a frequency variation of 1.703 MHz/m or a bandwidth of 170.3 kHz across a 10 cm object.

Slice Selection

The first step in MR image formation is often excitation of a thin slice (1.5-6 mm thick) of magnetisation through the sample. This is achieved by playing out a gradient in the direction perpendicular to the slice during the RF excitation pulse. The RF pulse is designed so that its excitation (or frequency) profile is as close to

a rect or "top hat" function as possible so that a sharp section of magnetisation is excited. For small flip angles, there is an approximate Fourier relationship between the excitation profile and the temporal envelope of the RF pulse, so a sinc shaped RF pulse should excite a rect function. However, the sinc shape must be truncated to limit the duration of the RF pulse (to limit the echo-time and phase accrual artefacts during the pulse) so the resulting excitation profile is not an ideal rect function, as shown in Fig. 2.1a. The time bandwidth product

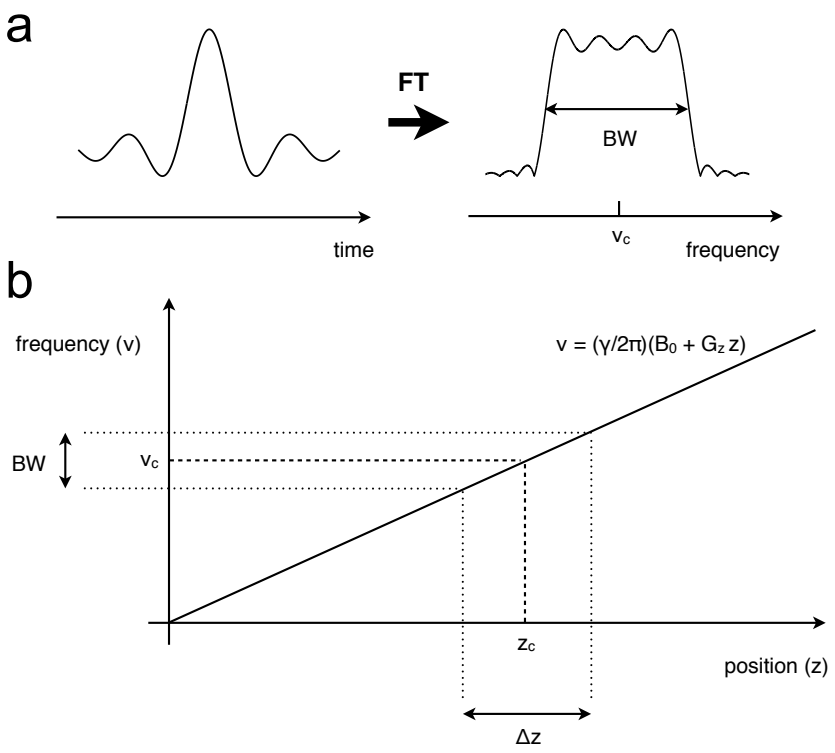

Figure 2.1: (a) Illustration of the approximate Fourier relationship between the RF pulse and the excited slice profile. The RF bandwidth (BW) is approximately the full width at half maximum of the slice profile. (b) Illustration of the conversion from frequency to position. The slice-select gradient amplitude G_z determines the slice position z_c and thickness Δz of a RF pulse with carrier frequency ν_c and bandwidth BW.

(TBWP), a dimension-less parameter of a RF pulse (the product of its duration and RF bandwidth) is a measure of the quality of the slice profile. For a sinc pulse the TBWP is proportional to the number of zero crossings in the RF envelope. Higher TBWP pulses excite sharper slices but for a fixed pulse duration they

require higher RF amplitude (B_1) as well as stronger gradients (which are limited by hardware) to achieve the same slice thickness. The peak B_1 is limited by the maximum voltage that can be played out in the transmit coil. If the requested B_1 amplitude exceeds the limit then the RF pulse will be "clipped" i.e., only the maximum available will be played out. Therefore to ensure that the desired RF envelope can be achieved, a trade-off between slice profile and pulse duration is necessary. Typical excitation pulses are in the range of several milliseconds and have a TBWP of about 5.

The carrier frequency (ν_c) of the RF pulse defines the slice position and the combination of the RF bandwidth (BW) and the gradient strength (G_z) determines the slice thickness, as shown in Fig. 2.1b.

Fourier Imaging

Fourier imaging is the basis for the zeugmatography [18] and spin-warp or 2DFT sequences [19]. The projection reconstruction method, as used in x-ray computed tomography scans [20], was the first MRI reconstruction but Cartesian Fourier transform methods are now more common.

The use of a gradient to separate signal from different locations in one dimension was mentioned previously. This is usually called "frequency-encoding" and the gradient is referred to as the readout gradient. If the readout gradient is applied along x, the 1D profile obtained after Fourier transformation of this signal is the image collapsed along the y direction. To achieve resolution in the orthogonal y direction this acquisition can be repeated but with incremental amounts of "phase-encoding" before the readout. The phase-encoding gradient is played out along y so that the different readouts have different phase variation as shown in Fig. 2.2. With a complete set of phase-encoded readouts a 2D image of the slice can be obtained via a 2D Fourier transform. This approach is called Fourier

imaging and can be understood by considering the signal under the influence of gradients.

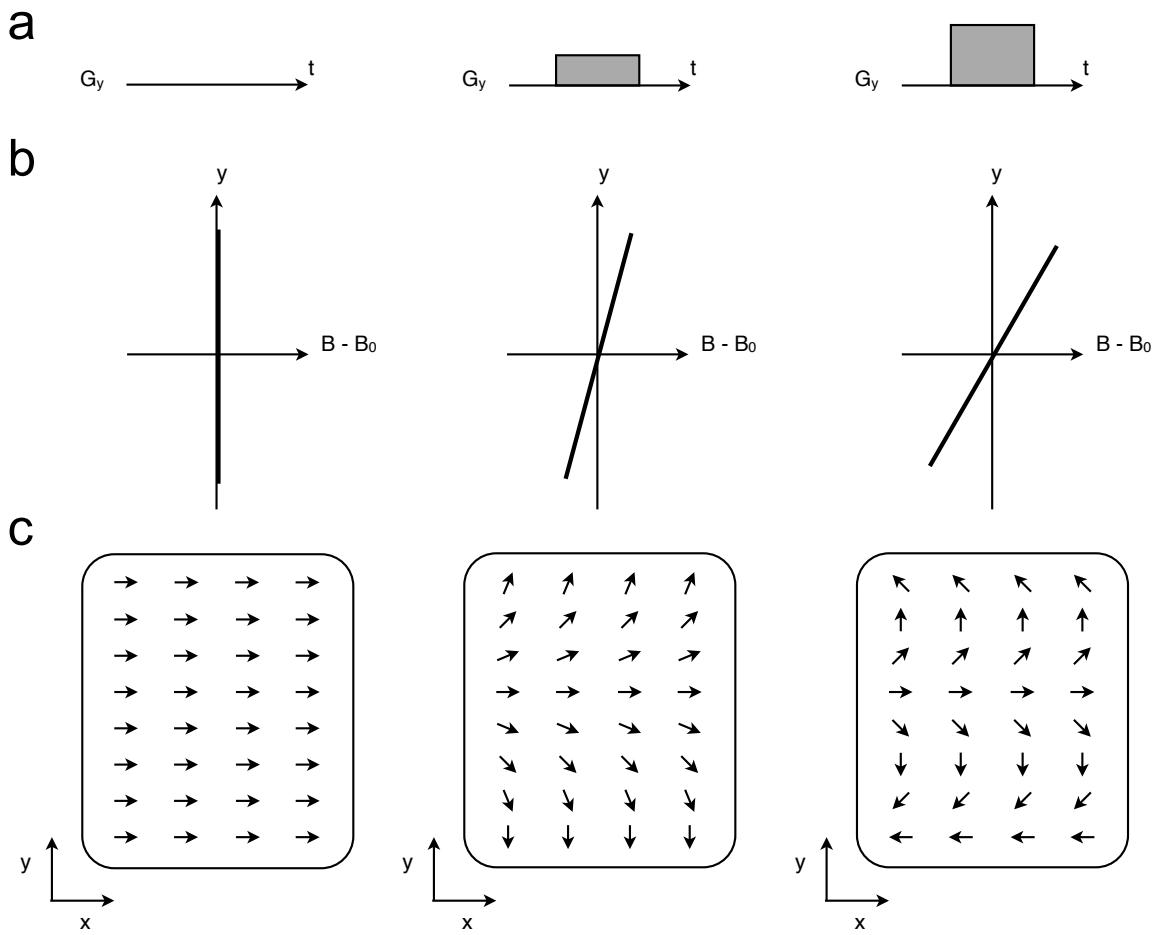

Figure 2.2: Demonstration of phase variation generated by a phase-encoding gradient, G_y, for three different gradient amplitudes: (a) plot of time-varying gradient waveform. Phase accumulation is proportional to the area of the gradient. (b) Variation in magnetic field, B, with position, y during the gradient. The magnetic field is given by $B = B_0 + G_y y$. (c) Phase of the transverse magnetisation. Spins that experience a stronger magnetic field accumulate more phase and the result is a variation in phase in the y direction. Spins at gradient isocentre (y = 0) are not influenced by the gradient.

The MR signal, demodulated at the Larmor frequency, is proportional to the spin density at each voxel location, $\rho(x, y, z)$, multiplied by the complex phase (relative to the Larmor frequency, ν_0) of the voxel. After slice selection (with rephasing so that all spins the slice have common phase), the signal in a slice of

magnetisation centred at z_0 with thickness Δz is given by:

$$S(t) \propto \int_x \int_y \left[\int_{z_0 - \frac{\Delta z}{2}}^{z_0 + \frac{\Delta z}{2}} \rho(x, y, z) \mathrm{d}z \right] e^{-i\phi(x,y,t)} \, \mathrm{d}x \, \mathrm{d}y \tag{2.13}$$

where $i = \sqrt{-1}$ and the accumulated phase, ϕ, is the integral over time of $\nu(x, y, z, t)$ relative to ν_0, which can be related to the applied gradients, $\mathbf{G}$ (using Equations 2.7 and 2.10):

$$\phi(x, y, t) = 2\pi\gamma \int (\nu(x, y, t) - \nu_0) \, \mathrm{d}t = \gamma \int (G_x x + G_y y) \, \mathrm{d}t \tag{2.14}$$

The gradient area can be defined generally as:

$$\mathbf{k}(t) = \gamma \int \mathbf{G} \, \mathrm{d}t \tag{2.15}$$

so that the signal equation can be expressed as a function of $\mathbf{k}(t)$:

$$S(k_x, k_y) \propto \int_x \int_y \rho(x, y, z_0) e^{-i2\pi(k_x x + k_y y)} \, \mathrm{d}x \, \mathrm{d}y \tag{2.16}$$

This equation (where the integral over z has been suppressed) has the same form as the 2D Fourier transform of the spin density. This shows that the signal is acquired in the Fourier domain of the image, which is referred to as k-space. The imaging gradients are used to control the frequency- and phase-encoding of the image and thereby define the so called k-space trajectory [21]. As shown in Fig. 2.3, the signal collected during a readout gradient, G_x, consists of a line of k_x points and readouts with different amounts of G_y phase-encoding correspond to lines of k-space data at different k_y values. Cartesian k-space trajectories are the most common but other sampling schemes such as radial and spiral are possible. Non-Cartesian trajectories can have advantages such as more efficient coverage

of the high signal region at the centre of k-space but require re-sampling onto a regular grid before the Fourier transform, which can be time consuming in image reconstruction, and the image artefacts are generally more complicated.

The inverse relationship between image space and k-space means that the image field of view (FOV) and resolution are determined by the spacing of k-space points (Δk) and the extent of k-space (k^{max}), respectively. The time between samples during the readout is called the dwell time, Δt, which is the inverse of the rate or frequency bandwidth of the sampling. For a constant readout gradient of magnitude G_x the k-space spacing can be calculated from Eq. 2.15:

$$\Delta k_x = \gamma G_x \Delta t \tag{2.17}$$

The Nyquist sampling theorem dictates that in order to represent a frequency spectrum of bandwith BW without aliasing, the signal must be sampled with:

$$\Delta t = \frac{1}{\mathrm{BW}} \tag{2.18}$$

The frequency range of an object of length L_x in a gradient G_x is:

$$\mathrm{BW} = \frac{\gamma}{2\pi} G_x L_x \tag{2.19}$$

Substituting Equations 2.18 and 2.19 into Eq. 2.17 gives:

$$\Delta k_x = \frac{2\pi}{L_x} \tag{2.20}$$

The extent of k-space sampled (i.e. from $-k_x^{max}$ to $+k_x^{max}$) is related to the image resolution δ_x by:

$$2k_x^{max} = \frac{2\pi}{\delta_x} \tag{2.21}$$

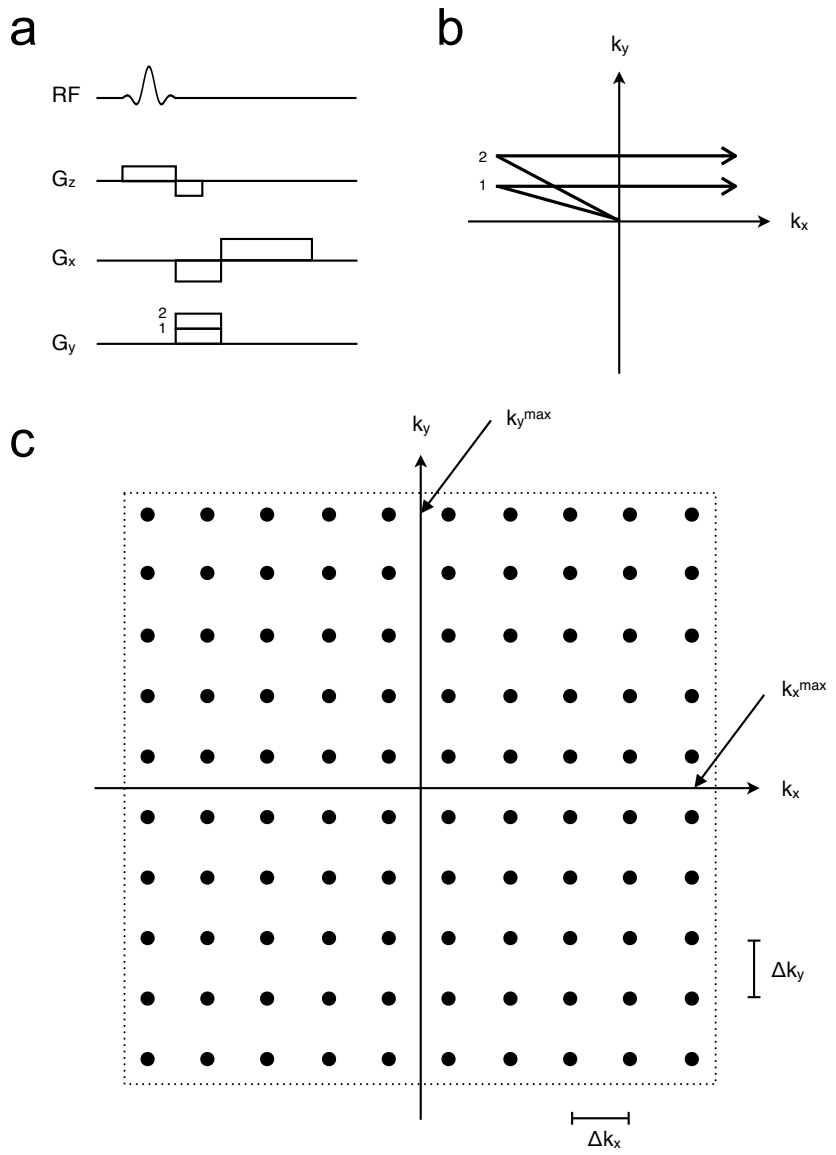

Figure 2.3: Illustration of (a) the basic 2D Fourier transform pulse sequence and (b) the corresponding k-space trajectory. The two different G_y amplitudes correspond to the labelled phase-encode lines which are at different k_y values. This pulse sequence is repeated, with varying G_y amplitudes until all the required k-space lines have been acquired. (c) Illustration of k-space sampling. The spacing between points, Δk, and the extreme k-space value, k^{max}, determine the image FOV and resolution, respectively.

Hence, given a desired FOV and image resolution, the appropriate k-space sampling and gradient parameters can be calculated. Analogous considerations apply to the y direction to determine the spacing between and number of phase-encode lines. In 3D imaging where a thicker slice or slab of magnetisation is excited, phase-encoding can be extended to the z direction so that a 3D k-space matrix of k_x, k_y and k_z points is acquired.

Image Contrast

As mentioned above, the MR signal is proportional to the number density of water protons. The water content of grey matter, white matter and CSF are relatively similar [22] so images weighted by the proton density do not provide much contrast between tissue types. However, contrast can be generated by choosing sequence parameters to accentuate differences in tissue relaxation times, T_1, T_2 and T_2^*. Other strategies include injecting a contrast agent to increase relaxation in particular tissues, exploiting physiological changes such as blood oxygenation that affect relaxation and sensitising the signal to diffusion or perfusion.

T_1-weighted images can be acquired by choosing a short repetition time (TR) between subsequent RF pulses. When TR is shorter than the T_1 value of a tissue, the longitudinal component of magnetization, M_z, does not fully recover. Hence, signal from different tissues will be weighted by their T_1 relaxation times. T_2-weighting is controlled by the choice of echo time (TE). In spin-echo sequences, the decay of the transverse component of magnetisation, M_{xy}, in different tissues depends on the T_2 value. By setting a long TE, short T_2 species will have low signal compared to those with longer T_2 values. Similarly, in gradient-echo sequences long TE will generate T_2^* contrast between tissues. The TE is defined as the time from the excitation pulse to the centre of the spin- or gradient- echo, or to the acquisition of the centre of k-space if multiple phase-encode lines are readout in

each TR.

2.2 Fast Imaging

This section discusses the echo-planar imaging (EPI) k-space readout, which is used extensively to acquire ~2 mm isotropic full brain images in ~3-10 s, and parallel imaging techniques, which can improve resolution, scan times and image artefacts.

2.2.1 Echo-Planar Imaging

Echo-planar imaging [23, 24] is a technique which can be used to acquire a 2D image in less than 0.1 s after a single RF excitation pulse. Multiple gradient-echoes are formed by repeatedly reversing the readout gradient, in contrast with RARE sequences [25], which generate a train of spin-echoes with RF refocusing pulses. Phase-encode blip gradients are applied between each echo so that each readout corresponds to a different line in k-space. The first implementation by Mansfield and Pykett [23, 24] used a constant phase-encode gradient during the whole readout thereby creating a zig-zag trajectory through k-space. Several modifications to the sequence were introduced, such as the blipped phase-encode gradients [26], time-reversal of data acquired during the reversed readout gradient [27] and non-linear time sampling during the ramps of the trapezoidal gradients to ensure regularly spaced k-space samples [27]. Later implementations encoded the complete k-space in order to allow reconstruction of more robust magnitude images [26, 28, 29]. Previously, when half of k-space was acquired, the imaginary part was discarded in the Margosian reconstruction [30] so the real images were sensitive to phase errors caused by field inhomogeneities and dephasing by flowing blood [29]. The encoding time is limited due to T_2^* decay, such that in order

to traverse k-space more quickly improvements in hardware were necessary to allow stronger gradients with faster switching. More powerful gradient current amplifiers were used and shielding coils were designed for the main gradient coils [31, 32, 33] to remove fringe fields, which generate unwanted eddy currents in metallic structures around the magnet and limit switching.

A major benefit of EPI is the reduced sensitivity to subject motion, which is effectively frozen during the timescale of the ~50 ms acquisition (unless the motion is of a similar time-scale). The high temporal resolution and robustness to motion artefacts made it suitable for diffusion imaging [34, 5, 35] and functional MRI [36, 37, 38, 39] as well as perfusion, cardiac and real-time imaging. For a more detailed account of the origins, methods and applications of EPI, see the book by Schmitt, Stehling and Turner [40].

EPI Pulse Sequence

The basic EPI pulse sequence for one slice and the corresponding k-space trajectory are shown in Fig. 2.4. After the slice selective excitation, pre-phasing gradi-

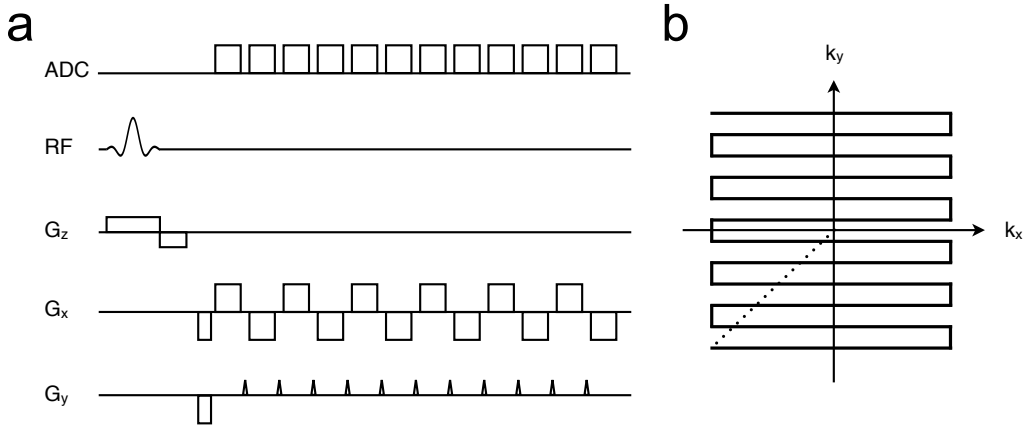

Figure 2.4: (a) Illustration of the EPI pulse sequence and (b) the corresponding k-space trajectory. After pre-phasing to $-k_x^{max}, -k_y^{max}$ (dotted line), the G_x readout gradient and the ADC are turned on to acquire one line of k-space. The G_y blip moves the k-space position to the next phase-encode line and the process is repeated until the whole of k-space' has been traversed.

ents are played out on the readout (k_x) and phase-encode (k_y) axes to move to the $-k_x^{max}, -k_y^{max}$ starting position. The positive readout gradient (G_x) moves the acquisition forward in k-space and the analog to digital converter (ADC) is opened to acquire the signal. After this first readout the ADC is turned off and a positive phase-encode gradient (G_y) "blip" is played out to move to the next k-space line. *k*-Space lines are acquired in this fashion during the train of positive and negative gradient-echoes until the full k-space path shown in Fig. 2.4b has been traversed. The number of echoes (or the number of positive and negative readout gradients) is called the echo train length (ETL) and is equal to the number of acquired k-space lines. The time between echoes is called the echo-spacing (t_{ES}) and is usually ~1 ms. The encoding time is limited by the T_2^* decay of the signal so the objective is to acquire as many k-space lines as possible, which can be achieved by sweeping through k-space quickly to minimize the echo-spacing. Minimizing the echo-spacing also reduces bandwidth-related artefacts (see Off-Resonance and Blurring sections below). However, the speed of k-space traversal is limited by the gradient slew rate, gradient amplitude, receiver bandwidth and the k-space matrix size along the readout direction (n_x). An accumulation of gradient events with high slew rates can cause unpleasant peripheral nerve stimulation (PNS) in subjects as well as gradient coil heating. The readout time per line T_{RO} is given by the product of the dwell time per k-space point (Δt) and n_x:

$$T_{RO} = n_x \Delta t = \frac{n_x}{\mathrm{BW}} = \frac{2\pi n_x}{\gamma L_x G_x} \tag{2.22}$$

using Equations 2.18 and 2.19. Signal-noise-ratio (SNR) is inversely proportional to $\sqrt{\mathrm{BW}}$ but having a high BW can also increase SNR by reducing signal loss due to T_2^* decay. A compromise between the echo-spacing, SNR and achievable gradient amplitude results in readout bandwidths in the range of 100 kHz to 1

MHz.

To increase efficiency of the encoding, signal can be sampled during the finite G_x rise-times [41, 42] as well as during the constant part (often called the "flat-top") of the trapezoid waveform. If ramp-sampling is used with linear temporal sampling, the k-space samples will be unevenly spaced and lead to a ripple artefact in the x direction of the image, if left uncorrected. The distance between k-space samples is the integral of the gradient over time, as shown in Eq. 2.15, so the irregular k-space samples can be "regridded" to a regularly spaced grid with interpolation techniques [43]. Alternatively, non-linear time sampling can be used to compensate for the changing k-space velocity during the gradient ramp [27]. Sinusoidal readout gradient waveforms can also be used in conjunction with either constant time sampling and regridding or with non-linear time sampling.

Ghosting

In EPI, even lines are acquired under positive G_x and odd lines under negative G_x. The data from odd lines are time-reversed [27] because they have negative velocities in k-space. The odd and even lines have slightly different characteristics due to background susceptibility gradients, imperfections in the gradient pulses or any gradient-induced eddy-currents. There can be relative shifts between k-space lines caused by image-space phase gradients (Fourier shift theorem). The mismatch between lines causes a faint repeat of the image shifted by half a FOV in the phase-encode direction, which is a called a Nyquist or N/2 ghost. A calibration scan where the central k-space line is repeatedly acquired under positive and negative readout gradients is used to determine the constant (ϕ_0) and linear (ϕ_1) phase difference (and higher order terms if required) between the odd and even lines. This phase difference can be added to even lines in reconstruction [44].

Off-Resonance

Spins that are not precessing at the Larmor frequency are said to be off-resonance. Two types of off-resonances generate artefacts in EPI: chemical shifts in frequency arising from different tissue types (e.g. fat) or main field inhomogeneity due to an imperfect shim or susceptibility differences, which are common at boundaries between tissue and air or bone, for example, above the frontal sinuses. The slow traversal of k-space in the phase-encode direction in EPI allows off-resonance phase to accumulate during the readout, which results in a constant shift of fat signal and geometric distortions due to main field inhomogeneity.

The EPI readout has high bandwidth in the readout direction $\mathrm{BW_{RO}}$ and low bandwidth in the phase-encode direction $\mathrm{BW_{PE}}$. The bandwidth is a measure of the k-space velocity and can be calculated from the inverse of the time between k-space samples in the particular direction:

$$\mathrm{BW_{RO}} = \frac{1}{\Delta t} \tag{2.23}$$

$$\mathrm{BW_{PE}} = \frac{1}{t_{ES}} \tag{2.24}$$

where Δt and t_{ES} are the previously mentioned dwell time and echo-spacing, respectively. Main field inhomogeneity ΔB_0 and chemical shift frequency offsets between different tissues $\delta\nu$ generate image encoding artefacts in the phase-encode direction because they modify the precession frequency. The frequency at position x under the influence of a gradient G_x is given by:

$$\nu = \frac{\gamma}{2\pi}(B_0 + G_x x + \Delta B_0) + \delta\nu \tag{2.25}$$

The position is given by:

$$x = \frac{\frac{2\pi\nu}{\gamma} - B_0}{G_x} - \frac{\Delta B_0}{G_x} - \frac{2\pi\delta\nu}{\gamma G_x} \tag{2.26}$$

which is different from the expected position that does not account for the ΔB_0 and $\delta\nu$ terms. The distortions due to field inhomogeneity and chemical shift in the readout direction are given by:

$$d_{cs} = \frac{2\pi\delta\nu}{\gamma G_x} = \frac{\delta\nu L_x}{\mathrm{BW_{RO}}} \tag{2.27}$$

$$d_{dist} = \frac{\Delta B_0}{G_x} = \frac{\gamma\Delta B_0 L_x}{2\pi\mathrm{BW_{RO}}} \tag{2.28}$$

Analogous expressions apply for the phase-encode direction. This equation shows how the low bandwidth ($\mathrm{BW_{PE}}$) in the phase-encode or k_y direction in EPI generates chemical shift and distortion artefacts (the artefacts in the readout direction are negligible in comparison). At 3 T, the frequency difference between water and fat $\delta\nu$ = 440 Hz and for Δt = 10 μs and t_{ES} = 1 ms, $\mathrm{BW_{RO}}$ = 100 kHz and $\mathrm{BW_{PE}}$ = 1 kHz. In this case the fat signal is shifted by $0.0044L_x$ and $0.44L_y$, i.e. a negligible shift in the readout direction and almost a half FOV in the phase-encode direction. The fat signal can be suppressed with a fat saturation module before the excitation pulse which excites at the frequency of fat and then dephases the signal with strong "crusher" gradients. Alternatively, a spatial-spectral pulse can be used, where both the slice location and the spectral content are defined, so that only water spins within the slice of interest are excited [45]. In spin-echo sequences the polarity of the 180° slice selection gradient can be reversed relative to that of the 90° so that only on resonance water spins are refocused [46, 47].

The component of ΔB_0 due to magnet imperfections varies slowly across the imaging volume and can be reduced by shimming. However the susceptibility

contribution to ΔB_0 can have sharp changes that can not be shimmed. The degree of stretching and compression (or signal pile-up) varies throughout the image and is worst in the vicinity of susceptibility boundaries, which limits EPI around the sinuses and lower in the brain where air/tissue boundaries exist. Distortions can be corrected to some extent in post-processing by acquiring a B_0 field map to characterise the distortions [48]. Also, the direction of the distortions is reversed if k-space is acquired with the opposite phase-encoding polarity (i.e. from $+k_y^{max}$ to $-k_y^{max}$ rather than from $-k_y^{max}$ to $+k_y^{max}$) so if two images are acquired with opposite phase-encoding they can be used to derive a correction [49, 50]. Large changes in resonance frequency within a voxel can also reduce signal through intravoxel dephasing. This means SNR in some voxels can increase by choosing a smaller slice thickness even though SNR generally scales with voxel volume. Susceptibility effects increase at higher field strengths so areas of distortion which are manageable at 3 T may become a severe challenge for EPI at 7 T.

Diffusion imaging with EPI introduces another source of distortions as a consequence of eddy currents that are generated during the strong diffusion-weighting gradients. These distortions are discussed in Section 3.1.3 in the next chapter on diffusion imaging sequences.

Blurring

The decay of signal during the long EPI readout weights the k-space data resulting in blurring in the phase-encode direction [51]. It is this blurring and the resulting signal loss which effectively limits the EPI encoding time. In gradient-echo (GRE) EPI the signal decay has the form:

$$S(t) \propto e^{-t/T_2^*} \tag{2.29}$$

where t = 0 at the excitation. In spin-echo (SE) EPI the signal has the form:

$$S(t) \propto e^{-(t+TE)/T_2} e^{-|t|/T_2'} \tag{2.30}$$

where t = 0 at the echo-time (TE). T_2^* is the combined effect of the transverse relaxation constant T_2 and the local main field inhomogeneity constant T_2':

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{2.31}$$

If the inhomogeneity is Lorentzian, T_2' is approximately

$$\frac{1}{T_2'} \sim \frac{\gamma \Delta B_0}{2} \tag{2.32}$$

In SE EPI the 180° pulse refocuses the T_2' decay so the signal peaks at TE according to Eq. 2.30. This signal modulation in k-space degrades the point spread function (PSF) in image-space. The PSF is a mathematical description of how a point is depicted by an imaging method. The effective phase-encode resolutions in GRE and SE sequences are given by Equations 2.33 and 2.34 [7]:

$$\Delta y'_{GRE} = \frac{\sqrt{3}}{\pi} \Delta y \frac{T_{acq}}{T_2^*} \tag{2.33}$$

$$\Delta y'_{SE} = \frac{\sqrt{3}}{\pi} \Delta y \frac{T_{acq}}{T_2} + \frac{1}{\pi} \Delta y \frac{T_{acq}}{T_2'} \tag{2.34}$$

where Δy is the nominal phase-encode resolution and T_{acq} is the duration of the readout. Improved shimming and higher phase-encode bandwidth reduce the inhomogeneity and hence signal loss and blurring.

Variants of EPI

Single-shot EPI (ss-EPI) is the most widely used form of EPI because of its excellent time resolution and robustness to motion. However, as discussed in the above the long readout time and low phase-encode bandwith cause artefacts and limit spatial resolution. The readout duration can be shortened to reduce TE by omitting k-space lines (as shown in Fig. 2.5a) and reconstructing them with a partial Fourier algorithm [30]. The SNR of a partial Fourier acquisition is lower than a

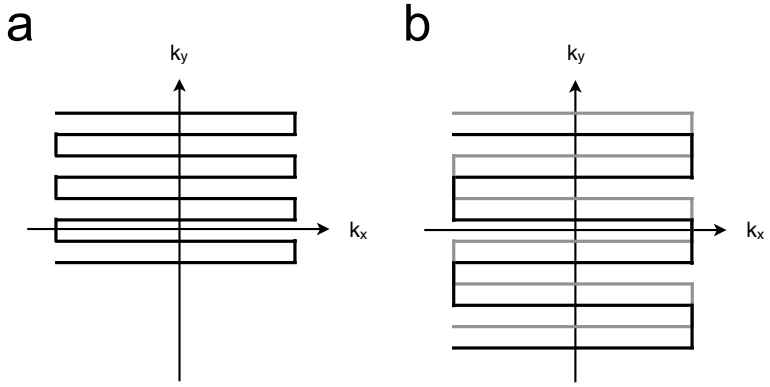

Figure 2.5: k-space trajectories of (a) partial Fourier and (b) interleaved EPI. In (a) the missing k-space lines are estimate with a partial Fourier reconstruction algorithm. In (b), the black and grey trajectories correspond to two different interleaves acquired in different TR periods. The k-space data are combined to form the fully sampled matrix.

full k-space acquisition by approximately $\sqrt{n_y^{PF}/n_y^{Full}}$ where n_y^{PF} and n_y^{Full} are the number of k_y lines in a partial Fourier and full k-space acquisition, respectively. However this loss in SNR is usually equal to or outweighed by the SNR gained from a lower TE, particularly in diffusion-weighted protocols where TE can be lengthy because of the diffusion preparation module before the readout. Partial Fourier reconstructions exploit the conjugate symmetry of k-space and the most common algorithms are homodyne [52] and projection onto convex sets (POCS) [53].

The spatial resolution can be increased if data from multiple k-space shots are combined. Interleaved EPI [54, 55] where k-space lines are skipped in each shot

(as shown in Fig. 2.5b) allows higher phase-encode bandwidth or lower effective echo-spacing because the effective temporal separation between k-space lines is smaller than in an otherwise equivalent single shot acquisition. This reduces off-resonance artefacts such as chemical shift and distortion but the scan time is increased by the number of shots. The longer scan times also make multi-shot EPI more sensitive to motion between the different excitations.

The EPI readout can be used in 3D volume excitation sequences, such as 3D SE EPI and magnetisation-prepared 3D GRE EPI, where the second phase-encode dimension (k_z) is acquired in multiple shots. Echo-volume imaging (EVI) [56, 57, 58] is a true 3D sequence which acquires multiple k_z planes after a single excitation. The T_2^* signal decay and gradient limitations only allow relatively modest acquisition matrices (~64×64×8) but by under-sampling k-space and estimating the missing data with parallel imaging techniques (see Section 2.2.2) higher resolution (3 mm isotropic 64×64×48 or 2 mm isotropic 128×128×16) has been demonstrated [59]. The advantage of the EVI approach is that all slices have the same phase behaviour and are acquired at the same time (rather than with delays as in multi-slice sequences), which can be important in functional studies. However, because encoding in the k_z dimension is usually slower than k_y encoding in 2D EPI, the distortion and blurring artefacts in the z direction can be severe.

2.2.2 Parallel Imaging

Parallel imaging is a technique for accelerating image acquisitions by reducing the number of phase-encoding steps. It is made possible by the encoding power of array coils [60] used for RF signal reception that arises from the spatial sensitivity variation between individual elements of the RF coil array. The originally intended application for array coils was to increase SNR in large field field of view (FOV) images. The basic principle behind this concept (which is discussed

in Refs. [9] and [61]) is that a small and a large coil would see the same signal (assuming that they are both close to the object) but the noise (dominated by thermal noise in the object), which is integrated over each coil's sensitivity profile, would be higher in the larger coil. If the object is further away then the larger coil would detect a stronger signal so on balance the SNR is higher. Hence an array of independent coils around the object can limit the noise and detect approximately the same signal thereby increasing the measured SNR. To reduce coupling between coils in the array they have individual receiver systems so separate coil images can be reconstructed. The optimal combination of coil images is weighted by the coil sensitivity profiles but in the high SNR regime sum-of-squares is a good approximation [60].

In Cartesian acquisitions that are accelerated with parallel imaging, the same extent of k-space (the same spatial resolution) is sampled but the distance between phase-encode (k_y) lines is increased by a factor R_{PE}, which reduces the encoding time by the same factor. The regular under-sampling of k-space reduces the size of the FOV by a reduction factor R_{PE} and leads to aliasing (signal wrap-around) when the Nyquist condition is not met because the reduced FOV is smaller than the object. The aim of a parallel imaging reconstruction is to remove image aliasing from under-sampled data. The SNR of the reconstructed images is less than a fully sampled image by a factor $\sqrt{R_{PE}}$ due to the reduction in Fourier averaging and by a factor g that describes the parallel imaging performance of the array coil configuration. The geometry-factor (or g-factor), which varies spatially in the image depending on the degree of under-sampling, the array coil and the reconstruction details, is lower when there is a large sensitivity variation between aliased pixels, which makes it easier to separate the signal into un-aliased components. As R_{PE} is increased, the distance in image space between aliased pixel locations, and hence the sensitivity variation between them, decreases. Thus the

g-factor is higher and residual aliasing and spatially varying noise amplification is evident in the reconstructed images. Noise amplification is usually most pronounced at the centre of the array coil where there is least sensitivity variation and where most of the aliased signal is collapsed onto itself. Clinical quality images can be obtained using standard array coils for acceleration factors of R_{PE}~2-3.

Partially parallel imaging (the "partially" is usually dropped but was initially used to distinguish combined coil and Fourier encoding from pure coil encoding [62]) was first demonstrated in a phantom in 1993 [63] and then *in vivo* in 1997 with the k-space-based SMASH technique [64]. The most commonly used techniques on clinical scanners are the SENSE [65] and GRAPPA [66] algorithms, which operate in the image and k-space domains, respectively. These algorithms are less computationally expensive and more robust than alternative methods so they can be used for fast reconstruction of diagnostic quality images.

SENSE

In SENSE (SENSitivity Encoding) signals overlap from R_{PE} equally spaced locations in the image. This is demonstrated for R_{PE}=2 in the following equations:

$$\begin{aligned} S_1(x,y) &= C_1(x,y)\rho(x,y) + C_1\left(x, y+\frac{FOV}{2}\right)\rho\left(x, y+\frac{FOV}{2}\right) \\ S_2(x,y) &= C_2(x,y)\rho(x,y) + C_2\left(x, y+\frac{FOV}{2}\right)\rho\left(x, y+\frac{FOV}{2}\right) \end{aligned} \tag{2.35}$$

where S_j is the aliased image from the j^{th} coil in the aliased image, C_j is the coil sensitivity (which can be estimated experimentally) and ρ is the desired spin density of the object. A diagram of this equation from Larkman and Nunes [61] is shown in in Appendix A as Fig. A.1, which demonstrates the role of the coil

sensitivity profiles in aliasing. Equation 2.35 can be generalised to matrix form:

$$\mathbf{S} = \mathbf{C}\boldsymbol{\rho} \tag{2.36}$$

This equation can be solved by inverting the matrix $\mathbf{C}$ as long as the number of coils is greater than or equal to the acceleration factor and the coil profiles are sufficiently different from each other (so that $\mathbf{C}$ is not rank deficient). The estimate of the spin density in the SENSE reconstruction is given by:

$$\boldsymbol{\rho}' = (\mathbf{C}^H \boldsymbol{\Psi}^{-1} \mathbf{C})^{-1} \mathbf{C}^H \boldsymbol{\Psi}^{-1} \mathbf{S} \tag{2.37}$$

where $\boldsymbol{\Psi}$ is the noise correlation matrix between receiver channels. Pruessman *et al.* [65] derived the SNR at pixel p of the SENSE parallel imaging reconstruction SNR_p^{PI}:

$$\mathrm{SNR}_p^{PI} = \frac{\mathrm{SNR}_p^{full}}{g_p \sqrt{R}} \tag{2.38}$$

where the g-factor at pixel p is given by:

$$g_p = \sqrt{[(C^H \Psi^{-1} C)^{-1}]_{p,p} [(C^H \Psi^{-1} C)^{-1}]_{p,p}} \tag{2.39}$$

The g-factor (which is greater than or equal to one) is a common measure of the performance of a parallel imaging reconstruction that takes into the acceleration factor and the coil array configuration. Estimated g-factor maps also show how the noise varies spatially in the reconstructed images. In highly accelerated acquisitions noise enhancement in images can be noticeable, particularly in regions close to the centre of the array coil where there is typically less sensitivity variation between aliased pixels. Due to the g-factor penalty, an acceptable acceleration factor is in general much smaller than the number of receiver coils in an array. In fact, even with highly dedicated 32-channel array coils, achieving good image

quality in EPI with R = 4 is challenging.

GRAPPA

The GRAPPA (GeneRalised Auto-calibrating Partially Parallel Acquisition) reconstruction estimates the missing k-space lines so that the data can be Fourier transformed to an un-aliased image. To estimate a missing data point in the accelerated (under-sampled) k-space data, surrounding points from all N coils are multiplied by a set of linear weights which relate these points to the missing data point. To determine the relationship between missing and acquired points, the set of weights is trained on fully-sampled calibration data using a least-squares fit to all the available calibration data (also called auto-calibration signal (ACS)). The number of surrounding points is defined by the GRAPPA kernel. The fitting approach in GRAPPA is different to the original k-space-based technique SMASH, which linearly weights the coil sensitivity profiles to synthesise spatial harmonics that mimic phase-encoding. The fitting equation for a typical 5×4 kernel is given by:

$$S_j(k_x, k_y) = \sum_{l=1}^{N} \sum_{b_x=-2}^{2} \sum_{b_y=-2}^{1} n_{j,l,m}^{b_x,b_y} S_l(k_x + b_x \Delta k_x, k_y + (b_y R_{PE} + 1)\Delta k_y) \tag{2.40}$$

which relates the k-space signal S_j in coil j to the points defined by the kernel indices b_x, b_y in all of the N coils. There are (R_{PE}-1) missing lines between acquired lines and m denotes the different kernels for each set of missing lines when R_{PE}>2. Equation 2.40 can be written in matrix form:

$$\boldsymbol{S}_j = \boldsymbol{n}\boldsymbol{S} \tag{2.41}$$

In the ACS data, the target points, $\boldsymbol{S}_j$, and the source points, $\boldsymbol{S}$, are known so $\boldsymbol{n}$ can be estimated:

$$\hat{\boldsymbol{n}} = (\boldsymbol{S}^H \boldsymbol{S})^{-1} \boldsymbol{S}^H \boldsymbol{S}_j \tag{2.42}$$

where $\hat{\boldsymbol{n}}$ is the linear least-squares estimate of the weight set $\boldsymbol{n}$. Hence, in the accelerated data, missing lines, $\boldsymbol{S}'_j$, can be estimated by applying $\hat{\boldsymbol{n}}$ to the acquired lines, $\boldsymbol{S}'$:

$$\boldsymbol{S}'_j = \hat{\boldsymbol{n}} \boldsymbol{S}' \tag{2.43}$$

Using this procedure fully-sampled k-space can be reconstructed for all coils in the array, which can then be Fourier transformed and combined using a sum of squares reconstruction. The auto-calibration reference data [67] is so called because the parallel imaging reconstruction parameters are determined directly from the k-space data. The ACS data is usually acquired at the start of the acquisition but can be integrated into the accelerated data to reduce sensitivity to motion after the calibration. In EPI a variable density acquisition where the centre of k-space is fully sampled and the outer parts are accelerated would have image artefacts due to variable off-resonance phase accrual during the echo train. Therefore separate calibration scans for all slices are acquired in one TR period at the start of the sequence.

GRAPPA can be understood by considering the effect of the coil sensitivities in k-space. In the image domain, a single coil image is the spin density multiplied by the coil sensitivity. The k-space signal seen by a single coil is a convolution of the Fourier transform of the spin density with the Fourier transform of the coil sensitivity. The fitting of a GRAPPA kernel to the k-space data tries to estimate the effect of this convolution, which in theory should be the same throughout k-space. Small kernels can be used because the coil sensitivities usually contain a small number of Fourier terms (slowly varying in image space translates to

closely spaced in k-space). To improve the accuracy of the fit the weight set is estimated at multiple locations in the ACS data (and the reconstructed line can be an average of different kernels or blocks). The g-factor in GRAPPA reconstructions can be estimated from the GRAPPA weights [68] or using a gold standard "pseudo multiple replica" method which is applicable to any parallel imaging reconstruction [69].

The SENSE reconstruction is an exact solution to the system of equations so if the coil sensitivities were perfectly known, then SENSE would perform better than GRAPPA, which is a least-squares fit. However experimental errors in estimating coil sensitivities and motion between calibration and acquisition degrades SENSE reconstructions. GRAPPA is generally more robust because auto-calibration removes the need for sensitivity maps and the least-squares fit can limit the errors due to motion.

Applications

Parallel imaging has had a significant impact on MRI and it is now routinely used to reduce scan time in conventional 2DFT sequences or to shorten the effective echo-spacing and the readout duration in EPI to reduce off-resonance and blurring artefacts, respectively. Phase-encode under-sampling can be applied independently to the k_y and k_z dimensions [70], which has been of great benefit in time-consuming 3D sequences.

Array coils and parallel imaging reconstruction can also be used to un-alias multiple slices which have been excited and encoded simultaneously, as demonstrated by Larkman *et al.* [71]. The benefit of this technique is that slices can be acquired simultaneously so fewer excitations are necessary to acquire the same number of slices. Hence TR can be reduced, which is especially valuable in multi-slice protocols with large numbers of thin slices that require a long TR (e.g. in

diffusion imaging and functional MRI). The reduction in TR translates directly into a reduction in scan time by the same factor. Also, as demonstrated by Fig. 3.8 in Chapter 3, decreasing TR will increase SNR efficiency until TR is ~T_1.

Multiband RF pulses are used to localise signal from multiple slice locations before image encoding with conventional readouts. The result of this is that the simultaneously excited slices are on top of each other or multiplexed in the reconstructed images. If there is sufficient sensitivity variation in the slice direction of the array coil, parallel imaging reconstructions can be used to separate the multiplexed images into separate images of each slice. The g-factor related SNR loss in these reconstructions can be minimized by using RF- [72] or gradient-induced [73, 74] phase to shift the slices relative to each other within the FOV. This means that slice-aliased pixels have additional in-plane sensitivity variation (as well as the sensitivity variation in the slice direction) between them and therefore the g-factor is reduced. This is important because often array coils have less sensitivity variation in the slice direction than in-plane and the slice FOV can be small, which means that aliased pixels can be close together. Larkman *et al.* [71] used a SENSE reconstruction to un-alias the slices and since then a combination of SENSE and GRAPPA [75, 76] and a GRAPPA-based algorithm [74] have been demonstrated.

2.3 Diffusion Imaging

An MRI experiment can be sensitised so that image contrast depends on the degree of self-diffusion of water molecules. Diffusion is caused by the collisions between molecules in liquids or gases and the random motion was first described by Robert Brown [77]. An expression for the root-mean-squared (RMS) displacement (r_{RMS}) of a particle in a random walk during time t was derived by Einstein

[78, 79] using the diffusion coefficient (D) in Fick's law:

$$r_{RMS} = \sqrt{2Dt} \tag{2.44}$$

Hahn showed that the spin-echo signal was reduced by the dephasing due to diffusion [80] and Carr and Purcell applied magnetic field gradients during the sequence to measure the effect [81]. Stejskal and Tanner built on this approach by using finite gradient pulses either side of the 180° refocusing pulse [82]. In a spin-echo sequence these pulses are of the same polarity but the same diffusion weighting can be achieved in a gradient-echo sequence (without a 180° pulse) by using a bipolar pair of gradients, that have opposite polarity, as shown in Fig. 2.6. When a gradient ($\mathbf{G}$) is applied along a particular direction, spins acquire a phase proportional to their position ($\mathbf{r}$) along the gradient direction and the gradient area:

$$\phi = \gamma \int_0^t \mathbf{G} \cdot \mathbf{r} \, \mathrm{d}t \tag{2.45}$$

During a bipolar gradient pair (shown in Fig. 2.6), static spins will be dephased

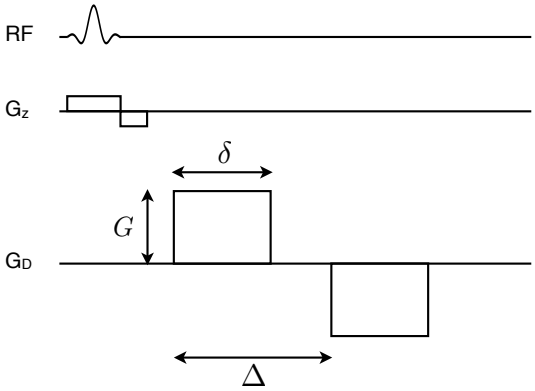

Figure 2.6: Illustration of diffusion-weighting gradients, G_D. The diffusion-weighted signal is acquired after this preparation module.

by the first lobe and refocused by the second lobe so they will accumulate no phase. Diffusing spins will have components of motion along the gradient direc-

tion and therefore their phase is not refocused by the second gradient lobe (because they experience a changing field during the gradients due to their motion). Thus a voxel that contains diffusing spins will experience signal attenuation from destructive interference if the spins move far enough to accumulate significant dephasing. The signal attenuation is described by the following equation:

$$S = S_0 e^{-bD} \tag{2.46}$$

where S_0 is the signal in the absence of diffusion and b is a measure of the degree of diffusion weighting. The "b-value" is calculated from the gradient amplitude (G) and timing:

$$b = \gamma^2 \int_0^{TE} \left| \int_0^t G(t')\mathrm{d}t' \right|^2 \mathrm{d}t \tag{2.47}$$

In a diffusion-weighted (DW) image, regions of fast diffusion, such as cerebrospinal fluid (CSF), have low signal intensity. A map of the diffusion coefficient can be estimated from the ratio of an image with diffusion gradients (finite b-value) and without (b = 0). From Eq. 2.46, the diffusion coefficient is given by:

$$D = \frac{ln\frac{S(b)}{S(b=0)}}{-b} \tag{2.48}$$

The diffusion coefficient of water in biological tissue is on the order of 10^{-3} mm^2/s, which is roughly a third of the value in pure water. So according to Eq. 2.44, a spin in tissue will diffuse ~10 μm in 50 ms. The b-value is usually chosen to be ~1000 s/mm^2 so that the product bD in Eq. 2.46 is on the order of 1 and the signal is attenuated by a factor e^{-1}. The diffusion weighting is adjusted by varying the gradient amplitude (G) and duration (δ) and the separation between lobes (Δ), which is called the diffusion time. If the finite gradient rise times are

neglected, the b-value of the gradients shown in Fig 2.6 is given by:

$$b = (\gamma G \delta)^2 \left(\Delta - \frac{\delta}{3} \right) \tag{2.49}$$

It is important to minimize TE so that signal loss due to T_2 decay (in spin-echo diffusion) is minimized. Therefore the diffusion preparation module should be as short as possible so typically the gradient duration and separation are minimised and the amplitude is maximised. The maximum gradient amplitude on clinical scanners is usually ~40 mT/m. High gradient slew rates can be used to quickly ramp up to the maximum amplitude but these increase gradient-related eddy currents that can cause artefacts and image distortion, especially in diffusion imaging (see Section 3.1.3), and can also induce peripheral nerve stimulation.

In biological tissue, the presence of membranes and other cellular structures decreases diffusion [83] and the reduction can be greater in some directions than others. For this reason, the apparent diffusion coefficient (ADC) in a particular direction is measured because the experiment depends on the direction of the diffusion-encoding gradient. The direction of diffusion sensitisation is the vector sum of the gradient amplitudes along the three gradient axes so the amplitudes can be chosen to generate weighting in the desired direction. In some tissue types in the brain, such as grey matter, the ADC is approximately the same in all directions hence diffusion can be described by a single ADC and the diffusion is said to be roughly isotropic. In white matter, diffusion appears anisotropic because water molecules are able to move more freely along the fibres than across them. In an anisotropic medium, diffusion can be represented by a second order tensor:

$$D = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix} \tag{2.50}$$

The tensor is a symmetric 3×3 matrix so it is described by six independent elements and can be thought of an ellipsoid whose three axes are the root-mean-squared diffusion lengths. If the coordinate system (x, y, z) is chosen to coincide with the principal axes of the diffusion tensor then the off-diagonal elements of the matrix are zero and the diagonal elements are the three diffusion coefficients along the primary axes. In a diagonalised matrix, the diagonal elements are the eigenvalues and the unit vectors describing the coordinate system are the eigenvectors.

In the remainder of this section, the two main branches of diffusion imaging are described: diffusion-weighted imaging (DWI), which estimates the mean diffusivity independent of angle and has been used extensively in diagnosis of ischemic stroke [84], and diffusion tensor imaging (DTI) [3], which is a technique for estimating the diffusion anisotropy that has many applications in neuroscience. In this thesis we refer to protocols that estimate the (non-directional) mean diffusivity as DWI and all multi-direction diffusion-weighted acquisitions (for tensor measurements, tractography etc.) as DTI. The standard method for acquiring diffusion-weighted data is to use a spin-echo sequence with a ss-EPI readout. Other sequences can be used but the major benefits of ss-EPI are the relative insensitivity to motion artefacts and the imaging speed. However, as discussed in Section 2.2.1, the resolution of ss-EPI images is limited by geometric distortions and other artefacts (current techniques acheiving ~2 mm isotropic at 3 T). The choice of pulse sequence and readout module is discussed in detail in Chapter 3.

2.3.1 Diffusion-Weighted Imaging

Diffusion-weighted MRI is used routinely for imaging the early stages of ischemic stroke, which is the most common neurological cause of severe disability and death [1]. Reduced diffusion in ischemic brain tissue was first demonstrated

in cats minutes after occlusion of the middle cerebral artery [84] and these results were later reproduced in humans [85]. The mechanism of the reduction of diffusivity is usually attributed to movement of extracellular water to intracellular water as a result of cell swelling in early ischemia [83]. Combination of diffusion-weighting with EPI [5] made robust DWI of stroke possible. Strokes can be divided into major, minor and "lacunar" strokes. It is important to identify patients who have suffered a major stroke after occlusion of large cerebral arteries because they are at risk of severe functional impairment or death if not treated quickly. Recanalisation treatments such as thrombolysis to dissolve blood clots are often successful in improving major strokes. Minor and lacunar strokes, which are caused by occlusions in branching arteries and small penetrating vessels, respectively, usually have good outcomes and mild symptoms. Transient ischemic attacks are brief clinical manifestations of brain ischemia caused by vascular insufficiency. They do not necessarily lead to lasting damage but are often an early indicator of more serious future strokes.

Ischemic tissue is characterised by hyperintense signal on raw diffusion-weighted images or low signal on spatial maps of the ADC. If an image with diffusion-weighting in one direction is acquired then the resulting voxel-wise ADC estimates are dependent on the orientations of white matter tracts relative to the gradient direction. This directional dependence can be removed by acquiring three DW images (at the same b-value) sensitised in orthogonal directions. These directions are usually the readout, phase-encode and slice-select directions denoted by x, y and z, respectively, but can be any three orthogonal directions. The signal

intensities for x, y and z diffusion weighting are given by:

$$S_x = S_0 e^{-bD_{xx}} \tag{2.51}$$

$$S_y = S_0 e^{-bD_{yy}} \tag{2.52}$$

$$S_z = S_0 e^{-bD_{zz}} \tag{2.53}$$

Note that the signal and diffusion coefficients are functions of position within each image (i.e., $S_x = S_x(x, y)$, $S_0 = S_0(x, y)$ and $D_{xx} = D_{xx}(x, y)$ etc.) but the spatial dependence is omitted here for clarity. The geometric mean of these three images generates an image weighted by the trace of the diffusion tensor (D_{trace}, which is the sum of the diagonal elements of the matrix):

$$S_{xyz} = \sqrt[3]{S_x S_y S_z} = S_0 e^{-b(D_{xx}+D_{yy}+D_{zz})/3} = S_0 e^{-bD_{trace}/3} \tag{2.54}$$

These trace-weighted images are useful for diagnosing diffusion changes caused by stroke because they measure the diffusion strength without confounding anisotropy information. Rather than acquiring three separate images (S_x, S_y, S_z), trace-weighted images can be acquired in a single DW scan with diffusion preparations that generate isotropic weighting [86, 87]. These preparation modules can be used to avoid errors in registering the images due to inconsistent eddy current distortions in three images with diffusion-weighting along orthogonal axes. However, weighting along all three axes necessitates an increased TE and therefore SNR is reduced.

ADC maps can be estimated from two images with and without diffusion-weighting (as described above) or from a pixel-by-pixel fit to images acquired at multiple b-values. An average ADC map can also be estimated using a trace-weighted image and a b = 0 image. ADC maps have inverted contrast compared to DW images because fast diffusion areas have high signal. Also there is less

contrast between gray and white matter on ADC maps because their diffusivities are similar. In raw diffusion-weighted images, there is residual T_2 contrast between gray and white matter and this increases with TE. Chronic strokes can have high T_2 so remaining "T_2 shine through" can cause high signal on DW images. When diagnosing hyperintense lesions it is important to compare the ADC map to ensure that the signal is a result of reduced diffusivity, and therefore an acute stroke.

2.3.2 Diffusion Tensor Imaging

Diffusion tensor imaging (DTI) estimates the diffusion tensor at each voxel [3] using a minimum of seven MR images: one image without diffusion weighting and at least six images with weighting in non-collinear directions. More DW sampling orientations are typically acquired to reduce noise in the estimation of the tensor but this increases the scan time. At least 30 unique orientations, which are distributed isotropically over a sphere, are required to obtain robust estimates of anisotropy and tensor orientation [88].

The tensor can be diagonalised to determine the eigenvalues λ_1, λ_2 and λ_3:

$$D' = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} \tag{2.55}$$

The eigenvector corresponding to the largest eigenvalue is the tensor orientation or the direction of maximum diffusivity. The eigenvalues can be used to calculate maps of scalar diffusion measures [89, 90], such as mean diffusivity (MD):

$$\mathrm{MD} = \bar{\lambda} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \tag{2.56}$$

which is equivalent to the trace-weighted image mentioned above, and fractional anisotropy (FA):

$$\mathrm{FA} = \sqrt{\frac{3}{2}\frac{\sum_{i=1}^{3}\left(\lambda_i - \bar{\lambda}\right)^2}{\sum_{i=1}^{3}\lambda_i^2}} \tag{2.57}$$

FA values range from 0 (completely isotropic diffusion) to 1 (completely anisotropic diffusion). The principal eigenvector can be used to generate vector maps of the principal diffusion direction in each voxel. Often directionally encoded colour schemes are used to show left-right, anterior-posterior and superior-inferior directions in red, green and blue, respectively [91]. The eigenvectors provide information about the direction of the maximum diffusion within a voxel that has been used to track white matter fibres in the brain [92, 93, 94]. Fibre tracking or tractography traces these diffusion orientations to reconstruct a pathway, which assumes that the fibre orientation in a voxel can be approximated by the principal diffusion direction. However, the simple ellipsoid described by a second order tensor is generally not an accurate model for the underlying fibre structure. With this model, multiple fibre bundles in a single voxel can lead to incorrect estimates of fibre geometry. Tractography models can be adapted to distinguish parallel, crossing, "kissing" and bending fibres. Fibre bundles can be segmented, reconstructed and visualised but it is difficult to validate tractography estimates in areas of the brain where the connectivity is not well understood. Despite the limitations, tractography has been a powerful technique for neuroscience research because it is the only way to non-invasively visualise white matter connections between brain regions *in vivo*. It has improved models of normal white matter anatomy in humans and has been used clinically as marker for degeneration and could guide surgical interventions [4].

The tensor model is well known to break down under large numbers of directions, high b-values and other conditions. Many groups are exploring these

regimes because they may be informative about interesting aspects of tissue macro- or microstructre. In this thesis, we will restrict ourselves to more conventional regimes of signal behaviour and focus on techniques for improving image quality, as described in the following chapter.

Chapter 3

Diffusion Imaging Pulse Sequences

3.1 Pulse Sequence Considerations

3.1.1 Motion-Induced Phase Artefacts

Diffusion imaging pulse sequences are sensitive to microscopic diffusive displacements but, as will be discussed below, larger-scale motions lead to signal changes and artefacts. For *in vivo* neuroimaging, subject motion consists of rigid-body motion, such as translation and rotation of the head, and non-rigid pulsatile motion of brain tissue, which is related to the cardiac cycle. Involuntary movements on a millimetre scale, for example due to respiration, cannot be avoided and there is generally more movement during scans on acute patients and children. Brain tissue deforms due to an increase in blood volume during systole that remoulds the central and lower regions of the brain and pulls the brainstem downwards [95, 96, 97]. The brain slowly returns to the equilibrium position during diastole. Tissue and CSF velocities have been measured up to 3 mm/s with the main component being in the superior-inferior direction. An illustration of how the brain is re-shaped during systole deformation is shown in Figure 9 of Greitz *et al.* [95] and is reproduced in Appendix A as Fig. A.2.

In sequences where k-space is acquired in multiple shots, motion during the diffusion-encoding gradients causes phase changes in the signal that generate artefacts in the reconstructed magnitude images. Rigid-body translations and rotations cause constant phase offsets and linear phase ramps, respectively, in the image domain [98]. A Taylor expansion of displacement from gradient isocentre ($x(t)$) of an object with initial displacement x_0 and velocity v_0 is given by:

$$x(t) = x_0 + v_0 t + ... \tag{3.1}$$

Using Eq. 2.45 the phase accrued during a gradient $G(t)$ is therefore:

$$\phi = \gamma \int G(t)x(t)\mathrm{d}t = \gamma x_0 \int G(t)\mathrm{d}t + \gamma v_0 \int G(t)t\mathrm{d}t + ... \tag{3.2}$$

The net area of the diffusion gradients ($\int G(t)\mathrm{d}t$) is zero so translational motion with velocity v_0 causes a phase change $\phi = \gamma v_0 \int G(t)t\mathrm{d}t$. Figure A.3 (reproduction of Fig. 1 from Butts *et al.* [99]) in Appendix A shows how the linear relationship between phase and velocity results in a linearly increasing phase during rotation. The linear phase is in a direction perpendicular to the axis of rotation and the direction of the applied gradient. From the properties of the Fourier transform, constant phase in image space adds a constant phase offset in k-space and a linearly increasing phase in image space shifts the signal in k-space. When k-space is acquired in multiple shots, the rigid-body motion at each shot will in general be different and therefore the k-space phase and shift will be different at each shot. For example, if one phase encode line is acquired at each shot, each line will be phased and shifted by different amounts. This irregular k-space sampling results in ghosting in the reconstructed image and phase cancellation can cause signal loss. One way to avoid these artefacts is to acquire all of the k-space data in a single shot so that the whole trajectory will be phased and shifted by the

same motion-induced phase. Therefore the k-space sampling is consistent and the reconstructed images are not corrupted by motion. Magnitude images can be calculated and the phase can be discarded to avoid destructive phase interference if images are averaged.

In addition to rigid-body translations and rotations, a significant component of physiological motion is non-rigid and the diffusion gradients create such strong motion sensitivity that brain deformation during the cardiac cycle leads to non-linear phase across the brain. This results in a warping of k-space because each local part of the brain has a local phase ramp in the image domain, so different parts of k-space shift in different directions and by different amounts. This creates more complicated artefacts in multi-shot sequences, but just as above, single-shot acquisition captures a snapshot of the non-linear phase corruption, enabling one to simply discard the phase and keep the uncorrupted magnitude image. Hence single-shot sequences such as EPI (Cartesian or spiral) and FSE that acquire all the k-space data needed for a full image after each diffusion preparation are less sensitive to motion artefacts. The first single-shot *in vivo* diffusion studies were performed with EPI [34, 5, 35] and had substantially reduced artefacts compared to previous diffusion imaging.

However, single-shot sequences have limited achievable spatial resolution and as discussed in Section 2.2.1, EPI is susceptible to off-resonance artefacts and blurring. Parallel imaging can be used to reduce image distortions and blurring or improve resolution by factors of 2-3 [100, 101, 102] but further improvements in resolution require a segmented k-space acquisition. Motion artefacts can be mitigated in a multi-shot acquisition if the phase at each shot is measured and the data is corrected in reconstruction. These phase correction techniques are discussed below in Section 3.1.2. The large non-linear phase errors that occur during systole can also affect single-shot acquisitions. Local in-plane phase ramps can

be large enough to shift the corresponding k-space signal outside the acquisition window causing focal signal loss in images. Signal can also be attenuated by through-slice dephasing when the phase gradient is in the through-slice direction. Cardiac-gated acquisitions monitor the subject's pulse rate and image during diastole, the quiescent part of the cardiac cycle, to minimize these artefacts. Gating improves stability of diffusion imaging but extends the duration of the acquisition, particularly if the subject has a fast heart rate and therefore a short diastolic period. Waiting for diastole also introduces a variable TR and hence variable T_1 signal recovery but usually this is not significant for the 5-10 s TRs typically employed in multi-slice diffusion imaging.

In FSE methods such as PROPELLER [103] (which is discussed as an alternative to EPI-based DW imaging in Section 3.2), motion-induced phase causes failure of the Carr-Purcell Meiboom-Gill (CPMG) condition when refocusing pulses are not exactly 180°. The CPMG condition requires knowledge of the phase of the transverse magnetisation so that the phase of the refocusing pulses can be matched but in DW experiments the transverse magnetisation has a spatially varying and unknown phase. When the refocusing pulses are less than 180° then signal is attenuated due to interference between odd and even echoes. Phase cycling of the refocusing pulses can be used to produce a fairly stable echo train for around 16 echoes [103]. However it is more difficult to achieve stability in the echo-train at higher field where the B_1 field is less homogeneous and also the refocusing pulses add significantly to the total SAR, which is proportional to B_0^2 and B_1^2.

3.1.2 Navigator Correction

To achieve full FOV, high-resolution DW images a segmented k-space acquisition is usually employed. However, as discussed in the previous section, the variation

between shots of the phase of the transverse magnetisation causes artefacts when the shots are combined to form the complete k-space matrix. A demonstration of the changes in phase between five different shots is shown in Fig. 3.1a. The basic principle of navigated multi-shot schemes, which is illustrated in Figs. 3.1b and c, is to acquire a high-resolution k-space matrix in multiple shots while at each shot acquiring the central region of k-space. A 2D low-resolution estimate of the phase (such as those shown in Fig. 3.1a) can be computed from the "navigator" region of k-space shown in Fig. 3.1c. By measuring the phase of each shot in this way, the individual shots of high-resolution data are corrected by removing the phase corruption from each shot.

The first corrections for motion-induced phase errors in multi-shot sequences acquired an additional navigator line of k-space without phase-encoding, at each shot, that allowed measurement of the constant and linear phase in the readout direction [104, 98, 105]. This 1D navigator technique only allowed correction of phase offsets and linear shifts of k-space and was problematic for shifts that are not exactly aligned to the direction of the navigator. This was extended to two orthogonal 1D navigators for interleaved EPI [99] so that rigid-body correction could also be applied in the phase-encode direction. Later 2D navigator images were used to apply linear phase corrections in interleaved EPI [106, 107]. Unless cardiac gating was used these early techniques suffered from residual artefacts due to non-linear phase variation caused by pulsatile brain deformation. In interleaved EPI, the phase-encode separation Δk_y in each shot is increased (by skipping k-space lines) to reduce off-resonance artefacts and blurring. However, this means that each shot will be aliased if the FOV defined by Δk_y is smaller than the object. If the multi-shot scheme samples contiguous k-space points in the readout and phase-encode directions then full FOV, unaliased (low resolution) images can be reconstructed from the subset of k-space data sampled at each shot. Miller and

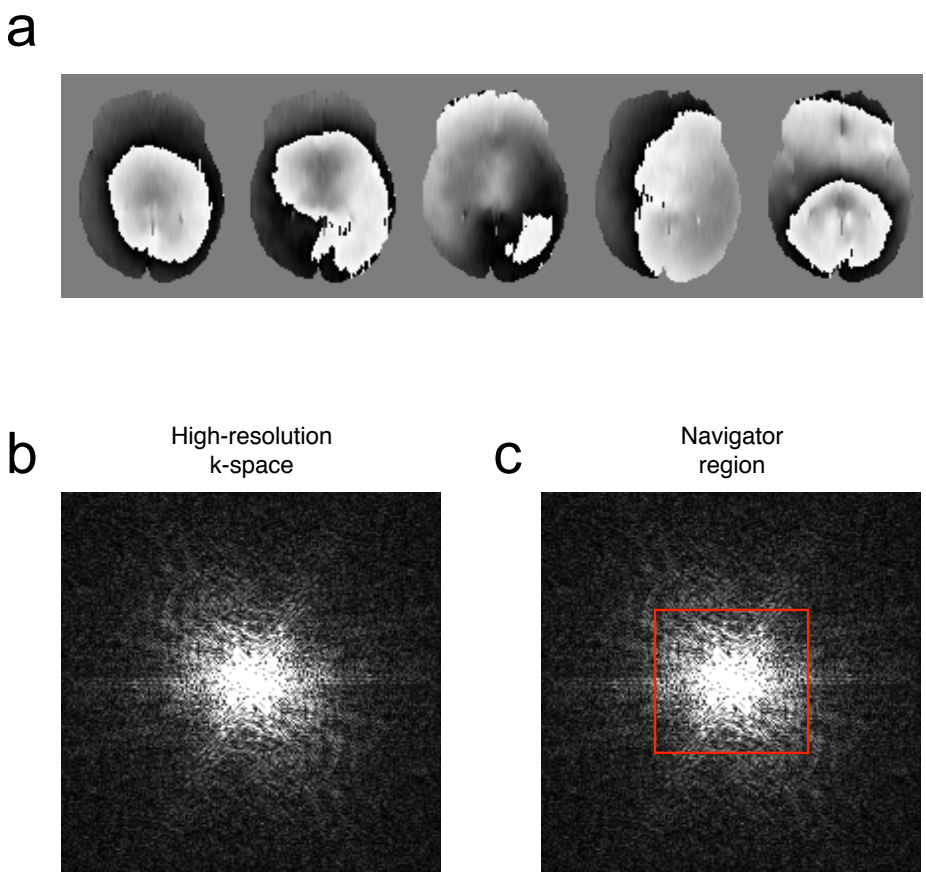

Figure 3.1: (a) Demonstration of the non-linear, spatial phase variation across the brain in five different shots. These images are from navigator data in a 2D multi-shot DW sequence (b) a complete high-resolution DW k-space map, which is acquired in multiple shots (c) the red box shows an example of the central region of k-space sampled by 2D navigator acquisitions. These low-resolution navigator acquisitions allow estimates of the spatial phase variation across the brain at each shot, such as those shown in (a). The figure of the navigator phase images in (a) is courtesy of David Porter.

Pauly derived a correction that uses a 2D navigator to apply a non-linear correction to each shot [108]. The full correction, which involves a computationally intensive matrix inversion, can be applied to aliased shots but the correction is simplified if each shot samples contiguous points in k-space and is therefore unaliased. The simplified version is termed the "refocusing" correction and can be implemented as an image-space subtraction of each navigator's phase estimate from its corresponding shot before the data from multiple shots are combined. The correction can also be applied in k-space as a deconvolution. The assumption is that the motion-induced phase caused by the diffusion gradients is slowly varying in image space so that it is well described by the low resolution navigator image. This assumption can be violated for severe motion corruption for example, when a shot is acquired during systole, but generally the correction performs well. It was also shown that the refocusing correction is improved when the acquisition is synchronised to the cardiac cycle. The motivation for cardiac synchronisation is so that adjacent regions of k-space have similar motion corruption to avoid discontinuities and periodicity of the signal phase and amplitude in the combined k-space data.

True 3D non-linear navigation has not yet been demonstrated due to the difficulty in acquiring a 3D navigator image with each shot in the limited encoding time. Unless navigator data can be incorporated into the final image, the efficiency of the pulse sequence decreases as more time is spent acquiring the navigator data. Pseudo-self-navigation was implemented in the 3D SSFP TURBINE sequence [109] by combining shots from similar points in the cardiac cycle to form a 3D navigator. However, the complicated motion dependency in SSFP, where motion phase is mixed over multiple shots, caused some residual artefacts. 2D corrections have been used in small slab 3D phase-encoded acquisitions with TSE and EPI readouts [110, 111, 112]. A linear correction using a 3D stack of spi-

rals navigator has been implemented but cardiac-gating was necessary to avoid systolic non-linear phase errors [113].

3.1.3 Eddy Current Distortions

EPI-based DW sequences are also affected by eddy currents in conducting parts of the scanner that are generated during the time varying part of magnetic field gradients. They are time varying magnetic fields that add to or subtract from the intended gradient fields. Eddy current fields interfere with the EPI encoding gradients and cause unwanted shifts of k-space. The rapidly switched gradients in the EPI readout generate eddy currents that can contribute to ghosting. In DW EPI, the ramps of the strong diffusion-encoding gradients cause eddy currents that persist during the readout and lead to image translations and shearing and scaling distortions. Eddy currents are proportional to the gradient slew rate, which has to be set high to minimize the duration of the diffusion encoding and hence TE. The distortion artefacts are typically different for different diffusion gradient directions and strengths so the mis-registrations cause problems for pixel-by-pixel analysis of diffusion tensor information.

Eddy current distortions can be minimised at source by means of a bipolar or doubly-refocused diffusion preparation [114], which is shown in Fig. 3.2a. The eddy currents produced by the bipolar scheme approximately cancel when a ramp up is closely followed by a ramp down because opposite polarity gradients have eddy currents with the opposite sign. The drawback of the bipolar scheme is that the minimum TE is increased compared to monopolar schemes. Post-processing techniques to correct the distortions can also be used [115, 116, 49, 50].

In the experiments in this thesis with readout-segmented EPI, a modified monopolar diffusion scheme is used where part of the second lobe is moved before the refocusing pulse and negated (see Fig. 3.2b) to achieve a shorter TE

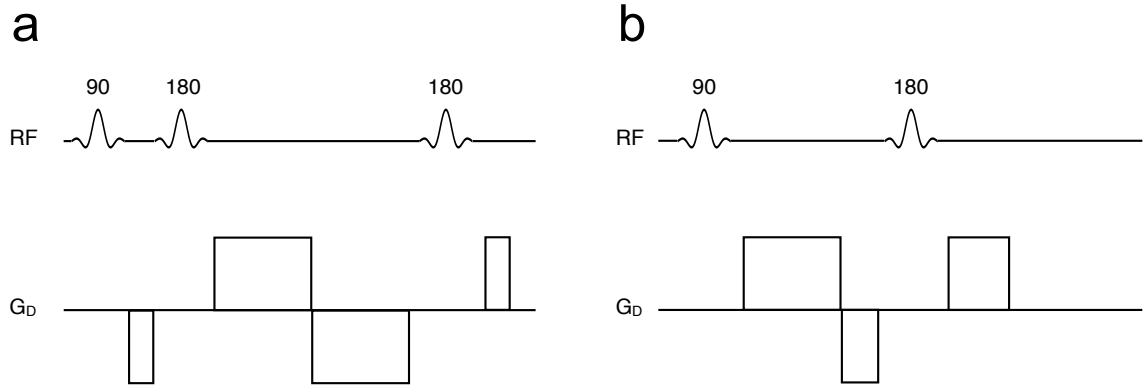

Figure 3.2: Illustration of diffusion-weighting gradient schemes. (a) bipolar scheme that has reduced sensitivity to eddy currents compared to monopolar schemes but lengthens TE. (b) a modified monopolar scheme where part of the second lobe is moved before the refocusing pulse and negated. This allows a shorter TE compared to conventional monopolar schemes.

[117]. The eddy current performance of the modified scheme is improved relative to conventional monopolar (Stejskal-Tanner) schemes because having a negative polarity gradient helps to cancel eddy currents. Readout-segmented EPI is less susceptible to eddy current distortions because it uses a short echo-spacing so it is acceptable to use the modified monopolar preparation rather than a bipolar preparation.

3.2 Pulse Sequences

A diffusion-weighted pulse sequence comprises a diffusion-weighting or preparation module to generate the diffusion contrast and a readout module where k-space is acquired. The preparation module is a set of pulses that imparts diffusion contrast to the signal between excitation and image acquisition and is to some extent, but not entirely, independent of the readout. Here, we briefly describe the most prominent techniques for creating diffusion contrast before considering classes of commonly-used readout strategies and their relative merits.

Spin-echo (SE) diffusion preparations have T_2 signal decay and therefore they are usually preferred to gradient-echo (GRE) preparations, which have faster T_2^* decay. Also, GRE in combination with an EPI readout will have worse off-resonance artefacts than SE EPI because the inhomogeneity-related dephasing is not refocused. A stimulated echo preparation [118, 119] stores the magnetisation in the M_z direction so that a long diffusion time between gradients can be set. However, half of the signal is spoiled and the scan times are generally long. Low-TR steady-state DW acquisitions are very efficient and have strong diffusion weighting but they suffer from complicated motion-induced phase artefacts because the diffusion weighting is mixed over multiple shots.

The most common readout modules in diffusion imaging acquire k-space in a single-shot, in which case the reconstructed images are largely insensitive to motion-induced phase artefacts. Acquisition schemes that acquire k-space in multiple shots either require phase navigation or some intrinsic insensitivity to motion artefacts, such as projection reconstruction (PR) [120] or line scanning [121] methods. However, artefacts arise in PR and line scan sequences when the motion-induced phase gradients are not parallel to the single line of k-space acquired [122]. These general classes of single- or multi-shot readouts are typically implemented in EPI- or FSE-based sequences, which have their respective advantages. EPI does not suffer from artefacts related to the phase stability of the FSE spin echo train but FSE sequences have low sensitivity to off-resonance artefacts. The choice of k-space trajectory (i.e., Cartesian, spiral, radial etc.) is important because it influences the image artefacts. For example, in radial or spiral sequences, off-resonance artefacts and residual motion-related errors manifest as streaking or blurring, which are more benign than the ghosting or geometric distortion artefacts encountered with Cartesian trajectories. In multi-shot schemes the trajectory also dictates the available motion correction strategies.

Spin-echo single-shot EPI (ss-EPI) [5] has been the dominant DW sequence for clinical imaging of human stroke [85] and DTI studies [3], due to its robustness to motion and high imaging speed (whole brain 2 mm isotropic data in ~10 s), which enables acquisition of large numbers of DW volumes. By accelerating the k-space acquisition with parallel imaging, distortions in 2 mm isotropic data have been reduced to acceptable levels at 3 T [101, 102] but a ceiling on resolution is imposed by the length of the readout. As discussed in Section 2.2.1, the length of the EPI readout makes it susceptible to blurring and off-resonance artefacts in the phase-encode direction, which increase when the readout duration is extended to encode higher resolution images. In practice the geometric distortions and blurring are often the most limiting artefacts, especially at higher field strengths where the field is less homogeneous and T_2^* is shorter. One approach that enables localised high-resolution DTI at 7 T [123] is to reduce the field of view so that a shorter readout with higher phase-encode bandwidth can be used.

In multi-shot sequences there are many ways to combine a navigator acquisition (which covers the centre of k-space at each shot) with a high-resolution acquisition (which covers the whole of k-space in multiple shots). The residual motion-related artefacts depend on the extent of the navigator and the readout details (and sometimes the diffusion preparation module being employed) but the correction methods have a common framework. The phase information in the navigator, which is effectively an estimate of the bulk motion, is used to undo the k-space phase offsets and shifts in the high-resolution data caused by the motion. A low-resolution 2D navigator image can be used to apply linear corrections for rigid-body motion [106, 107] or more complete non-linear corrections that account for non-rigid brain deformation [103, 108]. Trajectories where the high-resolution shots sample contiguous k-space points and are therefore fundamentally unaliased can be non-linearly corrected with a simple phase-subtraction

[103, 108] (without the need for an time-consuming iterative algorithm) if 2D navigator information is available. Self-navigated sequences [103, 124, 125] cover the centre of k-space in each high-resolution shot so they do not require a separate navigator and are therefore more efficient. The FSE PROPELLER sequence [103] acquires radial blades of k-space at each shot that are rotated about the centre in successive excitations to complete k-space sampling. Therefore, PROPELLER can combine self-navigation with a fast non-linear phase correction because each shot covers the centre of k-space and a low-resolution, full FOV (and therefore free from aliasing) image can be reconstructed from each blade. However, PROPELLER can be limited at high field strengths because the FSE echo-train requires good refocusing pulses to satisfy the CPMG condition and the inefficient radial coverage of k-space and high SAR (due to the large number of refocusing pulses) of the sequence can result in long scan times. 3D DW sequences have been a long standing target because they could allow high-resolution isotropic voxels without slice profile effects but it is difficult to acquire 3D navigator data in the short time available for image encoding. Recently small multiple slab acquisitions have been explored because it seems that phase errors can be controlled with 2D correction of each k_z encoded k-space plane [110, 112].

For this thesis, we focus on the multi-shot readout-segmented EPI (rs-EPI) method [6], which is attractive due to its close relationship to the widely used ss-EPI. As discussed in the next section, the rs-EPI sequence has been modified to include partial Fourier encoding, simultaneous multi-slice acceleration and 3D multi-slab DW imaging.

3.3 Readout-Segmented Echo-Planar Imaging

Readout-segmented EPI (rs-EPI) [6] is a multi-shot DW sequence that is capable of acquiring high-resolution images with reduced distortion and blurring compared to ss-EPI. Diffusion imaging with segmentation of k-space in the readout direction was first presented by Robson *et al.* without phase navigation [126] and this approach was modified by Porter *et al.* to include 2D non-linear phase correction [108, 127] and reacquisition of motion corrupted shots [128, 129]. The schematic in Fig. 3.3a shows how k-space coverage is completed in five shots, or readout segments, acquired in separate TR periods. Figure 3.4 gives examples of the k-space

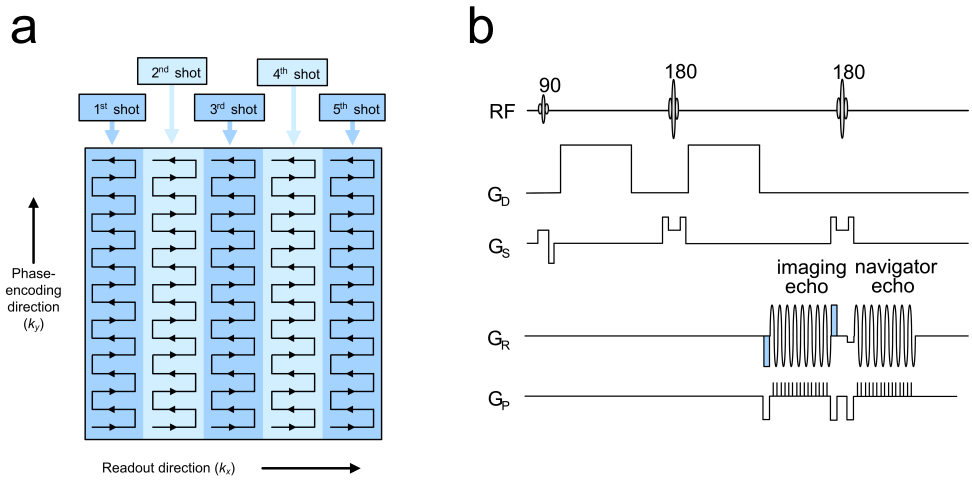

Figure 3.3: Illustration of (a) the k-space acquisition in rs-EPI and (b) the pulse sequence. The readout segments are acquired during the imaging echo and the navigator echo always acquires the central readout segment (the 3rd shot in this example). Pre-phasing gradients, shown in blue on the readout gradient (G_R) axis, are varied in different shots to move to the k_x location of each readout segment. Copyright 2009 Wiley. Used with permission from Porter and Heidemann [6].

and image-space representations of the readout segments. The smaller k_x width of a readout segment allows a shorter echo-spacing relative to ss-EPI where the full extent of k-space must be traversed. This short echo-spacing (or equivalently the high phase-encode bandwidth) reduces sensitivity to off-resonance artefacts and can be reduced further with parallel imaging. The length of the rs-EPI echo-

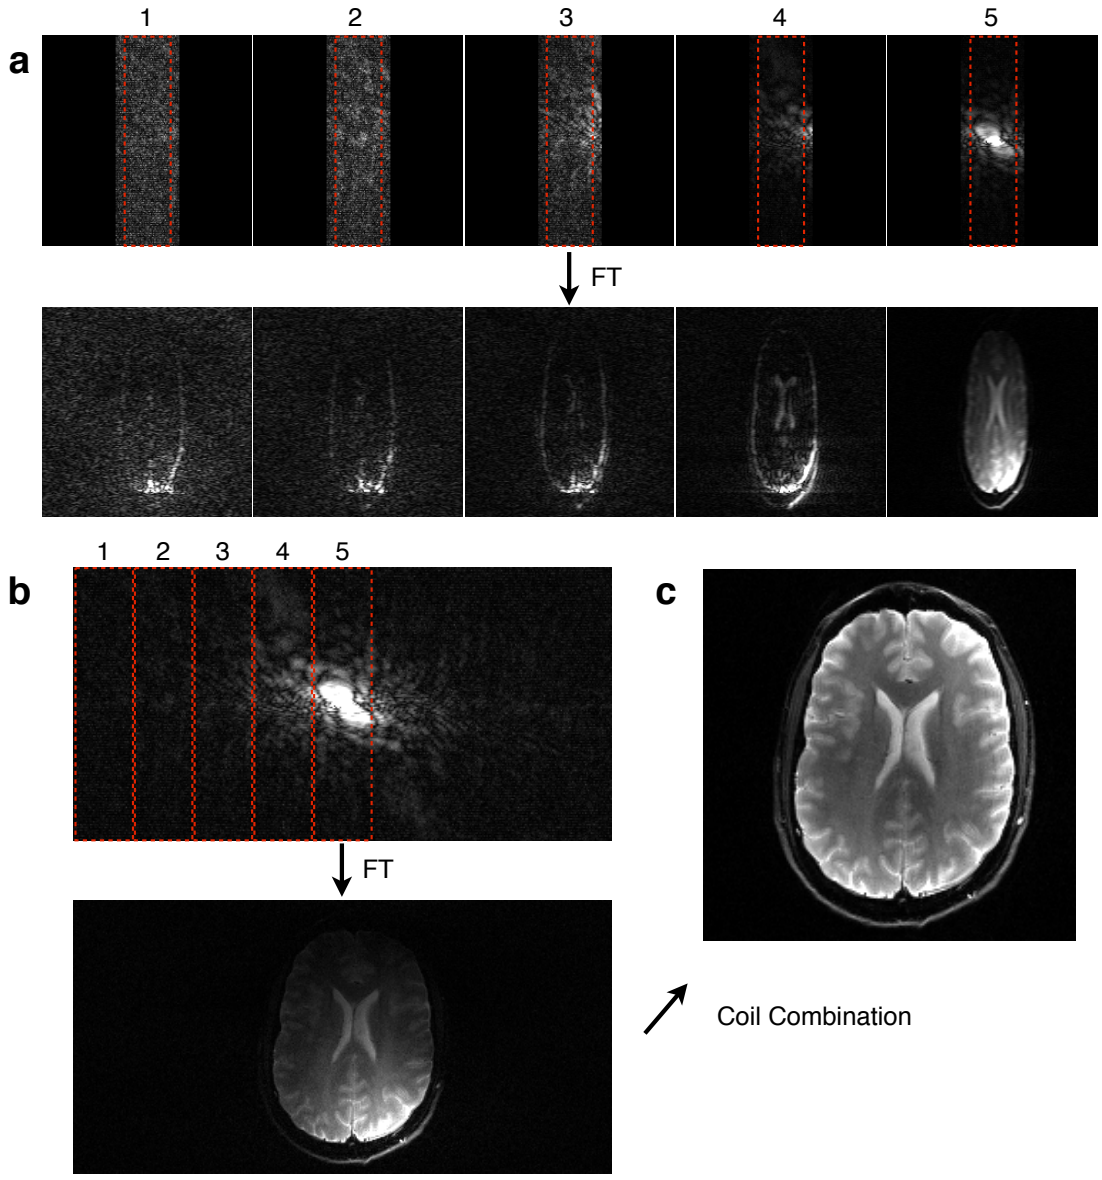

Figure 3.4: (a) Examples of k-space and image-space data from one channel in a 9 segment b = 0 rs-EPI acquisition. Note that the windowing of the data is different for each of the 5 segments. The k-space data outside the dotted red lines is discarded when the segments are spliced together to form the complete k-space. Typically 8 k_x columns are discarded on either side of each readout segment. (b) The complete k-space formed from all 9 readout segments and the resulting image from one channel. As can be seen from the wide readout FOV, the data are 2$\times$ k_x oversampled at this stage. (c) The final image after sum of squares combination of the data from all 32 channels.

train is generally shorter than in an equivalent resolution ss-EPI acquisition so there is less phase-encode blurring. This allows the acquisition of good quality high-resolution DW data, even at 7 T [130] where ss-EPI suffers from severe distortions and signal dropout. The improvements in distortions and blurring with rs-EPI are shown in the direct comparison with ss-EPI at 3 T in Figures 3.5 and 3.6.

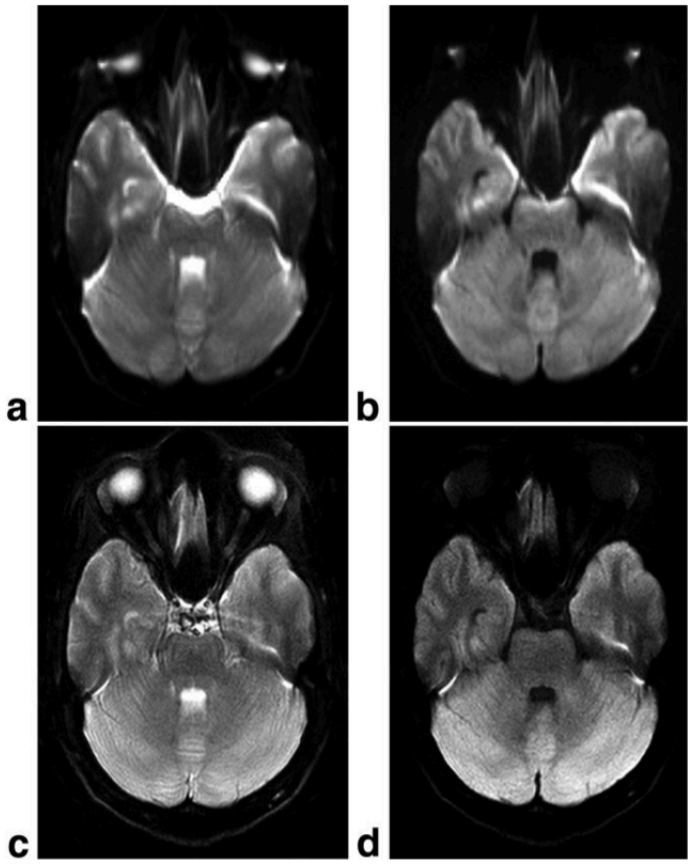

Figure 3.5: Comparison between single-shot EPI and readout segmented EPI at the base of the brain. Single-shot EPI acquisition with (a) b = 0 s/mm^2 and (b) corresponding trace-weighted image with b = 1000 s/mm^2. Readout-segmented EPI acquisition with (c) b = 0 s/mm^2 and (d) corresponding trace-weighted image with b = 1000 s/mm^2. Copyright 2009 Wiley. Used with permission from Porter and Heidemann [6].

The 2D navigator correction was implemented by acquiring the central readout segment (e.g., shot 3 in Fig. 3.3a) in a separate navigator echo after the imaging echo, as indicated in the pulse sequence diagram in Fig. 3.3b. The low k_x reso-

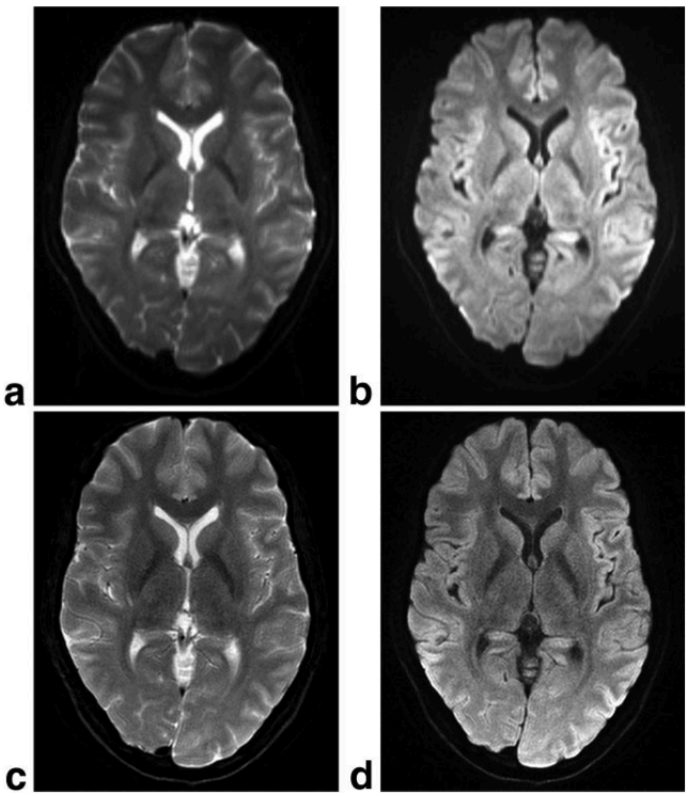

Figure 3.6: Comparison between single-shot EPI and readout segmented EPI at the level of the lateral ventricles. Single-shot EPI acquisition with (a) b = 0 s/mm^2 and (b) corresponding trace-weighted image with b = 1000 s/mm^2. Readout-segmented EPI acquisition with (c) b = 0 s/mm^2 and (d) corresponding trace-weighted image with b = 1000 s/mm^2. Copyright 2009 Wiley. Used with permission from Porter and Heidemann [6].

lution, full k_y resolution navigator phase estimate was used to apply a 2D non-linear image-space phase correction [108] to each readout-segment acquired with diffusion-weighting. Each readout segment samples contiguous k-space points and therefore image space representations of the segments are not aliased. Hence the correction can be applied as a simple complex multiplication that removes the phase of the navigator from the segment, which translates to a redistribution of signal in the k-space domain. As shown in Fig. 3.4, readout segments are overlapped so that there are more k_x columns than necessary (typically 8 on either side) and therefore signal which has been displaced by motion phase can be recovered into the true imaging window of the segment. These extra k_x columns also serve to improve regridding at the edges of readout segments, so that discontinuities in signal are avoided.

The navigator information is also used to guide a reacquisition of segments that had severe phase corruption caused by systolic pulsatile motion. If there is significant motion during the diffusion gradients then large parts of k-space can be shifted outside the segment window and the navigator correction is unable to recover the signal. These instances could be avoided with cardiac gating but as discussed in Section 3.1.1, this has associated problems. The solution in rs-EPI is to reacquire the shots with the worst motion corruption, a technique that was first demonstrated with 1D navigation in DW spin echo sequences [131, 128]. An example of how artefacts can be removed by reacquiring the corrupted shot (number 6) is shown in Fig. 3.7.

3.3.1 Thesis Motivation

This thesis is on the development of the rs-EPI pulse sequence, which has already been demonstrated at 3 and 7 T [6, 132, 130] and shown to be useful clinically [133, 134]. The high-resolution DW images are impressive however the segmentation

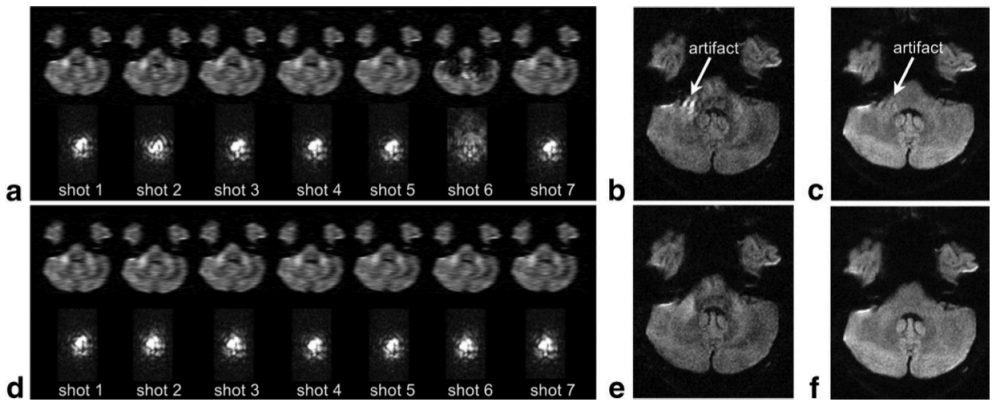

Figure 3.7: Example of artefact improvement with reacquisition. (a) The navigator images and k-space from 7 different shots. The navigator data is the central readout segment so the image data are low-resolution. The larger spread of k-space and the signal void artefact in shot 6 indicates that this shot has motion corruption. In the sequence this motion-corruption is detected by calculating the k-space distribution width, which quantifies the k-spacespread of the navigators. The resulting artefact can be seen in the reconstructed (b) diffusion- and (c) trace-weighted images. (d) The navigator data after reacquisition. Shot 6 and shot 2 have been reacquired and the artefact in the (e) diffusion- and (f) trace-weighted images has been removed. Copyright 2009 Wiley. Used with permission from Porter and Heidemann [6].

of k-space increases the scan time per volume. For equivalent matrix sizes, rs-EPI acquisition times are longer than ss-EPI roughly by the number of readout segments. Of course, the same data quality or resolution cannot be obtained with ss-EPI but the imaging time is an important consideration in clinical and research settings. Time may be limited when scanning acute patients who are in severe discomfort and also there are essential structural and diagnostic scans that require significant time in a scanning session. In DTI studies at least 30-60 DW volumes are generally required so the scan time per volume must be minimized. This motivates the work in Chapters 4 and 5, where strategies to reduce the scan time are explored.

In Chapter 4, readout segments on one side of k-space are omitted and the missing data are reconstructed with a partial Fourier algorithm [135]. The scan time is reduced by the number of segments omitted but the acceleration is lim-

ited to less than a factor of two. Two partial Fourier algorithms are compared to the original full k-space reconstruction in terms of image quality and estimates of diffusion parameters and it is demonstrated that a POCS reconstruction of partial Fourier encoded data gives reliable results. Chapter 5 presents a modification to the rs-EPI sequence where multiple slices are simultaneously excited and encoded [71]. A parallel imaging reconstruction is used to separate the multiplexed slice data into separate single-slice data. This translates into lower scan time because fewer excitations are required to achieve the same slice coverage within a TR period. TR and therefore scan times are reduced by the number of slices simultaneously imaged. Whole brain scans typically excite 60-70 slices during each TR resulting in TR values on the order of 10 s. In order to achieve the optimum SNR within a target scan time, the TR can be chosen to maximise SNR efficiency (signal/$\sqrt{\mathrm{TR}}$). Referring to the SNR efficiency plot in Fig. 3.8, the optimum TR is 1-2 s so reductions in TR will not only reduce scan times but also increase SNR obtained in a given scan time. The spin echo SNR efficiency plot in Fig. 3.8 used 3 T values for white matter tissue T_1= 1000 ms, T_2= 75 ms [136, 137] and echo time = 70 ms. Numerical simulation of the Bloch equations was used to calculate the spin echo signal. The compatibility of rs-EPI with slice-acceleration is demonstrated in trace-weighted and DTI acquisitions with reduced scan times [138]. An acceleration factor of 2 was used in the modified rs-EPI sequence but there is potential for faster scans because acceleration factors of 6-8 have been presented elsewhere.

Chapter 6 describes a 3D extension of the rs-EPI sequence, where multiple slabs were k_z phase-encoded in each TR. The k-space acquisition was also synchronised to the cardiac cycle to reduce motion phase artefacts, as has been shown previously by Miller and Pauly [108]. A 3D DW sequence would be very valuable because it may be possible to obtain high-resolution in the slice-direction,

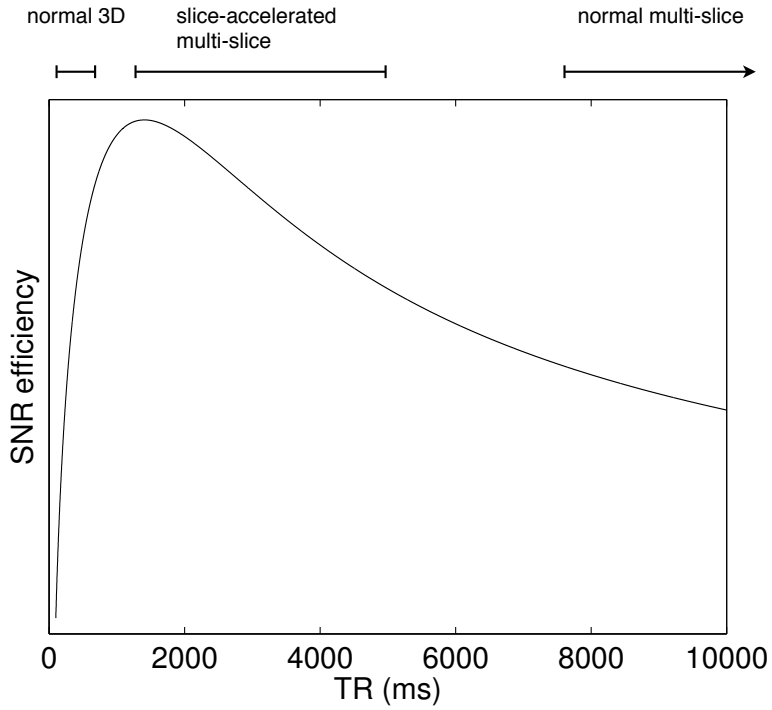

Figure 3.8: SNR efficiency (signal/$\sqrt{\mathrm{TR}}$) for a spin echo sequence is plotted as a function of TR. The spin-echo signal was calculated by numerical simulation of the Bloch equation with TE = 70 ms and white matter values of T_1= 1000 ms and T_2= 75 ms. In whole-brain multi-slice sequences, reductions in TR with slice acceleration to about 1-2 s will increase SNR efficiency. The optimal TR range is also achievable with multi-slab sequences, which only require about 10 slab excitations for whole-brain coverage.

which is limited in multi-slice sequences due to the difficulty in exciting thin, well-defined slices. A 3D multi-slab sequence would also allow imaging in the optimum TR range for SNR efficiency. Recent work in 3D diffusion imaging has used similar multi-slab approaches to the one presented here for phase-navigation and efficiency reasons.

Sequence modifications were programmed in the IDEA Siemens environment and images were reconstructed offline in Matlab. Future work would include translation of the Matlab reconstruction code into the scanner environment, which was beyond of the scope of this thesis. This and other remaining work is discussed in Chapter 7.

Chapter 4

Partial Fourier Readout-Segmented EPI

4.1 Introduction

The use of magnetic resonance imaging in assessing diffusion of water in tissue has found wide applicability in basic and clinical science. There is great interest in characterizing white matter integrity and pathways in brain imaging, with applications in basic neuroscience [4] and in assessment of tissue damage due to ischemia [139, 85, 140, 141, 142], but also in other applications such as in using diffusion characteristics to provide information on tumours [143]. However, the spatial resolution possible when using diffusion MRI has historically been limited by difficulties in correcting artefacts associated with bulk movement of the object during scan acquisition. To counter this, low-resolution single-shot echo planar imaging (ss-EPI) methods have been most commonly used as the readout module. However, sophisticated adaptations to the standard EPI acquisition, involving careful characterization of the confounding phase distortions using navigator echoes, have been used to enable multishot acquisition methods to be used,

with their consequent improvement in spatial resolution. One such extension is the use of readout-segmented EPI [126] (as opposed to the more conventional phase-encode segmented EPI). This sequence was adapted by Porter et al. [6] to include GRAPPA parallel imaging (to minimize B_0 inhomogeneity-related distortions), two-dimensional navigator correction (to account for motion-related phase corruptions), and navigator-based reacquisition (to allow for occasional uncorrectable k-space segments) [66, 108, 131]. That method uses a short-duration sinusoidal echo readout strategy to sample segments of k-space in the readout direction, thereby minimizing the echo spacing and hence image distortion. In contrast to the phase-encode segmented approach, this acquisition scheme samples contiguous k-space points at each shot, rather than interleaving them between shots. This approach results in Nyquist sampled data that are free from aliasing and convenient for applying a spatially variant phase correction. Other groups have demonstrated the effectiveness of readout-segmented EPI in preserving anatomical detail at 3 and 7 T [132, 130]. As noted by Heidemann *et al.*, the reduced echo-spacing allows a shorter echo time (TE) relative to ss-EPI, which has benefits in terms of high b-value diffusion encoding and finer spatial resolution when the shorter T_2 decay at higher field strength prohibits use of ss-EPI.

Readout-segmented EPI (rs-EPI) allows acquisition of high-resolution maps of apparent diffusion coefficient and other diffusion parameters, which promise application in evaluation of small infarcts in acute stroke. However, for rs-EPI to be used routinely in a clinical setting, and for application to multi-direction diffusion tractography studies, the imaging time must be reduced further to be more comparable with current ss-EPI protocols. This study demonstrates that the number of readout segments can be reduced from the number required for a full k-space acquisition, without significantly compromising image quality, by reconstructing missing k-space data with partial Fourier techniques. A prelim-

inary study investigated the combination of rs-EPI data with homodyne partial Fourier reconstruction in the readout direction [144] and a comparison of the homodyne and projection onto convex sets (POCS) algorithms for partial Fourier encoding in the phase-encode direction has been presented [145]. An evaluation of the effect of readout and phase-encode partial Fourier encoding on signal-to-noise ratio (SNR) efficiency in rs-EPI [145] concluded that phase-encode partial Fourier was advantageous in high-resolution (288×288) protocols when a twice-refocused diffusion preparation was used. However, the diffusion preparation used in our study was a modified monopolar scheme (see Methods section) so as to minimize the SNR penalty of the diffusion encoding duration. This chapter, which has been published elsewhere [135], compares homodyne [52] and POCS [146, 53] readout partial Fourier reconstructions to the full k-space acquisition. Comparison between these approaches is presented for the cases of estimates of the diffusion parameters: fractional anisotropy (FA), mean diffusivity (MD), and diffusion orientation (two-fibre model); T_2- and trace-weighted images; and SNR.

4.2 Theory

4.2.1 Partial Fourier Reconstruction

Partial Fourier reconstruction is possible because the spin density is a real function. The symmetry of k-space can be found by taking the complex conjugate of the signal equation:

$$\begin{aligned} S^*(k_x, k_y, k_z) &\propto \left(\int_x \int_y \int_z \rho(x,y,z) e^{-i2\pi(k_x x + k_y y + k_z z)} \,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z \right)^* \\ &\propto \int_x \int_y \int_z \rho^*(x,y,z) e^{i2\pi(k_x x + k_y y + k_z z)} \,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z \\ &\propto S(-k_x, -k_y, -k_z) \end{aligned} \tag{4.1}$$

This demonstrates that when the object is real, the k-space signal is Hermitian i.e., the real part is symmetric and the imaginary part is antisymmetric about the centre of k-space [17]. If this property holds then it is only necessary to acquire one half of 2D k-space and then the missing half can be reconstructed. In practice, the reconstructed image is complex due to phase caused by flow and motion or hardware imperfections such as offsets in resonance frequency, eddy currents and receive B_1 inhomogeneity. Therefore partial Fourier reconstructions attempt to make the image real by performing a phase correction step in image-space. Slightly more than half (e.g. 9/16 or 5/8) of k-space is acquired so that the symmetric region of low spatial frequencies (the central portion of k-space), where data have been acquired on both sides of k-space, provides a low resolution estimate of the image phase.

Homodyne Reconstruction

The homodyne reconstruction applies a filter to estimate the missing part of k-space. The origin of the filter is explained for the one-dimensional case, with reference to a derivation in Bernstein *et al.* [9]. The range of the k-space data is from $-k_{max}$ to k_{max} in a fully sampled acquisition and $-k_{ls}$ to k_{max} in a partial Fourier acquisition. The symmetric, low spatial frequency range used for phase correction is from $-k_{ls}$ to k_{ls}. The reconstructed real image is given by:

$$I(x) = \int_{-k_{max}}^{k_{max}} S(k) e^{-i2\pi kx} \, \mathrm{d}k \tag{4.2}$$

In our case with a real image, the data in the range ($-k_{max}$, k_{ls}) can be replaced by the complex conjugate of the data in the range (k_{ls}, k_{max}):

$$I(x) = \int_{-k_{max}}^{-k_{ls}} S^*(-k) e^{-i2\pi kx} \, \mathrm{d}k + \int_{-k_{ls}}^{k_{max}} S(k) e^{-i2\pi kx} \, \mathrm{d}k \tag{4.3}$$

In the first term in Eq. 4.3 can be written as a complex conjugate with a change of variable $k' = -k$:

$$\begin{aligned} I(x) &= \left[\int_{k_{ls}}^{k_{max}} S(k')e^{-i2\pi k'x}\,\mathrm{d}k'\right]^* + \int_{-k_{ls}}^{k_{max}} S(k)e^{-i2\pi kx}\,\mathrm{d}k \\ I(x) &= \left[\int_{k_{ls}}^{k_{max}} S(k')e^{-i2\pi k'x}\,\mathrm{d}k'\right]^* + \int_{-k_{ls}}^{k_{ls}} S(k)e^{-i2\pi kx}\,\mathrm{d}k + \int_{k_{ls}}^{k_{max}} S(k)e^{-i2\pi kx}\,\mathrm{d}k \end{aligned} \tag{4.4}$$

The identity for complex variables $z + z^* = 2\mathrm{Re}[z]$ is used to obtain:

$$I(x) = \int_{-k_{ls}}^{k_{ls}} S(k)e^{-i2\pi kx}\,\mathrm{d}k + 2\mathrm{Re}\left[\int_{k_{ls}}^{k_{max}} S(k)e^{-i2\pi kx}\,\mathrm{d}k\right] \tag{4.5}$$

The second term in Eq. 4.5 is real by definition and we assume that $I(x)$ is real. Therefore the first term must also be real and we can write:

$$I(x) = \mathrm{Re}\left[\int_{-k_{ls}}^{k_{ls}} S(k)e^{-i2\pi kx}\,\mathrm{d}k + 2\int_{k_{ls}}^{k_{max}} S(k)e^{-i2\pi kx}\,\mathrm{d}k\right] \tag{4.6}$$

This equation for the image is in terms of k-space data in the range ($-k_{ls}$, k_{max}) that is acquired in a partial Fourier acquisition. A filter $H(k)$ can be defined:

$$H(k) = \begin{cases} 0, & k < -k_{ls}, \\ 1, & -k_{ls} \leq k \leq k_{ls}, \\ 2, & k > k_{ls}. \end{cases} \tag{4.7}$$

The image can then be expressed as:

$$I(x) = \mathrm{Re}\left[I_H(x)\right] \tag{4.8}$$

$$I_H(x) = \mathrm{Re}\left[\int_{-k_{max}}^{k_{max}} H(k)S(k)e^{-i2\pi kx}\,\mathrm{d}k\right] \tag{4.9}$$

This shows that the homodyne high pass filter $H(k)$ can be used to reconstruct an image instead of calculating missing k-space data by conjugate symmetry. Smoother filters (see below) are used in practice to mitigate ringing caused by the sharp transitions.

As the image $I(x)$ is generally not real, taking the real part in Eq. 4.8 loses some signal. Before taking the real part, the image is phase corrected in an attempt to remove the phase. The phase correction uses an image calculated from the symmetric region of k-space in the range ($-k_{ls}$, k_{ls}):

$$I_L(x) = \int_{-k_{ls}}^{k_{ls}} S(k)e^{-i2\pi kx}\,\mathrm{d}k = \int_{-k_{max}}^{k_{max}} L(k)S(k)e^{-i2\pi kx}\,\mathrm{d}k \tag{4.10}$$

where

$$L(k) = \begin{cases} 0, & |k| \leq k_{ls}, \\ 1, & |k| > k_{ls}. \end{cases} \tag{4.11}$$

Again, a function with smoother transitions (e.g. Hanning filter) is used for the low pass filter $L(k)$. It is assumed that the low resolution $\phi_L(x)$, the phase of $I_L(x)$, is a good estimate of the true phase of $I(x)$. When the phase varies rapidly, e.g. due to sharp susceptibility changes, this assumption doesn't hold and the quality of the homodyne reconstruction is degraded. The full homodyne reconstruction can be expressed as:

$$I(x) \approx \mathrm{Re}\left[I_H(x)e^{-i\phi_L(x)}\right] \tag{4.12}$$

When implementing these phase calculations it is preferable to calculate the normalised complex image, shown in Eq. 4.13, to avoid problems with phase wrapping in the arctangent step.

$$e^{-i\phi_L(x)} = \frac{I_L^*(x)}{|I_L(x)|} \tag{4.13}$$

Iterative Reconstructions

Iterative approaches have two constraints: the image phase should be the low resolution estimate $\phi_L(x)$ and in k-space the estimated data should match the acquired data. This is achieved by iteratively applying phase correction and then transforming back to k-space. $S_j(k)$ is the complex k-space data at each step j and $I_j(x)$ is the image. The first image $I_0(x)$ is computed from the partial Fourier data, where the range of k-space values is ($-k_{ls}$, k_{max}). This image is then phase corrected and Fourier transformed back to k-space to generate $S_0(x)$, which contains an estimate of the missing part of k-space ($-k_{max}$, k_{ls}):

$$S_j(k) = \mathrm{FT}\left[I_j(x)e^{-i\phi_L(x)}\right] \tag{4.14}$$

The next iteration $S_{j+1}(x)$ is produced by replacing the values in the range ($-k_{ls}$, k_{max}) with the originally acquired data:

$$S_{j+1}(k) = W(k)S(k) + [1 - W(k)]\, S_j(k) \tag{4.15}$$

where $W(k)$ is a merging function, which can be smoothed if necessary:

$$W(k) = \begin{cases} 0, & k < -k_{ls}, \\ \\ 1, & k \geq= k_{ls}. \end{cases} \tag{4.16}$$

The next image $I_{j+1}(x)$ is:

$$I_{j+1}(x) = \mathrm{Re}\left[\mathrm{FT}^{-1}\left[S_{j+1}(k)\right] e^{-i\phi_L(x)}\right] \tag{4.17}$$

Images can be generated with the homodyne algorithm (iterative homodyne) or other approaches such as zero filling (POCS). The algorithms iterate until the

difference between successive images is small. Iterative approaches can perform better when there is rapid phase variation because they try to approximately match the k-space as well as constrain the phase.

4.3 Materials and Methods

Data were acquired with a diffusion-weighted readout-segmented EPI sequence from six healthy volunteers using a MAGNETOM Verio 3 T scanner (Siemens Healthcare, Erlangen, Germany) under an approved technical development ethics protocol. Additionally, to avoid confounding problems from subject motion, SNR was assessed for each acquisition/reconstruction combination using an agar gel phantom scanned 20 times during a repeated acquisition protocol. The diffusion preparation module was a modified version of the monopolar Stejskal-Tanner scheme [117]. The high phase-encode bandwidth resulting from the combination of the short rs-EPI echo-spacing with GRAPPA acceleration makes the sequence less sensitive to eddy current-induced phase changes (although we find some post-processing correction to be desirable), which allows monopolar schemes to be used rather than the eddy current compensated bipolar schemes. In the modified monopolar preparation, part of the refocusing gradient is moved prior to the 180° RF pulse (and negated) so that the full time between the 90° and 180° pulses is used for diffusion encoding, thereby allowing a shorter TE. Splitting the gradient also reduces eddy current effects relative to conventional monopolar encoding. A 32-channel head coil was used to acquire fully k-space sampled diffusion tensor images (DTI protocol) and diffusion-weighted images (DWI protocol), with the following parameters:

4.3.1 DTI Protocol

This protocol aims to acquire moderate-resolution data with low distortion for either tensor measurements or tractography [147, 88], for which we targeted a reasonably large number of diffusion directions (30 directions) and moderate scan time (17:19 min). The scan parameters were: FOV 220 mm, matrix 110×110, resolution 2.0×2.0×2.0 mm^3, 36 slices, 5 readout segments for full k-space reconstruction, GRAPPA parallel imaging with an acceleration factor of 2, repetition time (TR) = 5200 ms, TE = 75 ms, echo-spacing = 320 μs, five volumes at b = 0 s/mm^2 and 30 volumes of different directions at b = 1000 s/mm^2, imaging time for full k-space data 17:19 min (including magnetization preparation dummy scans, Nyquist ghost phase correction and GRAPPA reference scans, and time for the reacquisition of up to 20% of scans due to severe motion-induced phase errors that could not be adequately corrected using navigator data). The number of slices was restricted by a limit on the size of the raw data file imposed by the current implementation of the navigator-based reacquisition, but coverage of the major white matter tracts was possible with 36 slices. The set of 30 gradient vector directions was calculated using an algorithm that uses an analogy of electrostatic repulsion between vectors to give the optimal isotropic distribution [148].

4.3.2 DWI Protocol

This protocol aims to acquire high-resolution data for assessing focal ischemia in patients, dictating high-resolution trace-weighted measurements in a relatively short scan time. The scan parameters were: FOV 220 mm, matrix 256×256, resolution 0.9×0.9×4.0 mm, 23 slices, 13 readout segments for full k-space reconstruction, GRAPPA parallel imaging with an acceleration factor of 2, TR = 5400 ms, TE = 65 ms, echo-spacing = 320 μs, one volume at b = 0 s/mm^2 and three vol-

umes in orthogonal directions at b = 1000 s/mm^2, imaging time for full k-space data 4:48 min (including magnetization preparation dummy scans, Nyquist ghost phase correction and GRAPPA reference scans, and time for the reacquisition of up to 20% of scans due to severe motion-induced phase errors). For both DTI and DWI protocols, the acquisition sequence included a second echo that acted as a navigator echo for the imaging data.

4.3.3 Image Reconstruction

Images were reconstructed offline in Matlab (Mathworks, Natick, MA, USA) on an Intel Xeon E5640 processor (2.67 GHz) with 64-GB RAM via a computer cluster. The reconstruction was similar to the approach described by Porter *et al.* [6]. All readout segments were regridded to fit the sinusoidal readout gradient waveform, with each echo consisting of either 120 or 108 4$\times$ oversampled readout points (DWI and DTI protocols, respectively). The combination of the uniform time signal sampling and the sinusoidal readout gradient results in non-uniform k-space sampling i.e. slower k-space sampling at the edges of each line than at the centre. Therefore an oversampling factor of 4 was used to ensure that the Nyquist criterion is satisfied by the digital filter bandwidth, which is set on the assumption that the k-space is sampled uniformly. The data were then Nyquist ghost phase corrected: for all echoes and segments at each slice, lines collected under negative readout gradient were corrected in image-space for constant and linear phase terms using a double-echo reference scan for each slice, which was collected at the start of the acquisition. The same correction was applied to the GRAPPA reference scans for each slice before missing k-space lines were reconstructed with a GRAPPA algorithm. For DWI, a two-dimensional navigator phase correction was applied to all readout segments by performing an image-space phase correction on each imaging segment with its corresponding navigator seg-

ment. Navigator phase correction was not applied to b = 0 s/mm^2 images. The readout segments were then concatenated in the k_x direction after discarding the eight overlapping columns at the interfaces between segments (at this stage the data were 2$\times$ readout oversampled and the number of k_x points per segment, before additional columns were discarded, was 60 for the 110$\times$110 DTI protocol and 54 for the 256$\times$256 DWI protocol). These extra columns were acquired to aid the navigator phase correction by allowing recovery of k-space signal shifted outside the true k_x range of the readout segments by motion corruption and also to improve the regridding at the edges of readout segments. The k-space matrix of concatenated readout segments was subsequently Fourier transformed and individual coil images were combined using a sum-of-squares reconstruction to form an image, which was finally cropped for the remaining 2$\times$ readout oversampling to match the prescribed imaging matrix size.

For partial Fourier simulations, a varying number of readout segments on one side of the concatenated k-space matrix were zeroed. This method was chosen to simulate the achievable reduction in the total imaging time when using partial Fourier acquisition in the readout direction, because each unacquired readout segment reduces the acquisition time by one TR per slice. Hence, the imaging time (disregarding any dummy scans and reference scans) using a partial Fourier encoding approach would be reduced by the partial Fourier factor. Homodyne and POCS partial Fourier reconstructions using 3-4 of the total 5 readout segments in the 110$\times$110 DTI protocol and 7-12 of the total 13 readout segments in the 256$\times$256 DWI protocol were compared with their respective full k-space reconstructions. Further, the homodyne and POCS images were compared with the common strategy of simply zero-filling k-space to the required imaging matrix size.

4.3.4 Readout Partial Fourier Reconstruction

k-Space data acquired with readout (k_x) partial Fourier encoding consist of a symmetric region of low k_x spatial frequencies and a region of high spatial frequencies, which has only been acquired on one side of k-space in the k_x direction. In both regions, the full range of k-space is sampled in the phase-encode (k_y) direction. For the partial Fourier reconstruction approaches used in this study the extent of the symmetric low spatial frequency region was taken to be the maximum possible, given the partial Fourier factor used. For example, in the case of partial Fourier reconstruction with 12/13 segments, where the k-space data from segment 13 was omitted, the symmetric low spatial frequency region was segments 2-12. Before the phase correction, the low spatial frequency region was Hanning-filtered in k-space to reduce ringing in the final images. It was found that the effect of using different filters (triangular, Gaussian, Chebyshev, Blackman, Bohman) was very similar; only subtle changes in the location of the residual ringing were noticeable. The simplest reconstruction is to zero-fill the un-acquired portion of k-space to the required imaging matrix size, Fourier transform each channel image, and combine their sum-of-squares to form the final image. Zero-filled reconstruction results are shown in the comparison of images and absolute difference maps (Fig. 4.3) to demonstrate the induced blurring but they are excluded from the remainder of the analysis, because the images are of a lower resolution and therefore do not provide a useful comparison.

Homodyne Algorithm

The homodyne algorithm applies a weighting function, $H(k_x)$, in k-space which, in its basic form, doubles the acquired high spatial frequencies relative to the symmetric region of low spatial frequencies [52]. Two common homodyne weighting functions were compared: a linear ramp filter, $H_R(k_x)$, and an apodized step filter,

$H_{AS}(k_x)$, which are described in Eqs. 4.18 and 4.19, respectively.

$$H_R(k_x) = \begin{cases} 0, & k_x < -k_{ls}, \\ 1 + k_x/k_{ls}, & -k_{ls} \leq k_x \leq k_{ls}, \\ 2, & k_x > k_{ls}. \end{cases} \tag{4.18}$$

where $-k_{ls}$ to k_{ls} is the range of the low spatial frequency region.

$$H_{AS}(k_x) = \begin{cases} 0, & k_x \leq -k_{ls} - w/2, \\ \cos^2\left(\frac{\pi(|k_x| - (k_{ls} - w/2))}{2w}\right), & -k_{ls} - w/2 < k_x < -k_{ls} + w/2, \\ 1, & -k_{ls} + w/2 \leq k_x \leq k_{ls} - w/2, \\ 1 + \cos^2\left(\frac{\pi(|k_x| - (k_{ls} - w/2))}{2w}\right), & k_{ls} - w/2 < k_x < k_{ls} + w/2, \\ 2, & k_x \geq k_{ls} + w/2. \end{cases} \tag{4.19}$$

where w is equal to 0.9 times the width of one readout segment (after additional columns have been discarded). These weighting functions are preferred to the basic step weighting function because they reduce ringing in images; however, the smoothed filters lead to a further reduction in SNR [52]. The filters have different benefits in terms of image artefacts [52, 149], but it was found that the use of an apodized step filter was more accurate in all of our analyses, and therefore, the homodyne results presented were generated with this filter. The reduced k-space matrix for each channel was multiplied by $H_{AS}(k_x)$, and the result was Fourier transformed to the image domain. The channel images were phase corrected using a low-resolution phase image (calculated from the Hanning-filtered low spatial frequency data) and the real parts combined with sum-of-squares to form the final image. The symmetric component of $H_{AS}(k_x)$ is uniform across k-space and this fact is exploited in homodyne reconstruction by taking the real

component (which is symmetric) of the complex image at the final step (after multiplication by $H_{AS}(k_x)$ in k-space and phase correction in image-space) to estimate the unacquired k-space data.

POCS Algorithm

The POCS reconstruction method is an iterated reconstruction approach in which the full k-space image is estimated during each iteration but with the requirement that any lines for which true data are available match those of the estimated image [53]. For the implementation used here, the reduced k-space matrix for each channel was zero-filled and Fourier transformed to image space, then phase corrected using a low-resolution phase image (again, calculated from the Hanning-filtered

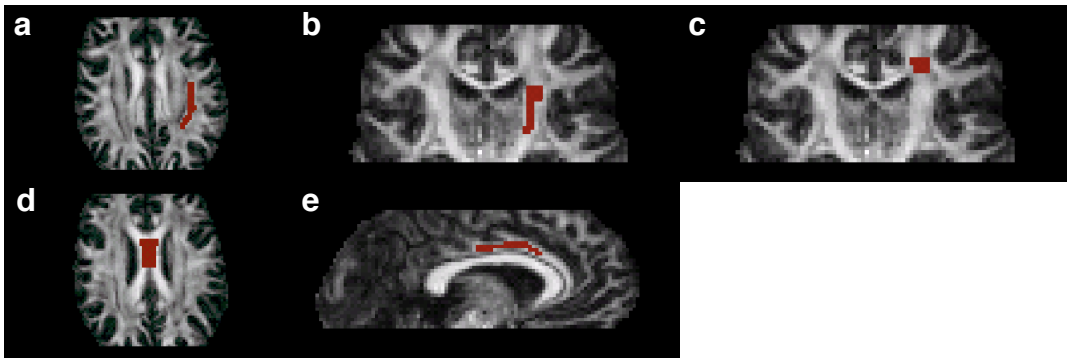

Figure 4.1: Locations of hand-drawn fibre tract ROIs used for analysis of the diffusion parameters (displayed on the full k-space FA map from a representative subject). These consisted of ROIs in (a) the superior longitudinal fasciculus, (b) the corticospinal tract, (c) the portion of the centrum semiovale where the corticospinal tract, the superior longitudinal fasciculus and a callosal radiation cross (crossing fibre region), (d) the corpus callosum, and (e) the cingulum bundle. Note that bilaterally placed ROIs are shown only on the right-hand side for simplicity.

low spatial frequency data described above). This estimated image was Fourier transformed back to k-space and the zero-filled portion of k-space in the original matrix was replaced with the corresponding part of the estimated k-space (no merging function between measured and estimated k-space was required). The voxel-by-voxel difference of POCS reconstructions from full k-space was found

empirically to asymptote after four to five iterations. In addition, the voxel-by-voxel difference from the image from the previous POCS iteration fell below the standard deviation of the signal in a background region after four to five iterations. A consistent increase in stability was not seen for more than five iterations. Hence, the algorithm was iterated five times and the resulting images for each channel were combined with sum-of-squares.

The times to reconstruct one brain volume of 110×110×36×32 (36 slices, 32 channels) data using the computing equipment described above were ~2.5 s for the homodyne algorithm and ~7.5 s for the POCS algorithm. These computation durations could be reduced by parallelization over slices and channels.

4.3.5 Image Processing and Analysis

DTI Protocol

Raw images were calculated using full k-space reconstruction and homodyne, POCS, and zero-filled partial Fourier reconstruction using 3/5 and 4/5 segments. Eddy current correction was applied to the raw data using FMRIB's Diffusion Toolbox [150], which consists of an affine registration (12 degrees of freedom) of all volumes to the initial b = 0 s/mm^2 volume. Non-brain tissue was removed from the raw images using Brain Extraction Tool [151], and the diffusion tensor was fitted to the reconstructed data using FMRIB's Diffusion Toolbox. FA and MD maps were calculated from the fitted tensors for each of the various reconstruction options. Bayesian estimation of diffusion parameters obtained using sampling techniques [152] was used to generate probability density distributions on two crossing fibre orientations at each voxel and from these the 95% uncertainty angle for the principal fibre orientation was estimated. Regions of interest (ROIs) were hand-drawn in the mid-sagittal plane for the corpus callosum, and

bilaterally for the cingulum bundle, the corticospinal tract, the superior longitudinal fasciculus, and the portion of the centrum semiovale where the corticospinal tract, the superior longitudinal fasciculus and the callosal radiations cross (crossing fibre region). Figure 4.1 shows ROIs drawn on a representative FA map from Subject 1. The mean number of voxels in the fibre tract ROIs were: 126±12 in the corpus callosum, 50±12 and 59±12 in the left and right cingulum, 66±14 and 67±12 in the left and right corticospinal tract, 92±14 and 91±8 in the left and right superior longitudinal fasciculus, 36±11 and 34±9 in the left and right crossing fibre regions. For each partial Fourier reconstruction, maps showing voxel-by-voxel difference and absolute difference in FA and MD were calculated relative to the full k-space reconstruction, from which ROI analyses could be performed.

DWI Protocol

Raw images were calculated using full k-space reconstruction and homodyne, POCS, and zero-filled partial Fourier reconstruction using 7/13, 8/13, 9/13, 10/13 and 12/13 segments. Non-brain tissue was removed from the raw images using Brain Extraction Tool and segmented white matter masks were calculated using FMRIB's Automated Segmentation Tool [153]. Trace-weighted images were calculated by taking the geometric mean of the b = 1000 s/mm^2 images acquired with three orthogonal diffusion weightings. The voxel-by-voxel difference and absolute difference in T_2- and trace-weighted signal intensity for voxels in segmented white matter in representative central slices of these brain-extracted images were calculated for each partial Fourier reconstruction relative to the full k-space reconstruction. These differences were normalised by dividing by the full k-space signal at each voxel.

SNR was estimated in the gel phantom for b = 0 s/mm^2 and b = 1000 s/mm^2 images for the 110×110 and 256×256 imaging matrices. The use of multichannel

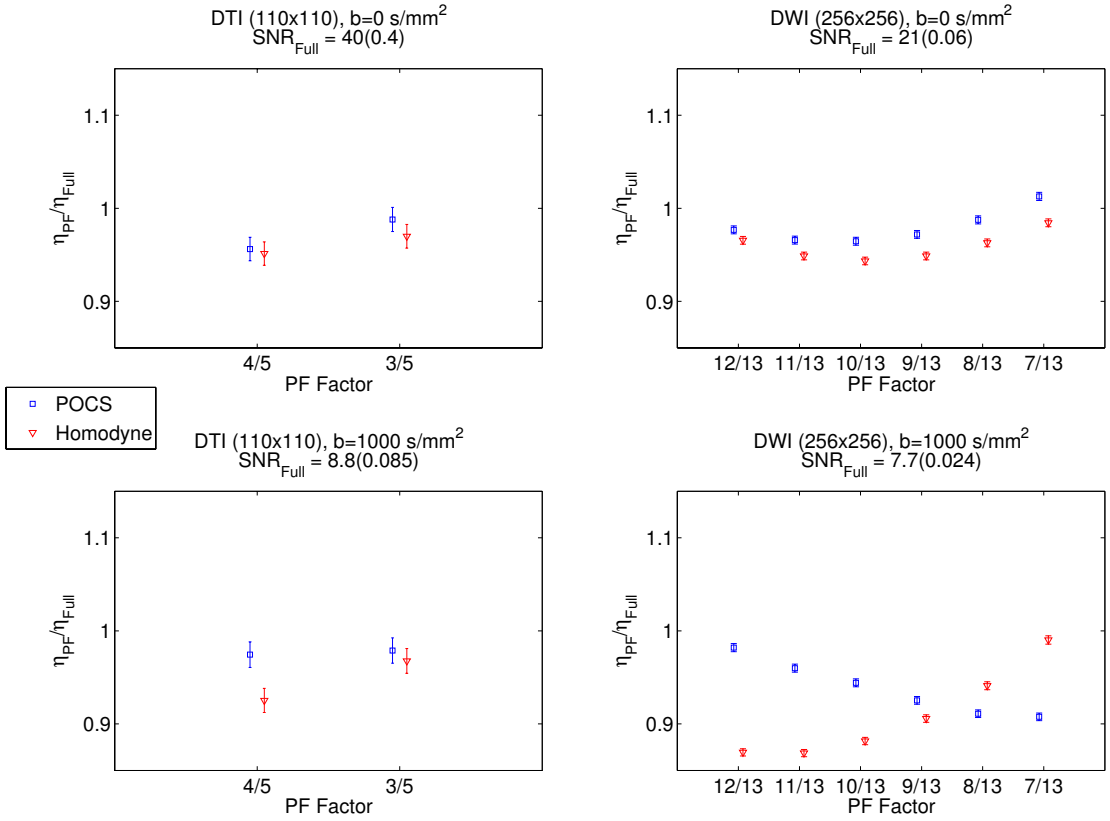

Figure 4.2: Ratios of partial Fourier SNR efficiency to full k-space SNR efficiency (η_{PF}/η_{full}) for the various reconstruction approaches are shown using SNR values calculated in a gel phantom. Temporal noise was estimated by acquiring 20 volumes and computing the square root of the mean within the ROI of the temporal variance in signal at each voxel. The mean of the signal within the ROI and over time was then divided by the temporal noise to give an estimate of SNR. The error bars show the standard error of the SNR efficiency ratio, which were calculated from the standard deviation of the signal within the ROI and over time divided by the square root of the number of voxels (standard error on the signal), and the square root of the standard deviation within the ROI of the mean temporal variance divided by the square root of the number of voxels (error on the noise). The standard error on the SNR was then calculated by quadrature error propagation.

coils and partial Fourier and parallel imaging GRAPPA reconstructions can cause the noise in the image to vary spatially so for this reason a measure of temporal noise was taken by acquiring 20 volumes and computing the square root of the mean within the ROI of the temporal variance in signal at each voxel. The mean of the signal within the ROI and over time was then divided by the temporal noise to give an estimate of SNR. For the 110×110 data, this was carried out for full k-space reconstruction and homodyne and POCS partial Fourier reconstruction using 3-4 of the total 5 segments. For the 256×256 data, this was carried out for full k-space reconstruction and homodyne and POCS partial Fourier reconstruction using 7-12 of the total 13 segments. Using these SNR values, the SNR efficiency, η, was calculated:

$$\eta = \frac{\mathrm{SNR}}{\sqrt{T}/N_{sl}} \tag{4.20}$$

where T is the acquisition time and N_{sl} is the maximum number of slices per TR. Omitting readout segments reduces the acquisition time by the partial Fourier factor, relative to the full k-space acquisition, and the maximum number of slices per TR is the same for partial Fourier and full k-space acquisitions.

4.4 Results and Discussion

Ratios of partial Fourier SNR efficiency to full k-space SNR efficiency (η_{PF}/η_{full}) for the various reconstruction approaches are shown in Fig. 4.2. The error bars in Fig. 4.2 show the standard error in the SNR efficiency ratio, which were calculated from the standard deviation of the signal within the ROI and over time divided by the square root of the number of voxels (standard error on the signal), and the square root of the standard deviation within the ROI of the mean temporal variance divided by the square root of the number of voxels (standard error on the noise). The standard error in SNR was then calculated by quadrature error

propagation. Taking the ratio of partial Fourier reconstruction to full k-space SNR efficiency using Eq. 4.20 and rearranging gives:

$$\mathrm{SNR}_{PF} = \frac{\eta_{PF}}{\eta_{full}}\sqrt{\frac{T_{PF}}{T_{full}}}\mathrm{SNR}_{full} \tag{4.21}$$

Overall the SNR follows the expected trends with η_{PF}/η_{full} being close to one and therefore demonstrating that the SNR approximately scales with the square root of the partial Fourier factor. In b = 1000 s/mm^2 results, where η_{PF}/η_{full} is less than one, a combination of slightly reduced signal values and increased noise values (not shown), relative to full k-space cause reduced SNR in partial Fourier reconstruction. In homodyne reconstructions (using apodized step and ramp filters, only the SNR efficiency data for the apodized step filter reconstruction is shown), there is an increase in noise with 12/13 segments, relative to POCS. The noise in homodyne and POCS reconstructions then increases approximately linearly with partial Fourier factor to reach similar levels at 7/13 segments. The signal in POCS reconstruction decreases whereas the signal in homodyne reconstructions increases marginally. The increase in signal in homodyne reconstruction arises from a higher level of artefactual background intensity, which is noticeable in the reconstructed images.

Examples of images from representative subjects are shown in Fig. 4.3, together with absolute difference maps between partial Fourier and full k-space reconstructions. Figures 4.3a and b compare FA and MD DTI maps reconstructed using full k-space with those collected using 3/5 and 4/5 partial Fourier reconstructions. Figures 4.3c and d compare T_2- and trace-weighted DWI images reconstructed using full k-space with those collected using 7/13 and 9/13 partial Fourier reconstructions. The images in Fig. 4.3 show qualitatively that the homodyne and POCS reconstructions are well matched to the full k-space reconstruc-

tion and are an improvement on zero-filling the k-space matrix. Figure 4.3 further shows that POCS produces the lowest level of artefact as judged by comparing the absolute difference images.

Mean FA and MD values in the listed fibre tract ROIs for the 4/5 and 3/5 POCS and homodyne partial Fourier reconstructions are compared with those for the full k-space acquisition in Fig. 4.4. Figure 4.4 also compares the mean values in each ROI of the 95% uncertainty angle on the estimated fibre orientation and the residual error on the tensor fit. The error bars in Fig. 4.4 are all 95% confidence intervals on the mean value in the ROI. Figure 4.5 shows box plots indicating the difference values in FA and MD within the different ROIs calculated from 3/5 partial Fourier reconstructions versus full k-space reconstruction of brain data acquired with the DTI protocol. The notches on the blue box show the 95% confidence intervals on the median difference values in each ROI. This reveals a good reproduction of the FA and MD values when using the POCS method (low median difference and low interquartile range). For the homodyne method, there is evidence of variable systematic bias across ROIs and higher variance in the FA and MD values relative to the POCS case. Similar results to these were observed for the case of 4/5 partial Fourier (data not shown), with the homodyne reconstruction showing a greater bias than the POCS method. The extreme difference values and outliers shown in the boxplots are caused by small shifts of voxels between adjacent regions of higher and lower FA/MD, caused by subtle differences between the partial Fourier and full k-space reconstructions. Plots showing the relationship between full k-space and partial Fourier reconstructions for T_2- and trace-weighted signal intensity (from brain data acquired with the DWI protocol) are shown in Fig. 4.6 for 7/13, 8/13, 9/13, 10/13, 11/13 and 12/13 partial Fourier factors. The Figure shows the mean of the absolute voxel-by-voxel differences in segmented white matter across subjects, with 95% confidence intervals, and

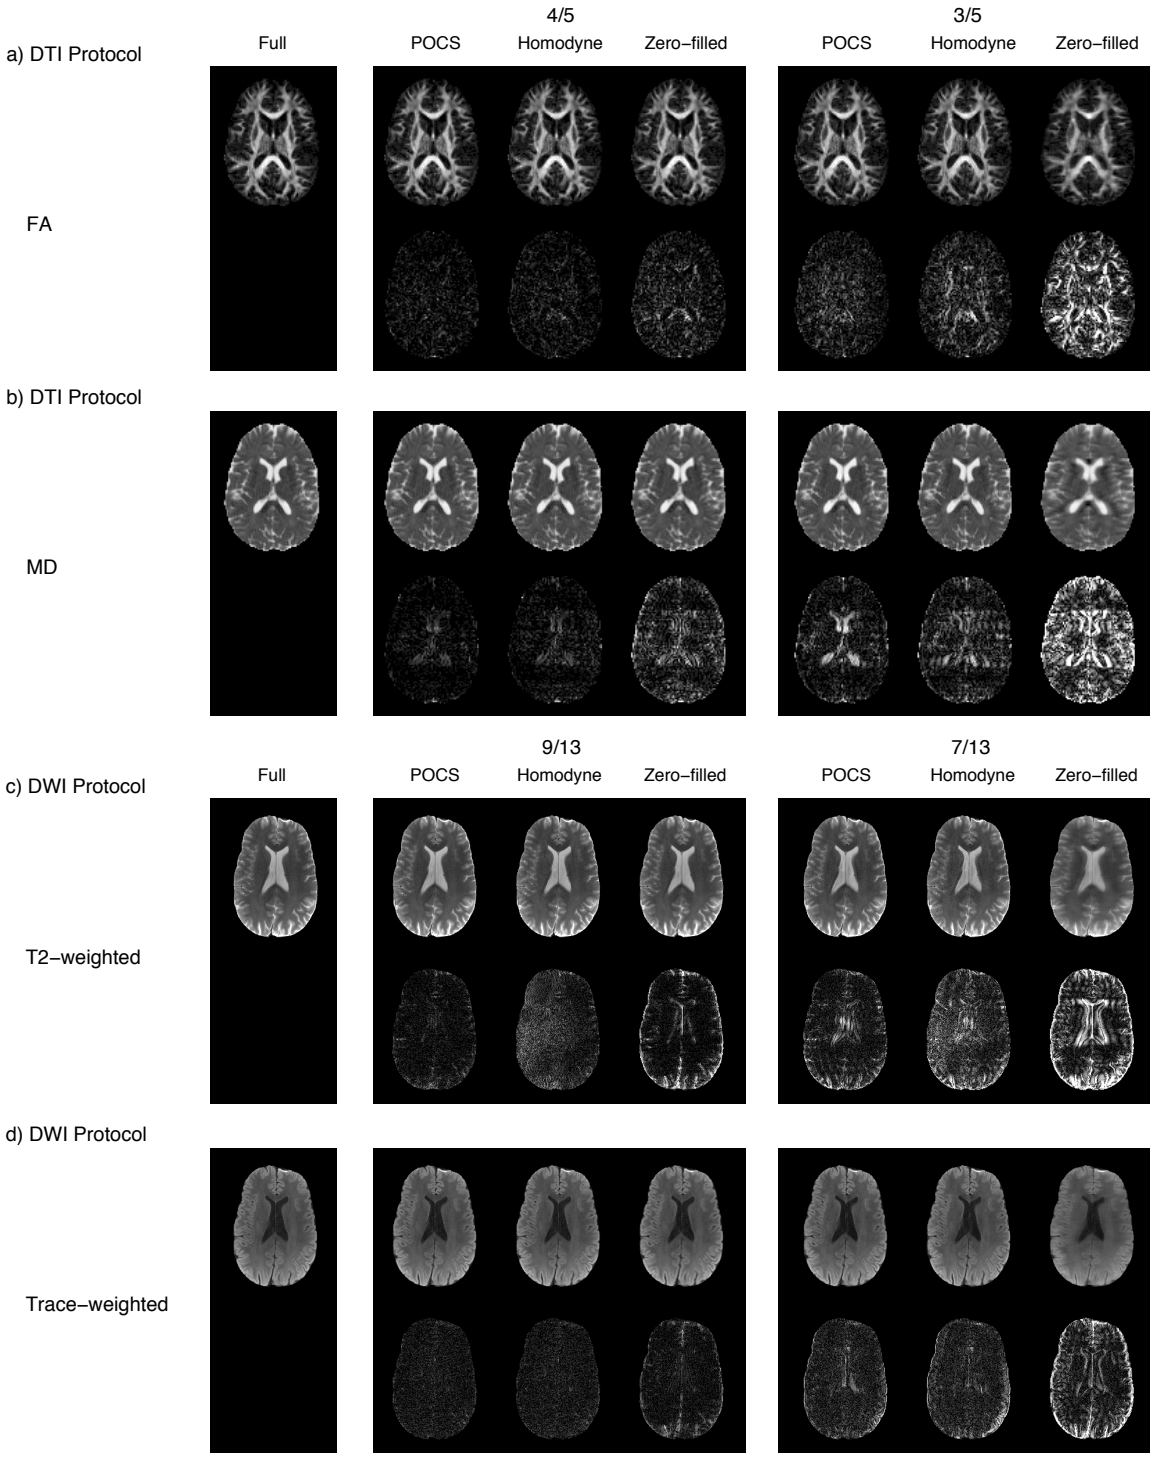

Figure 4.3: Upper rows: comparison with full k-space reconstruction of (a) FA and (b) MD maps reconstructed using 3-4 of the total 5 readout segments and (c) T2- and (d) trace-weighted images reconstructed using and 7 and 9 of the total 13 readout segments with POCS, homodyne and zero-filled reconstructions. Lower rows: maps of the absolute difference from full k-space for each of the partial Fourier reconstructions. The FA absolute difference maps are scaled by a factor of 5 and the MD, T2- and trace-weighted absolute difference maps are scaled by factor 10.

the standard deviation of the absolute difference (data are normalised to the full k-space signal intensity). Segmented white matter ROIs were used to detect differences in the reconstructions without influence from reconstruction artefacts at the edge of and outside the brain. This reduced the range of the differences and the number of voxels, which made the comparisons of the mean of the differences more robust. The mean number of voxels across subjects in the segmented white matter mask was 13,673$\pm$933 and the mean intensities in the segmented white matter of the full k-space reconstructions of T_2- and trace- weighted images were 515$\pm$139 and 251$\pm$68, respectively. Again, the POCS method shows the lowest systematic bias and variance, indicating that the POCS method best replicates the full k-space values. Finally, a comparison of maps of the principal diffusion direction for full k-space reconstruction and 3/5 POCS partial Fourier reconstruction is shown in Fig. 4.7.

In order to evaluate whether the errors imposed by a given technique are likely to adversely affect studies, it is useful to compare our errors to the variations expected, e.g., due to pathology. Differences in FA of 0.04–0.07 (mean standard deviation 0.06) in normal appearing white matter between controls (18$\pm$1) and patients (28$\pm$4) with multiple sclerosis have previously been reported [154]. In another study of 39 patients and 21 controls, patients with multiple sclerosis showed an FA in normal appearing white matter of 0.62 versus a significantly higher FA in volunteers of 0.65 (difference 0.03, with the mean range size [max - min] being 0.09) and a trend of increased diffusivity of 0.84×10^{-3} mm^2/s in controls versus 0.82×10^{-3} mm^2/s in patients (difference 0.02×10^{-3} mm^2/s, with the mean range size [max – min] being 0.17×10^{-3} mm^2/s) [155]. Looking at Fig. 4.5, the largest increase in median FA between 3/5 POCS and full k-space reconstruction across all subjects in a single ROI was 0.0094 in the corpus callosum and FA increased in all ROIs. This shows a consistent bias on the median FA, which is below the level of

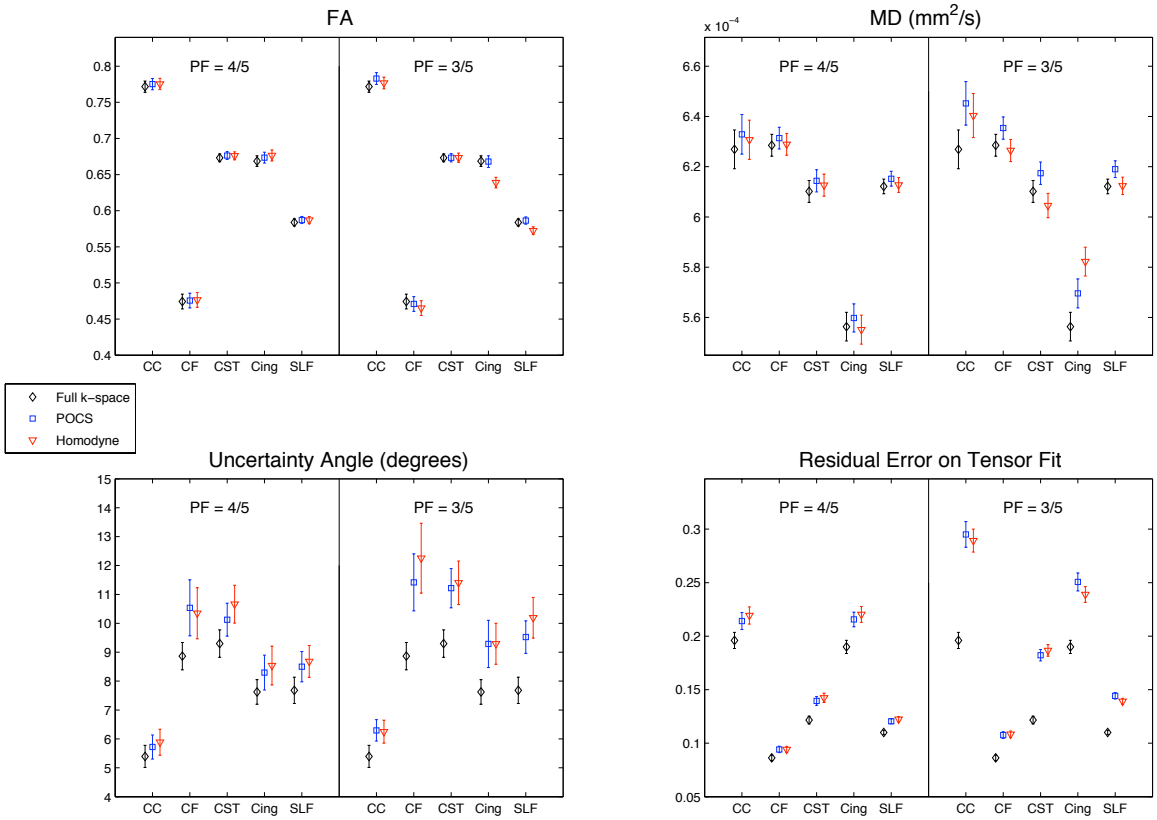

Figure 4.4: Plots of the mean values (with 95% confidence intervals) in the listed fibre tract ROIs of FA, MD, the angular uncertainty on the principal diffusion direction, and the residual error on the tensor fit. Results are shown for full k-space and 3/5 and 4/5 POCS and homodyne partial Fourier reconstructions. The results in left and right regions of bilateral ROIs (CF, CST, Cing, SLF) were not significantly different ($p < 0.05$) so the mean values reported are across both left and right regions. Key: CC = corpus callosum; CF = crossing fibre (the region in the centrum semiovale where the corticospinal tract, the superior longitudinal fasciculus and a callosal radiation meet); CST = corticospinal tract; Cing = cingulum; SLF = superior longitudinal fasciculus.

FA changes detected in clinical studies and previously reported intersubject and intersession variability [156, 157]. For MD, the bias is consistent and lower in 3/5 POCS reconstruction than in both homodyne and zero-filled reconstructions. The largest increase in 3/5 POCS in the median MD over all subjects in a single ROI was 0.0199×10^{-3} mm^2/s in the corpus callosum and MD increased in all ROIs. This means that using 3/5 partial Fourier-encoded rs-EPI with POCS reconstruction would allow reliable detection of FA and MD changes within subjects and across fibre tract ROIs and the ability to report accurate absolute FA and MD values relative to the reference full k-space values. However, it should be noted that the small bias found in this study could cause false positives (results that wrongly indicate pathology), which would be exacerbated by large sample sizes. The comparison of the maps of principal diffusion direction in Fig. 4.7 shows that there is little visible difference between the full k-space reconstruction and the 3/5 POCS reconstruction.

To evaluate if the reconstructions bias the mean angle, the angular differences between the partial Fourier reconstructions versus full k-space were calculated (in the fibre tract ROIs) for rotations about the L-R, A-P, and I-S axes. The angular differences from full k-space were found to be normally distributed about zero in all three planes with small mean differences and large standard deviations (mean differences of -0.2$\pm$8°, 0.4$\pm$20°, and -0.3$\pm$17° in the sagittal, coronal, and axial planes, respectively). This demonstrates that any bias will be small, i.e., it is much smaller than other sources of error in the calculation of the mean angle. The maps of 95% uncertainty angle in Fig. 4.7 and the mean uncertainty values in the fibre tract ROIs in Fig. 4.4 show a small increase ($< 3°$) in uncertainty on the estimate of the principle fibre orientation in $3/5$ POCS relative to full k-space which reflects the reduction in SNR of the raw images. We have compared the angular uncertainty we report in Fig. 4.4, to the uncertainty using our standard ss-

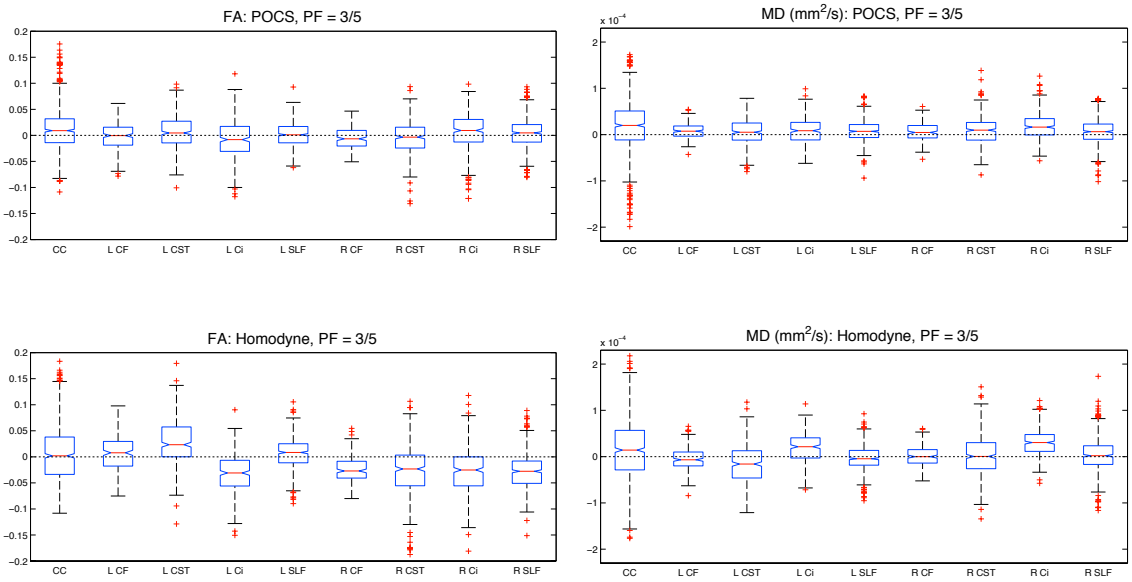

Figure 4.5: Boxplots of the difference in partial Fourier minus full k-space reconstruction, within all voxels and across all subjects for each ROI. The cases of 3/5 POCS and homodyne reconstruction of FA and MD are shown. The boxplots show median (red line), interquartile range (blue box), 95% confidence interval on the median (notches on blue box), the most extreme data points which arent outliers (black whiskers) and outliers (red crosses). A point is judged to be an outlier if it is above the upper quartile or below the lower quartile by 1.5×(interquartile range). Key: L = left; R = right. CC = corpus callosum; CF = crossing fibre (region in the centrum semiovale where the corticospinal tract, the superior longitudinal fasciculus and a callosal radiation meet); CST = corticospinal tract; Ci = cingulum; SLF = superior longitudinal fasciculus.

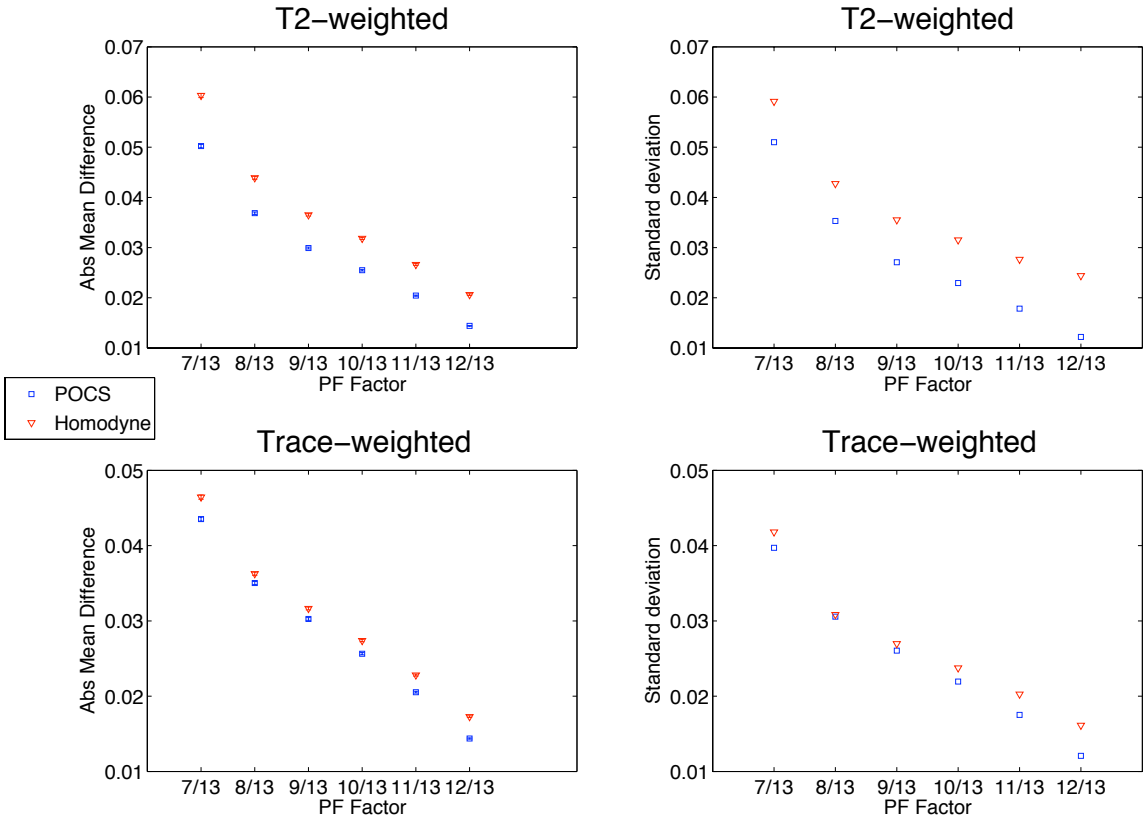

Figure 4.6: Plots of the mean absolute difference in segmented white matter with 95% confidence intervals (left column), and standard deviation of the absolute difference (right column), normalized to the full k-space signal intensity, for partial Fourier minus full k-space reconstruction of T_2- and trace-weighted images. POCS and homodyne reconstructions were used with 7-12 of the total 13 readout segments.

EPI protocol for the same resolution (2 mm isotropic, two averages, 60 directions, n = 14, 20 minute acquisition time). The ROI masks for all rs-EPI subjects were transformed to FMRIB Software Librarys standard FA space (derived from 58 subjects), combined (union), and then transformed to the native space of each ss-EPI subject. Uncertainty in the equivalent ROIs ranged from 5°-8° in ss-EPI data compared with 6°-11° in the rs-EPI data with partial Fourier reconstruction. The rs-EPI does therefore incur about 40% higher uncertainty than our highly optimized ss-EPI protocol [147] indicating that the reduced distortions in rs-EPI may come at some cost in uncertainty. The ss-EPI protocol used in our lab has been refined over a number of years using the considerations outlined in Smith *et al.* [147]. The 30 diffusion directions acquired in the rs-EPI protocol are within the suggested range, but it was not possible to match the 60 directions acquired in the ss-EPI protocol due to the lengthy scan time required. However, the rs-EPI uncertainty values are in good agreement with the uncertainties of around 4°-5° in the corpus callosum and 8°-10° in other white matter regions reported by Jones [158]. As the uncertainty on the fibre orientation dictates the performance of tractography, this demonstrates that 3/5 POCS reconstruction of rs-EPI data would not compromise probabilistic tractography analyses [150]. Acquiring 3/5 segments would reduce the scan time of the 30 direction DTI protocol from 17:19 min to 10:43 min and a 60 direction protocol from 33:49 min to 20:37 min.

Figure 4.6 shows that the POCS method has low systematic error and variance in T_2- and trace-weighted images. The normalised mean absolute differences in 9/13 partial Fourier reconstructions of T_2-weighted images were 0.030 and 0.037 and for trace-weighted images were 0.030 and 0.032 for the POCS and homodyne, respectively. These results suggest that 9/13 POCS partial Fourier reconstruction is a good compromise between scan time reduction and retained image quality. Acquiring 9/13 segments would reduce the scan time of the trace-weighted pro-

tocol from 4:48 min to 3:30 min.

4.5 Conclusions

This study shows that rs-EPI with a readout partial Fourier acquisition can be used successfully to reduce the scan time in diffusion MRI applications. A POCS reconstruction produced accurate images and estimates of diffusion parameters relative to full k-space acquisitions. Importantly, it has been shown that the difference from full k-space reconstruction is consistent across subjects and white matter tract regions, which enables comparisons to be made between patient and control groups in clinical studies. The differences between full k-space and partial Fourier reconstructions were found to be lower than effect sizes observed in previous patient versus control studies and intersubject and intersession variations. The partial Fourier rs-EPI scan times are longer than similar ss-EPI protocols by a factor approximately equal to the number of readout segments acquired (the DTI protocol with $3/5$ partial Fourier encoding is about three times longer than a similar ss-EPI protocol), and there is a small reduction in SNR efficiency in partial Fourier rs-EPI compared with a full k-space acquisition. However, the multishot approaches do have some important benefits, largely as a consequence of their short echo-spacing. The resultant reduced geometric distortions greatly improve anatomical accuracy of images. In addition, both stronger diffusion encoding and higher resolution are achievable due to the shorter TE in cases where ss-EPI is limited by T_2 signal decay of the echo-train. The reduction in imaging time of the DTI and DWI protocols possible with readout partial Fourier encoding and POCS reconstruction would allow rs-EPI to be used to acquire data with reduced geometric distortion and T_2^* blurring, relative to ss-EPI, for robust tractography studies and high-resolution clinical trace-weighted scans to identify small infarcts.

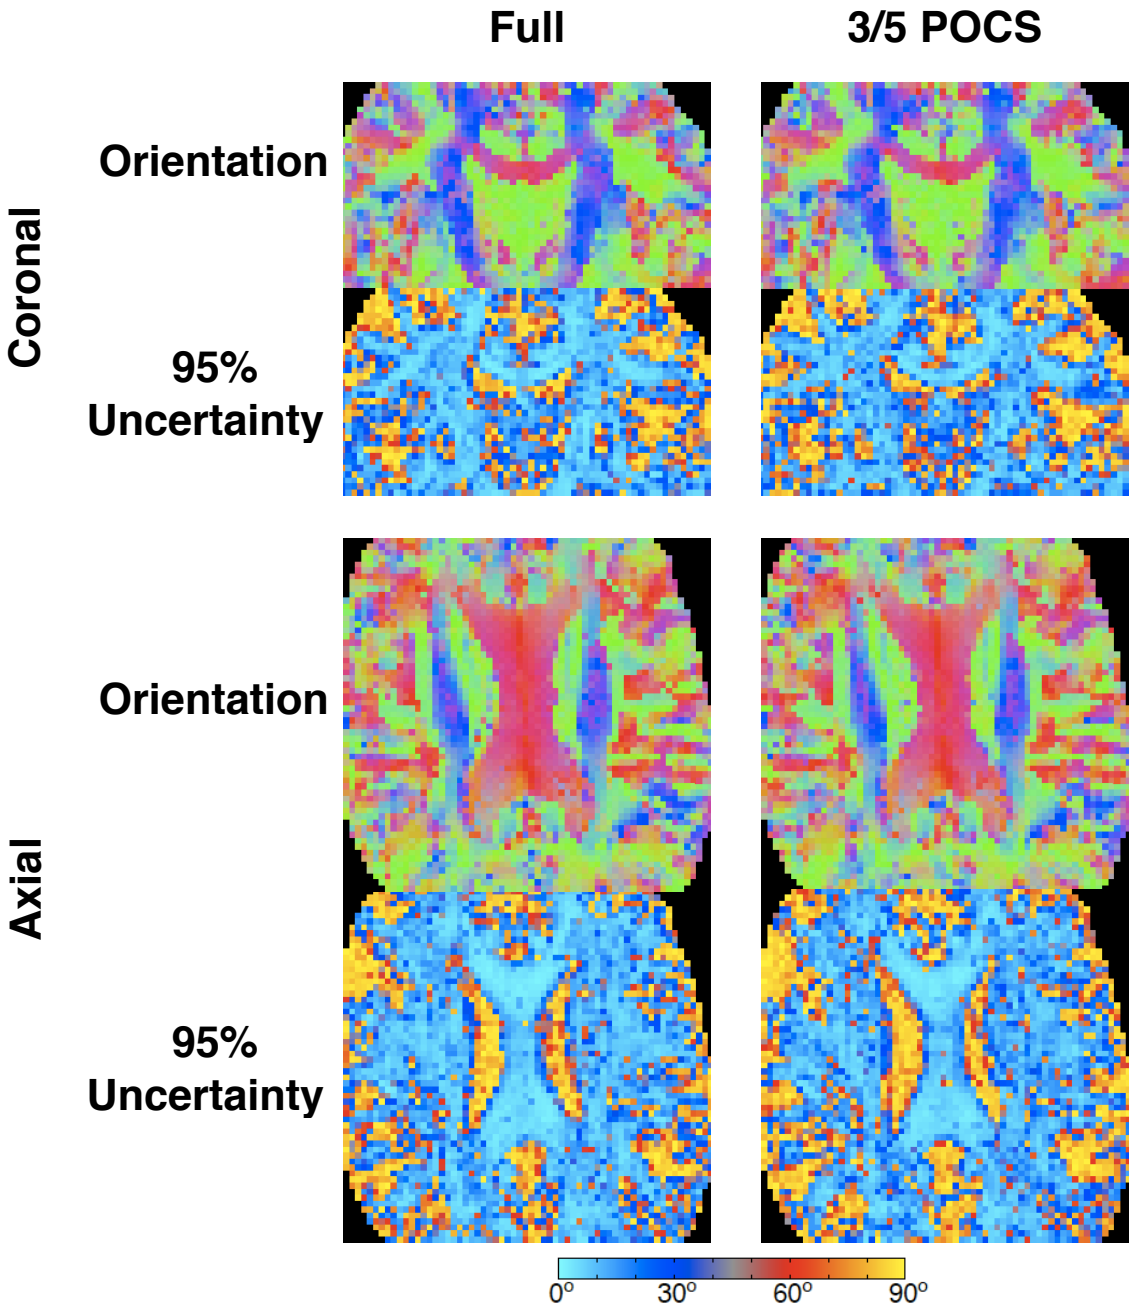

Figure 4.7: Comparison of color maps of the principle diffusion direction and corresponding maps of the 95% uncertainty angle for full k-space (left) and 3/ 5 POCS partial Fourier (right) reconstructions in representative coronal and axial slices. Left–right running tracts appear red, anterior–posterior tracts appear green, and inferior–superior tracts appear blue.

Chapter 5

Simultaneous Multi-Slice Readout-Segmented EPI

5.1 Introduction

The standard method for acquiring DW images is single-shot echo-planar imaging (ss-EPI) [35] due to its low sensitivity to motion artefact and high imaging speed. However, as discussed in Sections 2.2.1 and 3.2, ss-EPI suffers from limited achievable resolution, off-resonance artefacts and blurring that are increasing problems at higher field. The acquisition of k-space can be segmented to ameliorate these problems provided any unintended phase difference between multiple shots is accounted for with phase navigation [104, 98]. Motion of the brain during the diffusion encoding gradients creates a spatial phase variation across it that is non-linear due to the cardiac-related deformation of tissue [95, 96, 97]. The motion-induced non-linear phase in image space, or equivalently, the warping of k-space, is different for each shot so, when shots are combined, the irregular sampling in k-space causes image ghosting. The phase at each shot can be estimated with a low resolution navigator image (assuming the phase varies slowly in space)

and this estimate can be used to perform a 2D non-linear phase correction [108]. The navigator correction is simplest when the navigator satisfies the Nyquist condition so that unaliased images can be reconstructed at each shot. The PROPELLER sequence [103] and variants [159, 160, 161] fulfil this condition as does readout-segmented EPI (rs-EPI) [126, 6, 132], which can be modified to acquire an additional navigator echo at the centre of k-space after the imaging echo. In rs-EPI the navigator information is also used to identify shots that were acquired during systole and therefore have large phase errors that can not be corrected. The shots with the worst motion corruption are reacquired [131, 129] to ensure a robust acquisition of artefact-free DW images. The rs-EPI sequence segments the acquisition of k-space in the readout direction to increase the phase-encode bandwidth and shorten the readout duration, which reduces susceptibility artefacts and T_2^* blurring, respectively, in comparison to ss-EPI. However, for a given matrix size, the multi-shot acquisition makes the rs-EPI scan time significantly longer than a ss-EPI scan, approximately by the number of readout-segments, and this limits the number of slices and/or diffusion directions that can be acquired in a reasonable scan time. As discussed in Chapter 4, segments can be omitted and the missing portion of k-space can be reconstructed with a partial Fourier algorithm [162], if the associated SNR loss is deemed acceptable.

Another way to reduce scan time would be to simultaneously excite multiple slices and acquire the multiplexed signal so that fewer excitations are required to achieve the same slice coverage. This allows direct reduction of the repetition time (TR) by a factor equal to the number of slices simultaneously imaged. The simultaneous multi-slice (SMS) method was introduced by Larkman *et al.* [71] and uses the spatial sensitivity variation of multi-channel array coils [60] to separate the multiplexed slice images with a parallel imaging algorithm. The SNR of the images depends on the g-factor [65] of the reconstruction and there is no

$\sqrt{R}$ penalty (cf. conventional in-plane parallel imaging accelerated by factor R) because the data is not under-sampled. However, the g-factor related SNR loss for these slice unaliasing reconstructions can be significant, particularly in axial imaging where the slice direction usually has the smallest FOV. This means that the sensitivity variation over the 20-60 mm distance between slice-aliased voxels is usually less than what is encountered in conventional in-plane acceleration. For example, for axial slices with 190-220 mm FOV and in-plane acceleration factor R_{PE} = 2 the distance between aliased voxels is 95-110 mm. Also, coil designs tend to have fewer elements in this direction, which reduces sensitivity variation and hence parallel imaging performance. The CAIPIRINHA method [72] increases this sensitivity variation between slice-aliased voxels by alternating the RF phase of phase-encode lines to introduce a differential in-plane shift in the spatial domain of the slices, such that the aliased locations of two slices no longer contain the same coil sensitivity information. The same principle can be applied to EPI by using slice gradients (G_z) during the readout to create in-plane slice shifts [163, 73] but these methods suffer from voxel blurring artefacts (which increase with larger slice shifts) due to a tilted readout or phase-encode direction. The blipped-CAIPI method [164, 74] has addressed these artefacts by using alternating polarity G_z blips to avoid accumulation of phase at the edges of slices during the readout. In the absence of this voxel blurring the parallel imaging performance of the receive coil can be exploited in two directions to reduce g-factor related SNR loss and the improvements have been demonstrated with R_{slice} = 3 acceleration.

Recent studies have used SMS acceleration to sample higher angular resolution or more b-values for DTI or higher temporal resolution for functional MRI [76, 165, 166, 74, 167]. This chapter demonstrates that the blipped-CAIPI technique can be applied to rs-EPI in order to allow acquisition of high-resolution DW data with reduced distortion and blurring in reasonable scan times. We recently

presented a preliminary account of this work [138] at the annual meeting of the ISMRM. Simultaneous multi-slice has already been applied to ss-EPI to increase SNR efficiency further. However the problems of limited resolution, distortions and poor point-spread function of single shot acquisitions remain, whereas the application to rs-EPI would make scan times more feasible so that the improved data quality could be used for clinical stroke diagnosis and tractography.

This chapter explains the blipped-CAIPI method, and the modifications and considerations required for application to the rs-EPI sequence. Then R_{slice} = 2 slice-accelerated acquisitions are demonstrated in DTI and high-resolution trace-weighted images and rs-EPI image quality and tractography are compared to ss-EPI.

5.2 Theory

The blipped-CAIPI technique is described extensively in the recent paper by Setsompop *et al.* [74]. In this section, details of the method relevant to the implementation and reconstruction of blipped-CAIPI rs-EPI are discussed, including the slice-shifting gradient scheme, the reconstruction and the combination with the navigator phase correction. R_{slice} = 2 slice acceleration was implemented resulting in acquisitions with TR times in the range of 2.7-4.8 s. SNR efficiency (signal/$\sqrt{\mathrm{TR}}$) for a spin-echo sequence is plotted as a function of TR in Fig. 3.8 for white matter tissue T_1= 1000 ms, T_2= 75 ms [136, 137] and TE = 70 ms. At 3 T, literature values of T_1 for white matter and grey matter (the tissues of interest in DTI and trace-weighted acquisitions) are ~850-1100 ms and ~1300-1800 ms [136, 137], respectively, so we expect the slice acceleration to increase SNR efficiency. The TR times are also longer than the T_1 values so the efficacy of interleaved slice excitation should not decrease.

5.2.1 Slice-Shifting Gradient Scheme

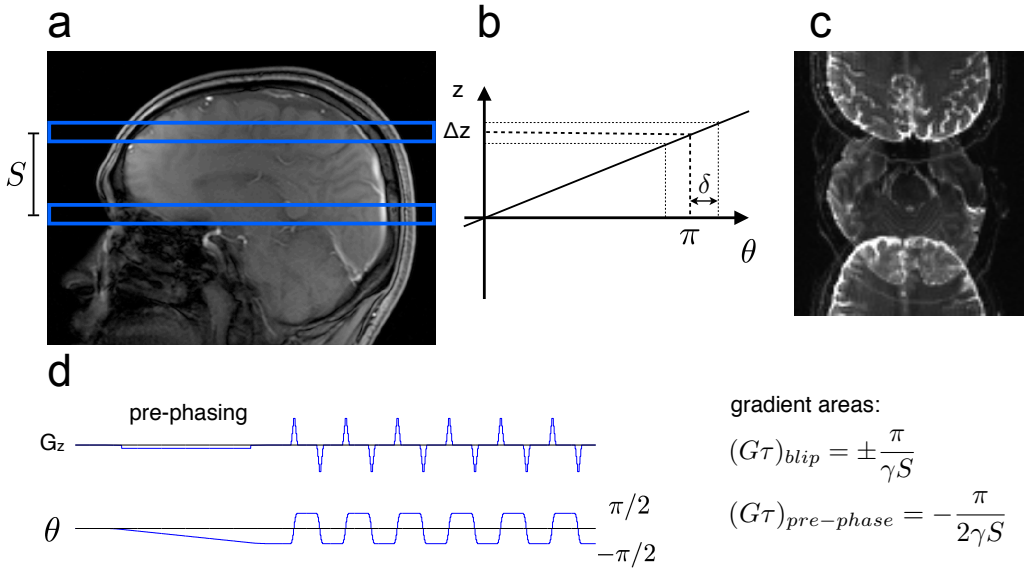

Figure 5.1: Illustration of how G_z blip gradients played out in conjunction with phase-encode blips (not shown) during the EPI readout are used to shift the upper slice by FOV/2. The lower slice is at gradient isocentre. A π phase shift is created between phase-encode lines for signal from the upper slice. The alternating polarity ensures that the phase does not accumulate across the finite thickness of the slice during the readout. In the FOV/2 gradient scheme the pre-phasing blip is half the gradient area of the G_z blips.

An image space slice-shift in the phase-encode direction can be achieved by creating a linear phase ramp across k-space. G_z blips can be played out together with the phase-encode blips to impart different amounts of phase at different slice locations as illustrated in Fig. 5.1 for two slices where one is at gradient isocentre. The slice at isocentre experiences no phase variation due to G_z while the upper slice experiences a π phase variation between phase-encode lines which is plotted in Fig. 5.1d. This results in a FOV/2 shift of the upper slice while the slice at isocentre remains in its original position in the FOV. The gradient area for a FOV/2 shift is calculated by requiring the phase difference, θ, at slice position

$z = S$ to be π. Rearranging Eq. 5.1, the blip gradient area is given by $G\tau = \pi/\gamma S$.

$$\theta = \gamma \int_0^t G_z z \, \mathrm{d}t = \gamma G \tau S = \pi \tag{5.1}$$

As illustrated in Fig. 5.1a and b, due to the finite slice thickness Δz, the phase at the edge of a slice differs by $\pm\delta$ from the phase at the centre. In the method of Nunes *et al.* [73], the polarity of the blips was constant so phase at the edge of slices accumulated during the readout, causing a blurred spatial point-spread function. In blipped-CAIPI, the alternating polarity and the pre-phasing gradient (which is half the gradient area of the G_z blips) before the readout, shown in Fig. 5.1d, ensures that the phase at the edges of slices alternates between $\pm\delta/2$, which avoids voxel blurring and only causes minimal signal loss [74].

5.2.2 Blipped-CAIPI Reconstruction

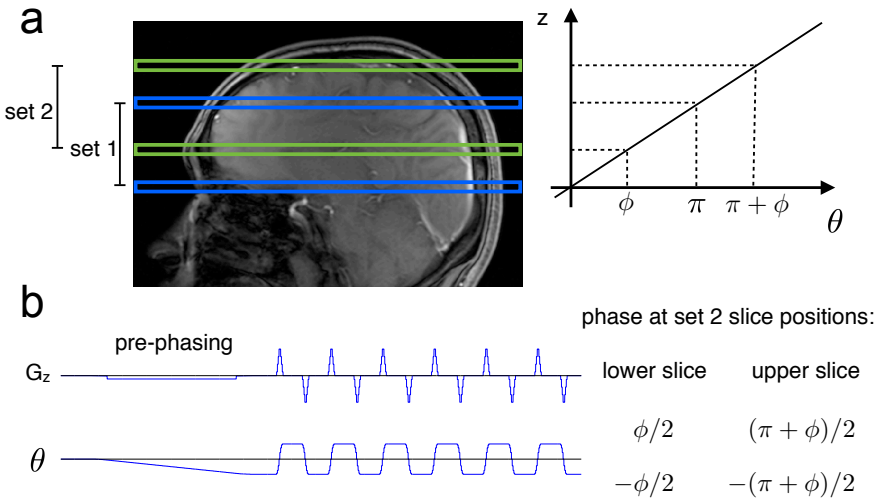

Figure 5.2: Phase correction for off-isocentre slices. $\phi/2$ phase is added to odd lines and subtracted from even lines at the start of the reconstruction. The lower slice of set 1 is at G_z isocentre (z = 0).

In the previous section, the application of G_z blips for a FOV/2 shift gradient scheme was demonstrated for the case where one slice is at gradient isocentre.

Figure 5.2 shows the phase-variation for a set of slices which is not at isocentre (set 2, coloured green) when the G_z blip gradients are applied. The lower slice of set 2 should have zero phase so that it is unshifted and the upper slice should have $\pm\pi/2$ phase so that it is shifted by FOV/2, which is the case for set 1. From Fig. 5.2b, the difference between on- and off-isocentre slice sets is thus a $\phi/2 = (z/2S)\pi$ phase shift (where z is the distance of the lower slice from isocentre and S is the slice separation of simultaneously excited slices) that needs to be added to odd lines and subtracted from even lines. This phase can be simply removed in reconstruction.

To separate the multiplexed slices into single-slice data a SENSE [65] parallel imaging reconstruction can be applied. However SENSE reconstructions can be more susceptible to motion between acquisition of the reference data (coil sensitivity maps) and the under-sampled data than GRAPPA [66]. A SENSE-GRAPPA combination [75] has been used to unalias simultaneously acquired slices but this reconstruction results in artefacts when there is an in-plane shift that causes slices to wrap around the FOV. Setsompop *et al.* [74] explained this artefact and introduced the slice-GRAPPA reconstruction that cleanly unaliases shifted slices. The slice-GRAPPA reconstruction was used in this study and a schematic of the reconstruction is shown in Fig. 5.3 with a comparison to the more common in-plane GRAPPA reconstruction. Separate kernels, or weight sets, are fit in k-space according to Eq. 5.2 from the multiplexed data to single-slice reference data of each of the simultaneously excited slices. In the reconstruction of blipped-CAIPI rs-EPI data, the kernels were fit using the central segments of the b = 0 data and were then applied to all segments of subsequent b = 0 and diffusion-weighted volumes of the slice-accelerated, multiplexed data to reconstruct single-slice images.

$$S_{j,z}(k_x, k_y) = \sum_{l=1}^{L} \sum_{b_x=-B_x}^{B_x} \sum_{b_y=-B_y}^{B_y} n_{j,z,l}^{b_x,b_y} S_l(k_x - b_x\Delta k_x, k_y - b_y\Delta k_y) \tag{5.2}$$

In Eq. 5.2, each point in the single-slice reference at (k_x, k_y) in the jth coil is

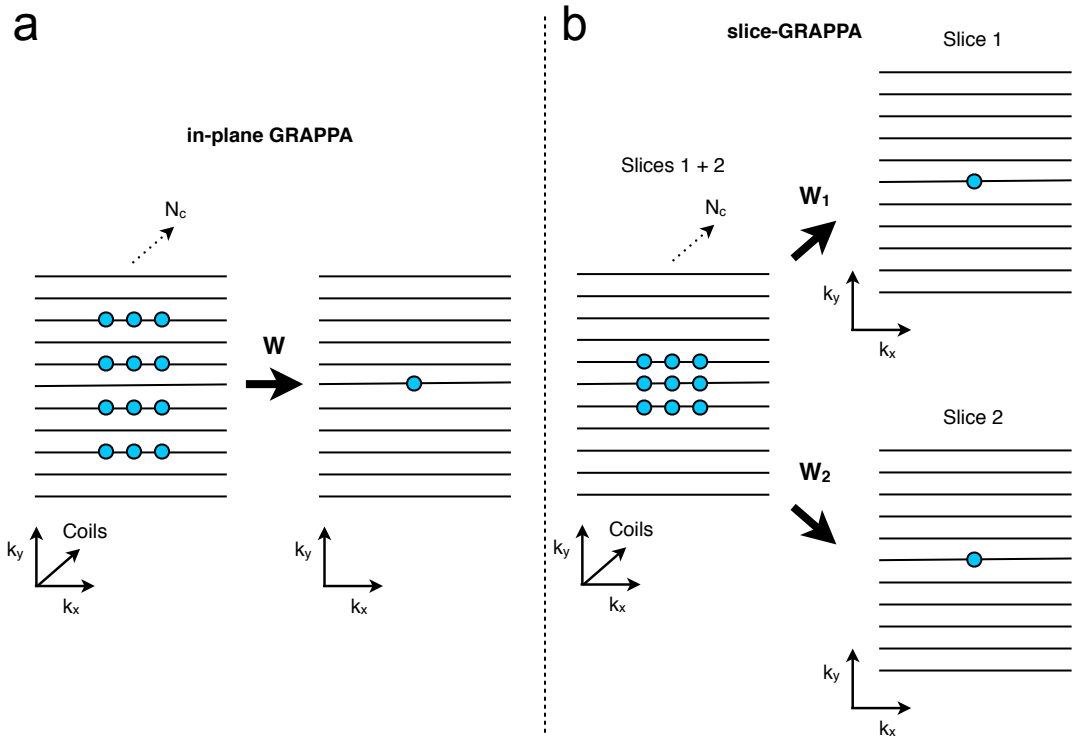

Figure 5.3: Illustration of (a) in-plane GRAPPA and (b) slice-GRAPPA parallel imaging reconstructions. In (a) each point is estimated from the surrounding points, in all N_c coils, using a kernel, W, that is estimated from fully sampled reference data acquired in a pre-scan. In (b) separate slice data are reconstructed from surrounding points, in all N_c coils, of the multiplexed data (slices 1 + 2) using separate kernels, W_1 and W_2. The kernels are estimated in separate fits from the multiplexed data to each single slice reference data acquired in a pre-scan, i.e., slices 1 + 2 are fit to slice 1 to estimate W_1 and slices 1 + 2 are fit to slice 2 to estimate W_2.

fit to the points surrounding (k_x, k_y) (defined by the kernel size B_x, B_y) in all L channels of the multiplexed data, for each single-slice reference z (i.e. $z = 1, 2$ for two simultaneously excited slices). The kernel $n_{j,z,l}^{b_x,b_y}$ is applied separately for each z to unalias the multiplexed data. Schematics of the reconstructions for in-plane and slice-GRAPPA are shown in Fig. 5.3a and Fig. 5.3b, respectively.

5.2.3 Multi-Shot Diffusion Corrections

The normal rs-EPI navigator correction can be performed after unaliasing the readout segments and their navigator echoes into single-slice data. The segments and navigator echoes are Fourier transformed to image-space and the phase of the navigator is subtracted from the corresponding segment's phase. The segments are Fourier transformed back to k-space so that they can be spliced together to form the full k-space.

It was necessary to test the navigator-based reacquisition because motion corruption is judged at run-time on the scanner. A metric called the distribution width [6] judges the degree of motion corruption by quantifying the k-space spread of the navigator. The distribution width, given in Eq. 5.3, was chosen because it judges motion corruption robustly and it is a fast calculation that can be performed on the raw k-space data in real time without image reconstruction.

$$W_x = \sum_i \sum_{k_y} \sum_{k_x} |S_i(k_x, k_y)| \cdot |k_x - \Omega_x| \tag{5.3}$$

In Eq. 5.3, S_i is the complex k-space signal in each receive channel i and Ω_x is the k_x coordinate of the location with the maximum signal amplitude across all channels. In blipped-CAIPI rs-EPI, this k-space data is slice multiplexed, which could complicate detection of motion corruption in either of the multiplexed slices. As demonstrated in Section 5.4 below, it seems that the sum over channels allows the corruption to be detected in either of the slices.

5.3 Methods

5.3.1 Pulse Sequence and Image Reconstruction

The rs-EPI sequence was modified as shown in Fig. 5.4 to play out multiband RF pulses and the FOV/2 slice gradient scheme described above. Windowed sinc RF

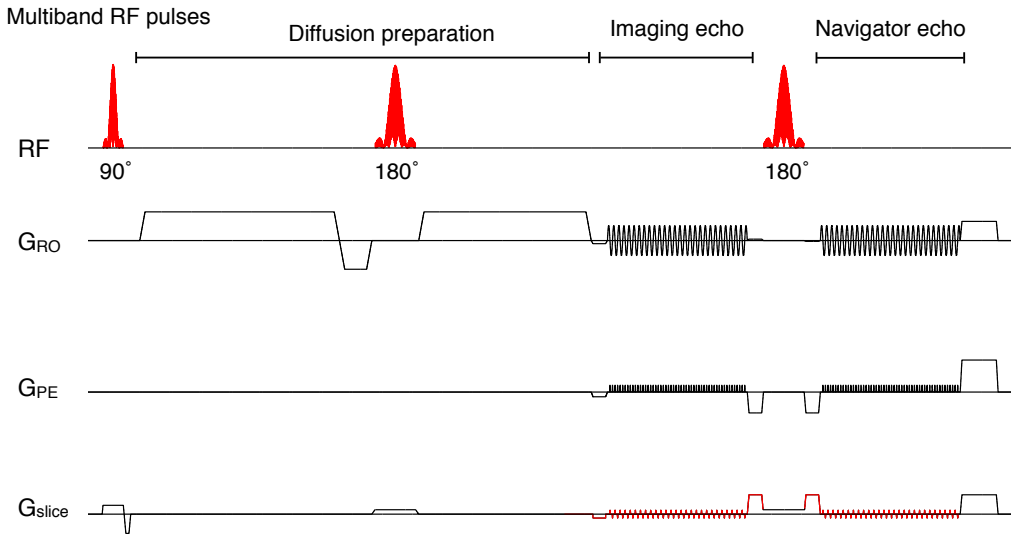

Figure 5.4: Blipped-CAIPI rs-EPI sequence with modifications to the original rs-EPI coloured red. These are multiband RF pulses and the slice-GRAPPA FOV/2 slice-gradient scheme.

pulses were phase modulated and summed in the sequence code to select multiple slices based on a slice separation parameter. The excitation and refocussing pulse durations were 2.56 ms and 5.12 ms, respectively, which ensured that the pulses did not exceed the subject-specific maximum RF voltage and SAR. The diffusion preparation module was a modified monopolar scheme [117] to allow a shorter TE.

Images were reconstructed off-line in Matlab (Mathworks, Natick, MA, USA) from raw scanner data. The data acquisition and reconstruction pipeline is shown in Fig. 5.5. The basic rs-EPI reconstruction used was the same as described in Chapter 4 [162] except for the addition of the blipped-CAIPI isocentre phase

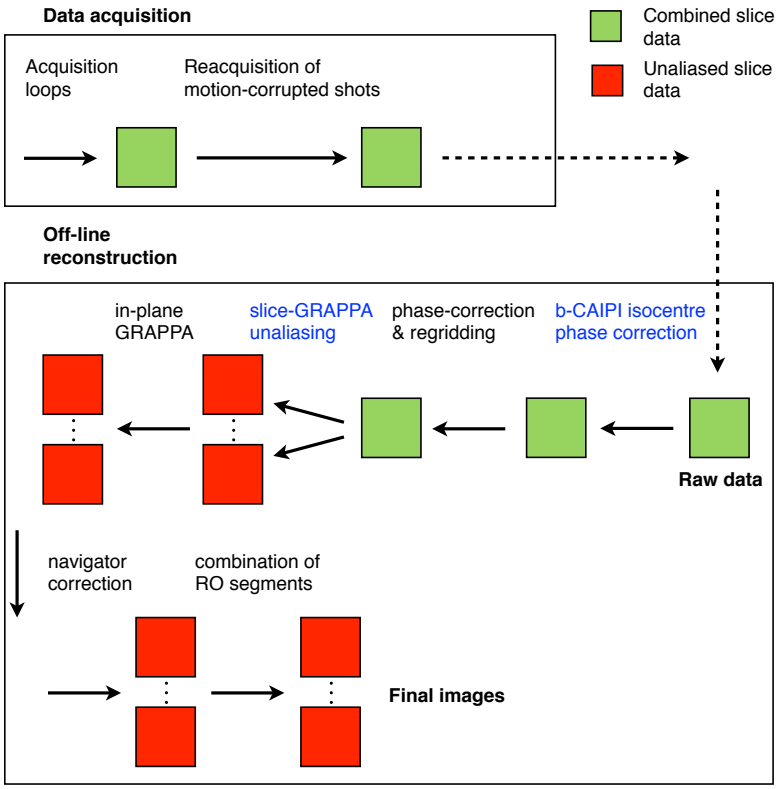

Figure 5.5: Reconstruction flow chart with blipped-CAIPI modifications coloured blue. Data was acquired as the combined signal (coloured green) from multiple slices and the reacquisition was based on this multiplexed navigator data. The multiplexed data were separated into single slice data (coloured red), with the slice-GRAPPA unaliasing reconstruction, before the in-plane GRAPPA reconstruction, navigator correction and combination of readout segments.

correction and slice-GRAPPA unaliasing, where a 3×3 kernel was used. Also, Nyquist ghost correction was performed before regridding, rather than after, which was found to be a slightly more robust correction.

5.3.2 Data Acquisition

The diffusion-weighted blipped-CAIPI rs-EPI sequence and a ss-EPI sequence were used to acquire data from two healthy volunteers with a MAGNETOM Verio 3 T scanner (Siemens Healthcare, Erlangen, Germany) under an approved technical development ethics protocol. Slice acceleration is necessary to acquire rs-EPI DTI data in reasonable scan times so it was implemented in the rs-EPI sequence. Slice-acceleration in the ss-EPI sequence can be used to increase the number of diffusion directions and/or SNR but it was not implemented because the motivation for the DTI acquisitions was to compare the rs- and ss-EPI data quality at higher resolution. To achieve increases in resolution with ss-EPI, larger matrix sizes or more in-plane acceleration can be used, which degrade images further. Higher data quality can be obtained by improving the time efficiency of rs-EPI. The blipped-CAIPI rs-EPI sequence was used to acquire two DTI protocols with 60 isotropically distributed b = 1000 s/mm^2 diffusion directions [148] at 2 mm and 1.5 mm isotropic resolution on one subject. Protocols with the same resolution and diffusion directions were also acquired with the ss-EPI sequence for comparison and three ss-EPI averages were acquired to match the scan time of the blipped-CAIPI rs-EPI. Two blipped-CAIPI rs-EPI DWI protocols were acquired with 1.1×1.1×4 mm and 0.9×0.9×4 mm in two subjects. A 1.1×1.1×4 mm blipped-CAIPI rs-EPI DWI protocol was acquired in a patient who had suffered a transient ischemic attack (TIA) the previous day. The patient was 82 years old and suffered from advanced white matter disease. The scan was at the end of a full clinical protocol that included a standard ss-EPI DWI scan for 1.8×1.8×5

Sequence	Resolution (mm)	FOV (cm^3)	matrix	ES (ms)	R_{slice}/R_{PE}	TR/TE (s)/(ms)	PF	Scan time (min)
rs-EPI (5 seg)	2 iso	22×22×12.6	110×110×58	0.32	2/2	4.1/73	-	36
ss-EPI	2 iso	19×19×12.6	96×96×58	0.95	-/2	8.4/87	6/8	33 (3 av)
rs-EPI (7 seg)	1.5 iso	22×22×9.3	144×144×62	0.32	2/2	4.8/79	-	54
ss-EPI	1.5 iso	19×19×9.3	128×128×62	0.70	-/2	13.9/88	6/8	54 (3 av)
rs-EPI (9 seg)	1.1×1.1×4	22×22×12	192×192×30	0.32	2/2	2.7/70	-	2:30
rs-EPI (11 seg)	0.9×0.9×4	22×22×12	256×256×30	0.32	2/2	3/76	-	3:20

Table 5.1: Acquisition parameters for all sequences

mm resolution trace-weighted images. To compare phase-encode blurring 1.5 mm isotropic ss- and rs-EPI b = 0 images of a resolution phantom were acquired. Acquisition parameters of all the sequences used are shown in Table 5.1.

5.3.3 Image Processing and Analysis

For the two blipped-CAIPI rs-EPI and two ss-EPI DTI datasets the following post-processing pipeline was applied. Eddy current correction was applied to the images using FMRIB's Diffusion Toolbox [150], which consists of an affine registration (12 degrees of freedom) of all volumes to the initial b = 0 s/mm^2 volume. Non-brain tissue was removed from the images using Brain Extraction Tool [151], and the diffusion tensor was fitted to the reconstructed data using FMRIB's Diffusion Toolbox. Fractional anisotropy (FA) maps were calculated and for tractography Bayesian estimation of diffusion parameters obtained using sampling techniques [152] was used to generate probability density distributions on two crossing fibre orientations at each voxel. Tractography was seeded in the subject's structural space to generate tracts for several association, commisural and projection fibres. Seed masks were defined in the superior longitudinal fasciculus (SLF), cingulum bundle, optic tract, fornix, corpus callosum (CC) and corticospinal tract (CST). Also, SNR and blurring was quantified in the raw images from the DTI protocol . SNR was estimated from the 12 b = 0 s/mm^2 images by dividing the temporal mean signal by the temporal standard deviation in signal. The mean

and standard deviation within the brain extracted mask were taken as the SNR estimate and the error, respectively. Resolution was estimated on the raw diffusion weighted images using the *smooth_est* tool in the FMRIB Software Library (www.fmrib.ox.ac.uk). The mean estimate across all diffusion-weighted images was calculated (one average of ss-EPI was analysed) and normalised to the rs-EPI estimate within each resolution. Echo-train parameters were used to calculate the theoretically expected percentage of voxel blurring and the effective phase-encode resolution, $\Delta y'$, using the following formula for spin-echo EPI [7]:

$$\Delta y' = \frac{\sqrt{3}}{\pi}\Delta y\frac{T_{acq}}{T_2} + \frac{1}{\pi}\Delta y\frac{T_{acq}}{T_2'} \tag{5.4}$$

where Δy is the nominal voxel size, T_{acq} is the duration of the readout and T_2 and T_2' are the decay constants of the thermodynamic and external field induced effects, respectively. The total loss of resolution is approximated as the result of two separate T_2 and T_2' filters. A white matter T_2' value of 114 ms was calculated using Eq. 5.5 with T_2^* = 47 ms and T_2 = 80 ms [136].

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{5.5}$$

For the DWI data, trace-weighted images were calculated by taking the geometric mean of the three b = 1000 s/mm^2 images acquired with orthogonal diffusion weightings.

5.4 Results

An example of detection of motion corruption with $R_{slice} = 2$ acceleration during the on-line reacquisition is shown in Fig. 5.6. In Fig. 5.6a, before reacquisition the navigator for readout-segment number 3 (the central readout segment)

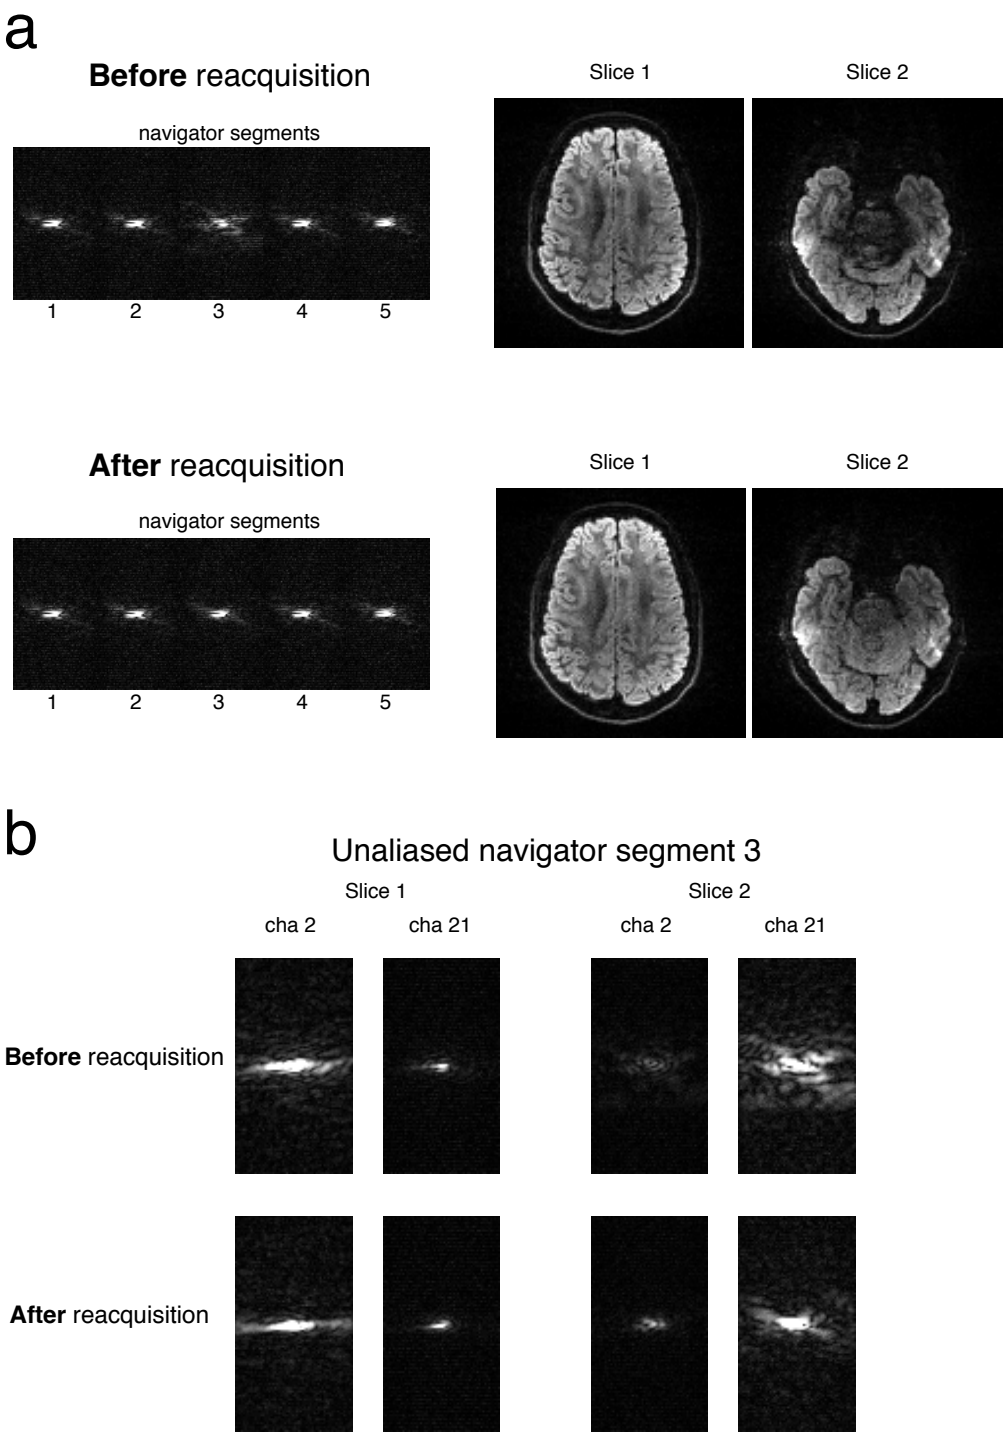

Figure 5.6: Navigator-based reacquisition with two simultaneously excited slices with b = 1000s/mm^2 diffusion-weighting. (a) Before reacquisition, corruption in segment 3 is indicated by the dispersed k-space in the navigator, which contains multiplexed data from the two slices. When segment 3 is reacquired artefacts are removed from the images. (b) Single-channel maps (all at the same scale) of the unaliased k-space from navigator segment 3. Two channels are chosen which are close to slice 1 (channel 2) and 2 (channel 21) to demonstrate the variation in signal.

shows more dispersed k-space signal than the other navigators. This segment was judged to be one of the 20% most corrupted of all the segments so it was marked for reacquisition. When this segment was reacquired, there was less motion corruption indicated by the reduced navigator spread and the dispersed signal in the centre of slice 2 is restored. This example, which was acquired with diffusion encoding close to the superior-inferior axis of the brain, highlights the nature of motion-phase artefacts in diffusion which are mainly due to the pulsatile deformation of the brain during systole [95, 96, 97, 108]. It also shows the importance of having cleanly acquired data close to the centre of k-space to minimize artefacts and for this reason the navigator reacquisition is preferentially weighted to central segments. In Fig. 5.6b, the segment marked for reacquisition (navigator segment number 3 from Fig. 5.6a) has been unaliased and the signal is shown from coil elements 2 and 21 which are close to slices 1 and 2, respectively. For example, in the unaliased slice 2 navigators the proximity of the slice to the coil element results in high signal in channel 21. These unaliased navigators are included to show that navigator correction could be judged separately on each of the simultaneously excited slices. With higher slice acceleration factors, identification of motion corruption may only be possible in unaliased slice data and this would be the subject of future work. If the slice-GRAPPA unaliasing weights were calculated at the start of the scan, all navigators could be separated into single-slice data so that separate navigator k-space distribution widths could be calculated for each slice in a time-efficient, and hence real time manner. The standard reacquisition algorithm sums all coil contributions when evaluating the distribution width. As an alternative for unaliasing the navigators it might also be possible to make use of coil-specific k-space distributions widths to consider the signal (and hence corruption) from more localised regions of the coil.

Raw images acquired with the DTI protocols are shown in Fig. 5.7 and 2 mm

and 1.5 mm blipped-CAIPI rs-EPI FA maps are shown in Fig. 5.9. These demonstrate the reduced blurring and distortion in the rs-EPI data due to the shorter readout duration but also the higher SNR which is achievable with three ss-EPI averages in the same scan time. The increased blurring in ss-EPI is quantified in Table 5.2, which presents a comparison of b = 0 s/mm^2 SNR and blurring in the raw diffusion-weighted images. The SNR is dependent on the point-spread function because more blurring corresponds to higher SNR. A further comparison of phase-encode blurring at 1.5 mm nominal resolution is shown in Fig. 5.8. The increased blurring and distortion in ss-EPI is evident in the images shown in Fig. 5.8a. The profiles plotted in Fig. 5.8b and c highlight the ability of rs-EPI to resolve small structures, such as the comb in profile A and the dip in signal in the structure at y voxels 5-10 in profile B. Signal that is represented by one voxel in the rs-EPI image is blurred into the neighbouring two voxels in the ss-EPI image.

Trace-weighted images with 0.9 mm and 1.1 mm in-plane resolution are shown in Fig. 5.10. Comparison of the ss-EPI and blipped-CAIPI rs-EPI trace-weighted images of a hyper-intense lesion in a TIA patient are shown in Fig. 5.11. The ss-EPI data was acquired with 5 mm slices with 1.5 mm gaps between slices and the blipped-CAIPI rs-EPI data had 4 mm slices with no inter-slice gaps. This explains the different appearance of the slices in the comparison of the images. Also, the ss-EPI data was acquired with a 12-channel head coil and the blipped-CAIPI rs-EPI data was acquired with a 32-channel coil (which is necessary for improved parallel imaging performance).

Tractography generally improved at 1.5 mm compared to 2 mm in both the rs- and ss-EPI, producing more true positive tracks and fewer false positive tracks (true/false positives were determined with the help of an expert anatomist, data not shown). For example, unlike for 2 mm rs- and ss-EPI, tractography seeded in the anterior pillars of the fornix correctly tracked into the hippocampus and

Sequence	Nominal resolution (mm)	b=0 SNR	Estimated normalised resolution	Theoretical % voxel blurring	Calculated effective PE resolution (mm)
rs-EPI	2	13(8)	1.00	17	2.34
ss-EPI (1 av)	2	14(8)	1.51	24	2.49
rs-EPI	1.5	8(5)	1.00	22	1.83
ss-EPI (1 av)	1.5	13(9)	1.78	44	2.16

Table 5.2: SNR and resolution analysis on the raw DTI images. SNR calculated from the repeated b=0 images and resolution in the diffusion-weighted images from a smoothness estimator normalised to the rs-EPI value. Theoretical % voxel blurring and the effective phase-encode resolution were calculated from echo-train parameters.

did not mistrack into the anterior commisure in the 1.5 mm data. Overall, tractography in the 1.5 mm rs- and ss-EPI data sets was comparable in terms of the ratio of true- to false-positives (correct vs. incorrect tracts). However, the streamlines generally died out earlier in the rs-EPI data, such that the total number of streamlines reaching cortex (whether correct or incorrect) was greater in ss-EPI. This effect is thought to be attributable to the higher SNR in the ss-EPI data. One example where the rs-EPI yielded improved tractography is the CST. Maximum intensity projections (within-hemisphere) of the CST are shown in Fig. 5.12 overlaid onto representative slices of the structural image. Streamlines from the right and left hemisphere seeds are shown in red-yellow and blue, respectively, at the same threshold (>100 streamlines). The ss-EPI mistracks into anterior regions via the SLF (arrows), while the rs-EPI delineates the CST projections into M1 and S1 cortex much more cleanly. The difference between rs-EPI and ss-EPI tractography (both thresholded at >100 streamlines) of the left and right CST in representative slices is shown in Fig. 5.13. Regions where there are more rs-EPI or ss-EPI tracts are shown in red-yellow or blue-light blue, respectively. In regions that are plausible segmentations of CST rs-EPI tracts are dispersing less than ss-EPI. Therefore, generally and in the example of Fig. 5.13, in the centre of tracts there are more rs-EPI streamlines and at the periphery there is a tendency to find more ss-EPI tracts, which could be a consequence of the lower blurring in rs-EPI.

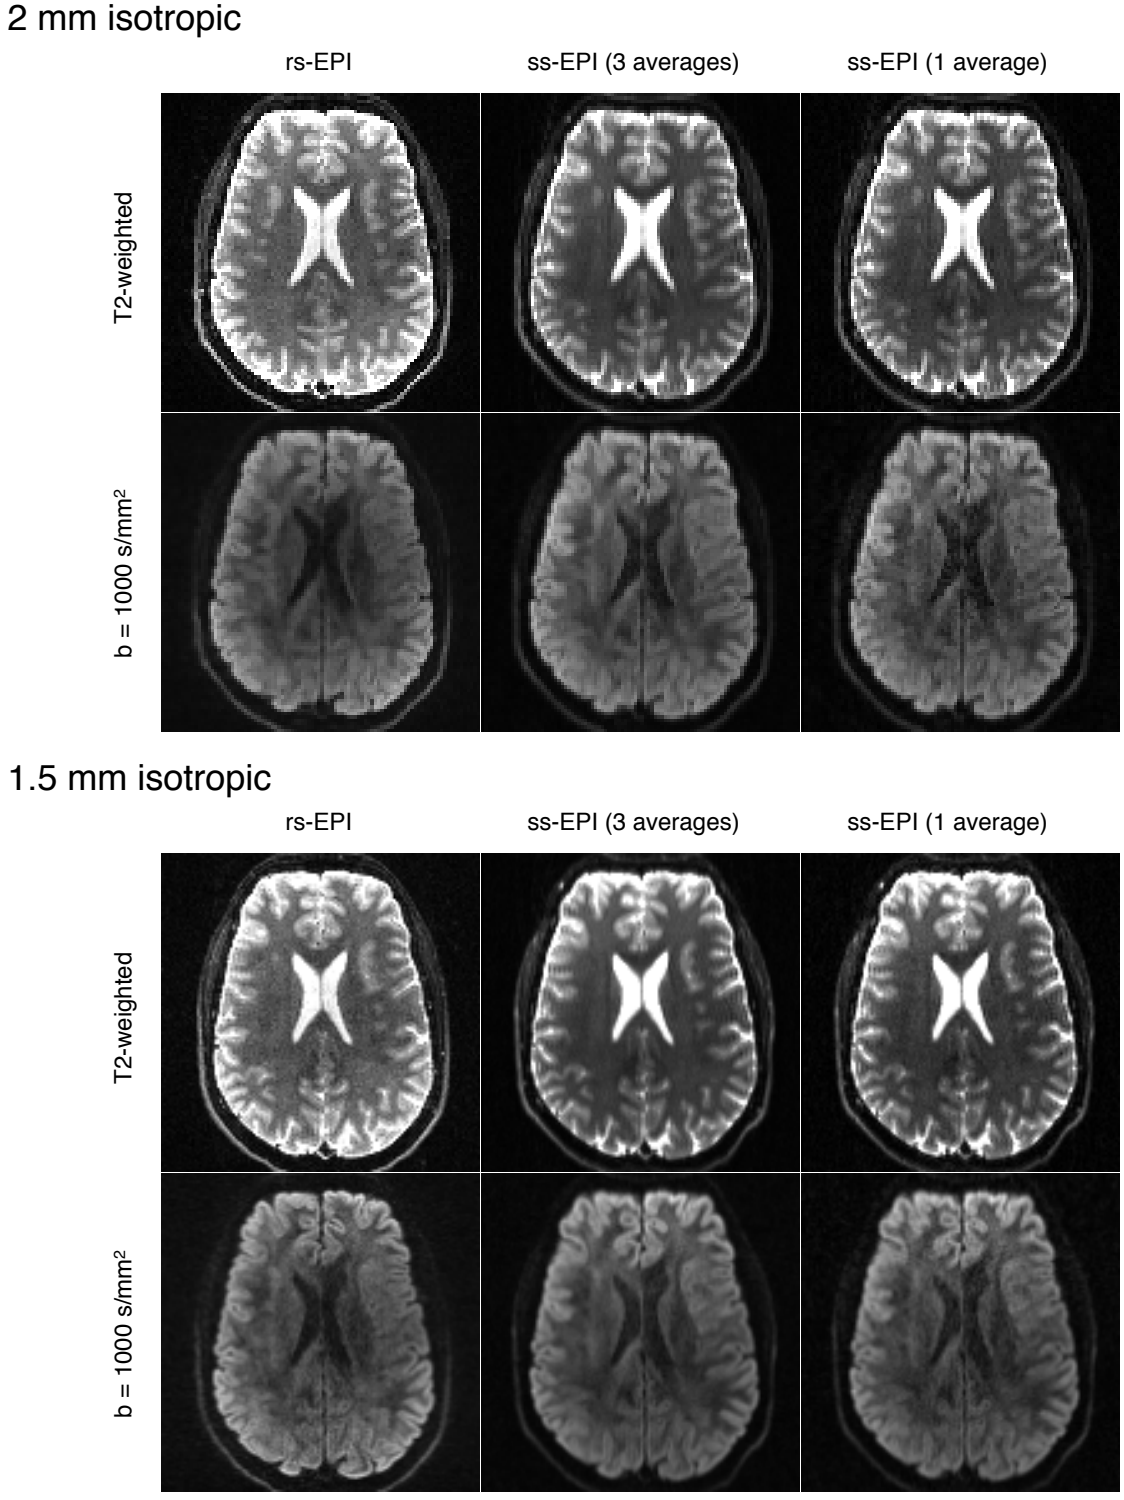

Figure 5.7: Examples of raw images T_2-weighted and DW images acquired with the blipped-CAIPI rs-EPI and ss-EPI DTI protocols with 2 mm and 1.5 mm isotropic resolutions.

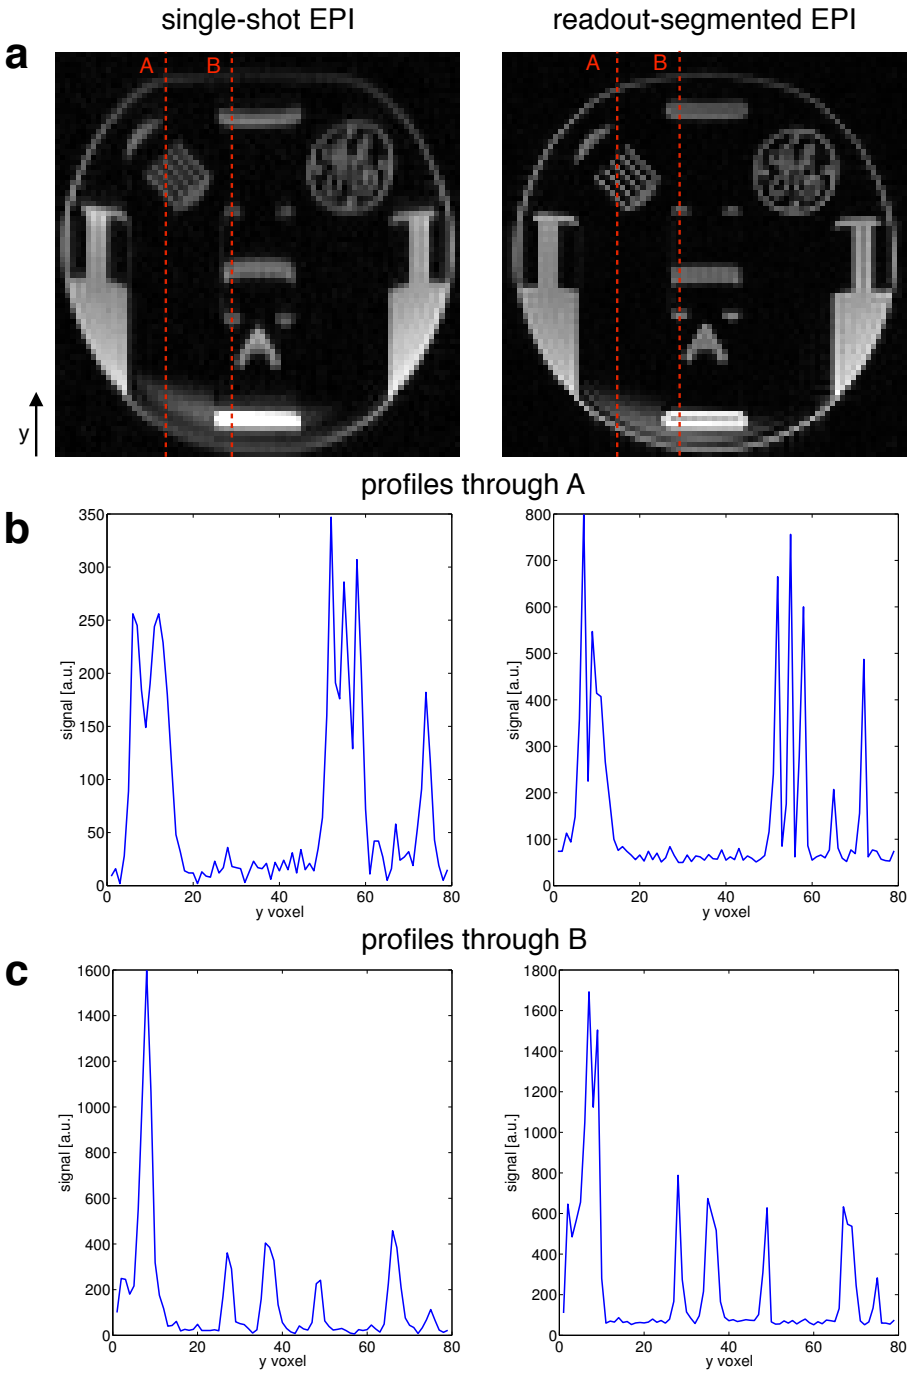

Figure 5.8: Comparison of blurring in the phase-encode (y) direction at nominal resolution of 1.5 mm between ss-EPI (left column) and rs-EPI (right column). (a) ss- and rs-EPI images of a resolution phantom indicating the positions of the profiles plotted below. Profiles in y through lines A and B are plotted in (b) and (c), respectively.

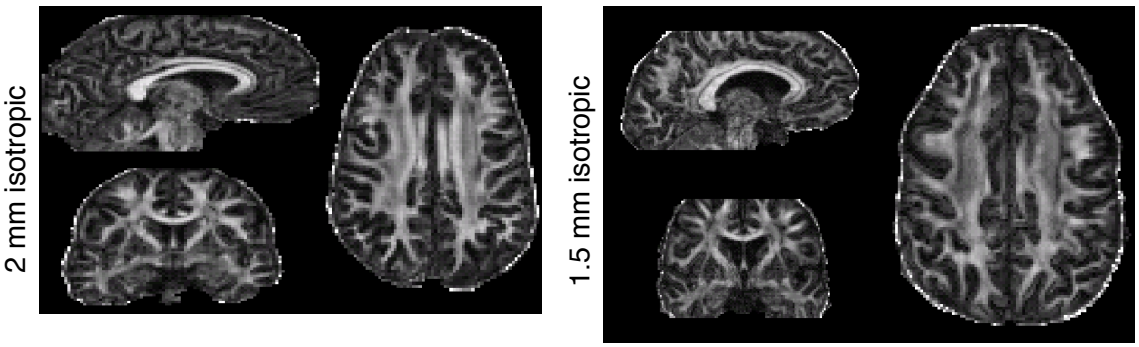

Figure 5.9: blipped-CAIPI rs-EPI FA maps with 1.5 mm isotropic resolution.

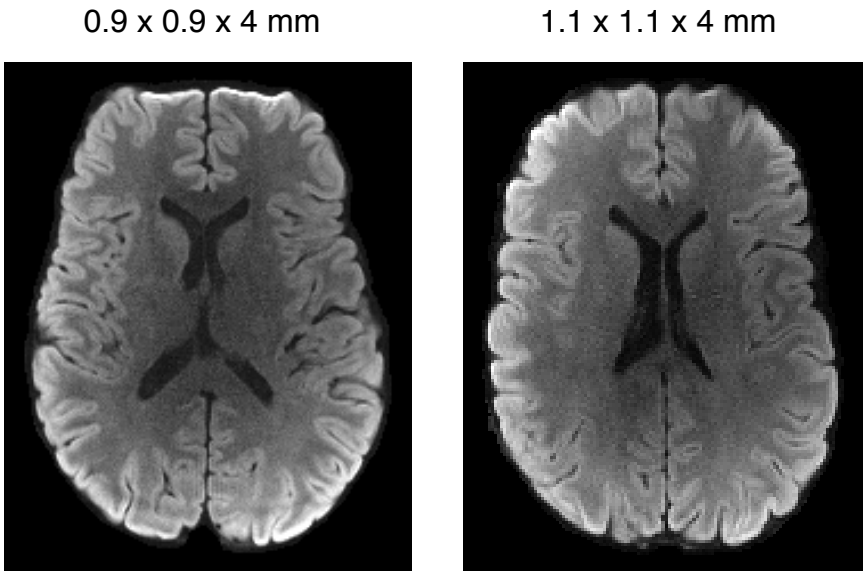

Figure 5.10: blipped-CAIPI rs-EPI trace-weighted images with 0.9 and 1.1 mm in-plane resolution.

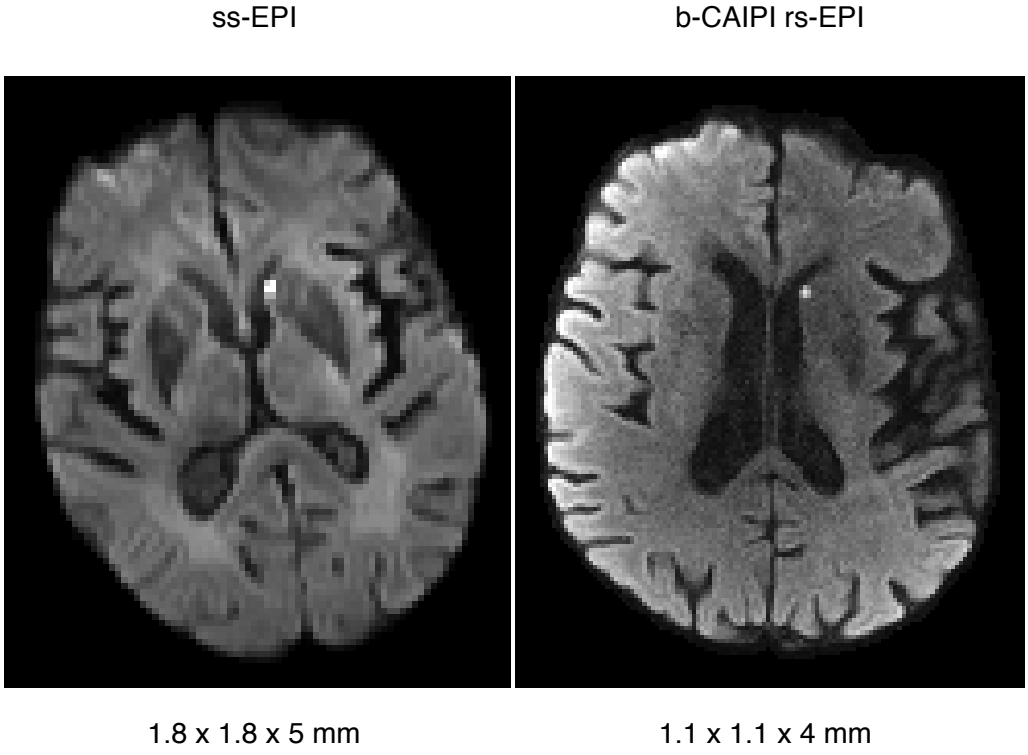

Figure 5.11: Comparison of a 1.8×1.8×5 mm ss-EPI and 1.1×1.1×4 mm blipped-CAIPI rs-EPI trace-weighted images showing a lesion in a patient who had suffered a TIA the previous day. The ss-EPI protocol acquired one T_2-weighted volume and 12 diffusion directions and 3 averages in 2:30 min. The blipped-CAIPI rs-EPI acquired one T_2-weighted volume and 3 orthogonal diffusion directions in 2:30 min. The ss-EPI image was acquired with a 12-channel head coil and the blipped-CAIPI rs-EPI image was acquired with a 32-channel coil.

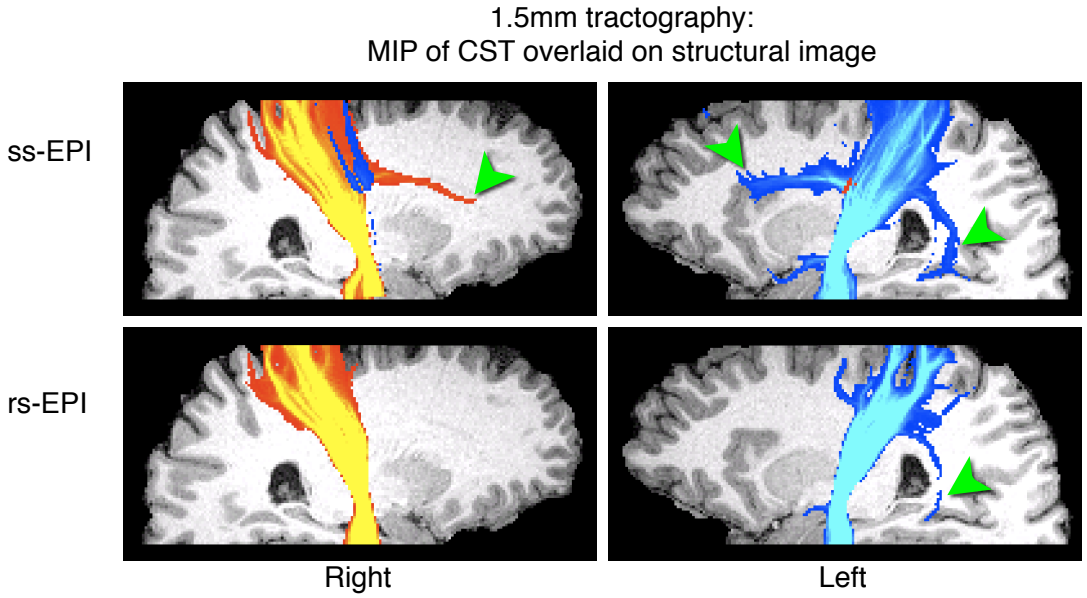

Figure 5.12: 1.5 mm tractography comparison in the CST. Green arrows indicate mis-tracking into anterior and posterior regions. Streamlines seeded in the left and right hemispheres are shown in blue-light blue and red-yellow, respectively.

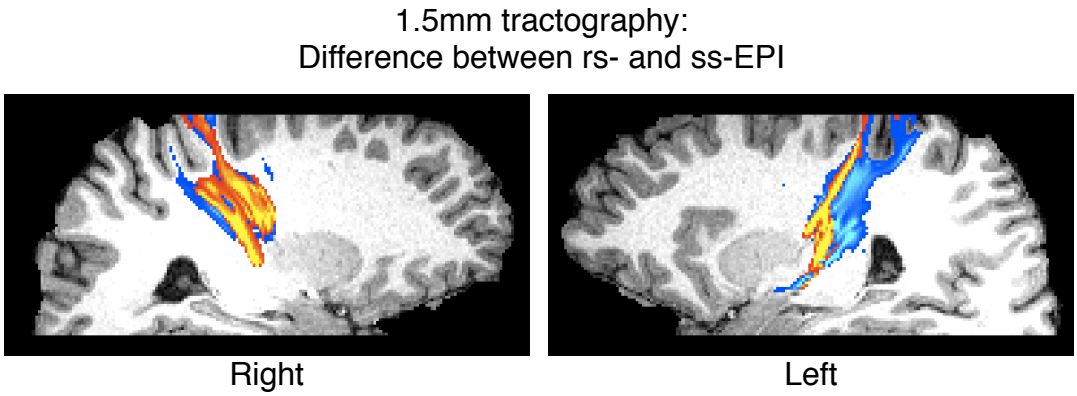

Figure 5.13: Difference between rs- and ss-EPI tractography in the CST. Regions where there are more rs-EPI streamlines are shown in red-yellow. Regions where there are more ss-EPI streamlines are shown in blue-light blue. The tractography differences are from representative single slices.

5.5 Discussion

This study demonstrates that a blipped-CAIPI multiband modification is compatible with the rs-EPI sequence and in particular, with multi-shot navigator corrections. In the rs-EPI approach, minor motion corruption (outside of systole) is removed by the non-linear navigator correction, while segments with more significant corruption (during systole) are simply reacquired. The navigator-based reacquisition was able to identify the shots with the most motion corruption based on the sum of the navigator signals from the simultaneously acquired slices. Lower slices in the brain suffer the worst motion artifact [95, 96, 97, 108], so as the number of simultaneously acquired slices is increased it may be necessary to modify the on-line detection of motion corruption to sensitively detect dropout in any of the slices. In the R_{slice} = 2 experiments in this study DW images did not exhibit serious motion artefacts when the reacquisition was judged on the combined k-space signal from two slices. With higher acceleration factors, the signal from individual receive channels could be considered in a more detailed way based on their proximity to the slices. Another option would be to implement on-line slice-GRAPPA reconstruction code (which has been written by the MGH group for ss-EPI data) with further modifications to unalias the data on-line. The slice-GRAPPA kernels would be estimated immediately after the pre-scan so that the data could be unaliased during the scan and the distribution widths calculated for the k-space signal from each slice.

The blipped-CAIPI multiband modification proposed here is expected to address one of the major current shortcomings of rs-EPI. Acquisition of whole-brain high-resolution trace-weighted data in clinically relevant scan times is now possible, as are measurements with a large number of diffusion directions. The tractography resolution demonstrated in this study clearly demonstrates the feasibility of tractography with rs-EPI. Although ss-EPI streamlines had a greater tendency

to reach cortex, likely as a consequence of higher SNR, the rs-EPI streamlines were more faithful to the local anatomy, as seen in Figs. 5.12 and 5.13. The deviations are in the phase-encoding direction (anterior-posterior) so could be attributed to higher blurring in ss-EPI from the longer readout, simple PF reconstruction and averaging. Also, higher resolution in all cases improved the accuracy of tractography, and it is thus significant that rs-EPI does not suffer from the distortion-resolution tradeoff that limits ss-EPI. Smaller voxels reduce the number of different tract orientations per voxel resulting in a less complicated diffusion-weighted signal.

High resolution does still pose an SNR challenge, which motivates higher acceleration factors to increase SNR efficiency and/or higher field strength. Limits on SAR and peak B_1 are a problem for RF pulse design, especially at higher field, because SAR scales with B_0^2 and B_1^2. RF pulses could be designed with the SLR algorithm [168] and VERSE [169] so that sharper slice profiles could be selected within peak RF voltage and SAR limits. In the simple addition of pulses used in this study, the peak RF voltage and SAR were usually close to the scanner imposed limits (which are calculated on a subject-specific basis). Lengthening the pulses reduces SAR but increases sensitivity to off-resonance effects in the excited slices. The recently introduced PINS technique [170] enables pulse design where the SAR is independent of the number of excited slices. An RF pulse envelope is multiplied with a comb function (the result is a sampled pulse envelope) so that an infinite series of repeating slices is excited, within the range of the transmit coil. In axial scans the series of repeating slices will generate unwanted excitation below the head that is detected by the receive coil and therefore sagittal or coronal orientations are necessary. Another option to reduce the RF power of pulses that has been suggested is to sum the pulses but shift their peaks slightly in time relative to each other so that the peak B_1 of the composite pulse is reduced.

Another avenue of future work would be to implement recently published improvements to the slice-GRAPPA reconstruction [171]. Also, rs-EPI data could be acquired with blipped-CAIPI slice acceleration and a reduced number of readout segments, with POCS partial Fourier reconstruction of the missing data [135], to further reduce the scan time per diffusion direction.

5.6 Conclusions

In this study we have applied the blipped-CAIPI method to rs-EPI to realise substantial reductions in scan time with potential for further acceleration. Slice acceleration is compatible with the multi-shot diffusion approach and allows rs-EPI to be used for high-resolution clinical trace-weighted imaging and tractography. Further work on details of the sequence could improve acceleration factors and data quality and implementing on-line scanner reconstruction code would be more practical for clinical diagnosis and research studies.

Chapter 6

3D Fourier Readout-Segmented EPI

6.1 Introduction

Diffusion-weighted imaging has been an important tool in the diagnosis of stroke [84] and for diffusion tensor imaging (DTI) and white matter fibre tracking [172, 93]. Images have typically been acquired with 2D multi-slice sequences with a single-shot echo-planar imaging (ss-EPI) readout [35]. This has proved a robust method for efficient acquisition of snapshot diffusion-weighted (DW) whole-brain images with ~2 mm isotropic resolution, particularly for DTI where large numbers of diffusion-encoded volumes are required. However, in ss-EPI, higher resolution is tied to more off-resonance and T_2^* blurring artefacts. Parallel imaging can be used to mitigate these artefacts [100, 101, 102] by reducing the effective echo-spacing and the readout duration by factors of 2-3, however the artefacts are worse at higher field. Resolution improvements and artefact reduction can be achieved by segmenting the k-space acquisition and using navigator techniques [104, 98] to account for the phase inconsistency between shots caused by motion during the diffusion-encoding gradients. Such approaches have been demonstrated in several sequences [106, 107, 103, 108, 124], including readout-segmented EPI (rs-EPI)

[126, 6, 132], which de-couples resolution from distortion and blurring artefacts by segmenting k-space in the readout direction to allow a short echo-spacing and a short within-shot readout. In the method of Porter *et al.* [6] the navigator information is used to apply a 2D navigator correction [108] for mild non-linear phase corruption and also to identify shots with large (uncorrectable) motion-induced phase errors so that they can be reacquired [131, 129]. Scan times are lengthened compared to ss-EPI by the multi-shot nature of the acquisition and two acceleration strategies have been presented: 1) using partial Fourier techniques to reduce the number of repetition time (TR) periods [135] and 2) reducing TR with a simultaneous multi-slice acquisition [138].

Although in-plane spatial resolution challenges can be overcome, the resolution in the slice direction is is usually not much higher than around 1.5 mm due to the difficulties in generating sharp radio-frequency (RF) pulse profiles for thin slices. Also, for small voxels low signal-to-noise (SNR) can be problematic, especially when the signal is further attenuated due to diffusion weighting. For these reasons a 3D DW sequence offers a potential solution; 3D Fourier encoding within an excited volume could allow high-resolution, well-defined voxels in the slice-direction. 3D sequences can often have higher SNR, although whether this is achieved in practice depends crucially on the achievable timing of the 2D and 3D sequences being compared. Isotropic high-resolution 3D DW data would facilitate identification of small infarcts and lesions and comparison with anatomical 3D images. The data would also be valuable for tractography because fewer fibre tract orientations will be present in smaller voxels thereby simplifying the DW signal [173, 174].

However, several challenges face 3D acquisitions that must first be addressed. In a multi-shot DW acquisition with volume excitation the motion-induced phase variation in the through-slab direction is significant and must be accounted for.

An extension to 3D of the non-linear navigator correction in 2D rs-EPI would require 3D navigator data which is difficult to realise. We refer to the k-space data that is used to form the high-resolution image as "imaging" data and the k-space data that is used to correct the imaging data as "navigator" data. The resolution of the imaging and navigator data is limited by the long time required to cover even a small 3D volume, which reduces the efficiency of the sequence. Also, the imaging data has susceptibility artefacts in the slow phase-encoding direction and point-spread function (PSF) problems caused by amplitude and phase modulations in the combined k-space signal. Some previous sequences have avoided multi-shot correction by using methods such as driven equilibrium diffusion preparation [175], acquiring a restricted volume [176, 177, 178], imaging in areas where motion phase errors are small [176, 179], and by cardiac gating the acquisition [176, 180, 113]. 3D steady-state sequences have been explored with pseudo-self-navigation using a 3D navigator compiled from multiple TRs [109] and with an additional rigid-body correction at each TR [181], but to date those methods still exhibit residual motion corruption that would compromise diffusion measures, particularly for diffusion-encoding in the superior-inferior direction where brain pulsation is most pronounced and causes the most severe artefacts [108].

The optimum TR with respect to maximal SNR efficiency is 1-2s (see below), which isn't particularly compatible with the most straightforward implementations of 2D multi-slice (TR=6-10 s) or 3D (TR=100-500 ms) volume acquisitions of the whole brain. One possibility is to sequentially excite a series of thin 3D slabs (with k_z phase-encoding within each slab over multiple shots) designed in such a way as to achieve the optimum TR. One can think of this approach as a hybrid between 2D multi-slice and conventional 3D imaging, which achieves this intermediate, optimal TR. As we shall see below, this has the further benefit

that motion artefacts in thin slabs are well-approximated by a 2D navigator (see simulations and experiments below), enabling navigation at no more cost to efficiency than in standard 2D readout-segmented EPI. Indeed, DW 3D multi-slab acquisitions have been demonstrated previously in TSE and EPI sequences with 2D in-plane navigator corrections [110, 112]. Multi-slab acquisitions with low TR (~1 s) will have increased slab-joining artefacts but these can be mitigated by overlapping slabs and by k_z oversampling [182].

In this study, we consider alternatives for navigator-corrected 3D rs-EPI data. We used a realistic simulation environment to compare two possible acquisition schemes that acquire 2D or 3D trajectories within 3D k-space for each shot of high-resolution imaging data. Based on these results, we selected the 2D trajectory for further study due to its favourable point-spread function properties, that predict high-fidelity images. The implemented scheme uses 2D imaging and navigator data in a 3D Fourier-encoded multi-slab extension of the rs-EPI sequence. This method offers reduced susceptibility artefacts and T_2^* blurring compared to similar multi-slab methods which encode a full 2D k-space partition at each excitation [183], at the expense of a longer scan time per volume. Non-linear in-plane correction (zeroth order correction in z) is applied to each k_z-encoded readout segment using its navigator (acquired at the centre of 3D k-space). In addition, simulations are presented which show that ordering the k-space acquisition with respect to the cardiac cycle [108, 184] reduces further any motion-induced phase artefacts. Shots close to the centre of 3D k-space are preferentially acquired in diastole to obtain cleaner data and the effect on image quality is demonstrated in simulations and experiments.

6.2 Theory

6.2.1 Assessment of 3D rs-EPI Acquisition Schemes

Two possible k-space trajectories for extension of rs-EPI to 3D encoded acquisition are shown in Fig. 6.1a. Both trajectories build on rs-EPI by acquiring 2D planes within the full 3D k-space (where a given k_z plane is achieved through addition of an appropriate k_z phase-encoding gradient). One trajectory (Fig. 6.1, top row) is based on single 2D k_z planes for both imaging and navigator data. We refer to the separate parts of this trajectory as IM_{2D} and NAV_{2D}, respectively, and the combined trajectory as IM_{2D}+NAV_{2D}. The second trajectory (Fig. 6.1, bottom row) acquires several contiguous k_z planes for both the imaging and navigator data. This trajectory is based on the single-shot echo-volumar imaging (EVI) technique [56, 185], in which spatial encoding can be applied in all three directions during a single readout. We refer to the separate parts of this trajectory as IM_{3D} and NAV_{3D}, respectively, and the combined trajectory as IM_{3D}+NAV_{3D}. The IM_{2D}+NAV_{2D} method can only achieve a 2D navigator correction, while IM_{3D}+NAV_{3D} is aimed at "proper" 3D navigation. For both schemes, as in rs-EPI, the imaging trajectories (IM_{2D} and IM_{3D}) acquire different (2D or 3D) segments of k-space from one TR to the next (indicated by the arrows in Fig. 6.1) so that the high-resolution k-space matrix is filled, while the navigator remains at the centre of k-space for all shots.

The imaging and navigator acquisitions are independent so, for example, IM_{2D} data could be acquired with NAV_{3D} data etc. From a motion correction perspective, acquiring IM_{3D} and NAV_{3D} data at contiguous k-space locations (i.e., at adjacent k_z planes) would allow an extension to 3D of the non-linear navigator correction in 2D rs-EPI. Alternatively, IM_{2D} data could be acquired with NAV_{3D} data so that phase information in z could be used to correct phase offsets and

shifts in the 2D imaging data.

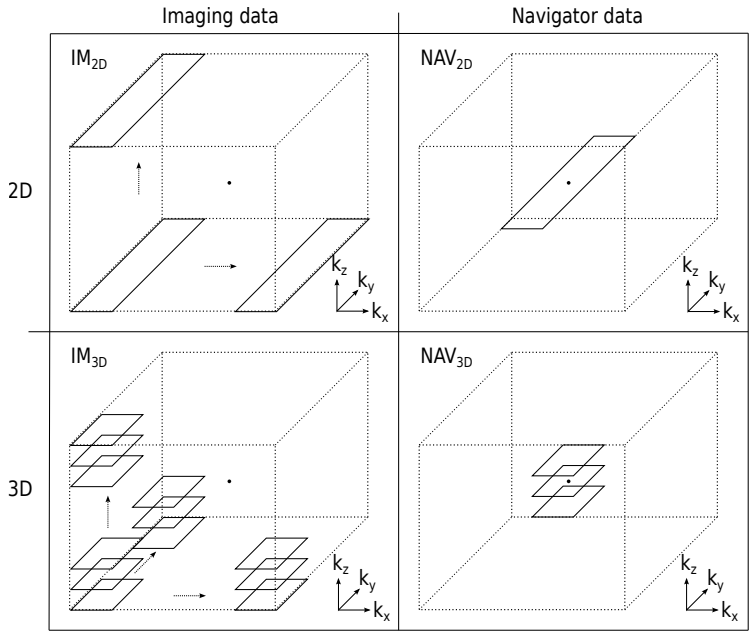

Figure 6.1: (Illustration of the possible 2D and 3D k-space acquisition schemes in a 3D version of rs-EPI. In each of the four panels, the dotted line represents the boundaries of the complete k-space matrix and the dot signifies the centre of 3D k-space. Solid lines represent acquisitions of one shot of k-space data. In the 2D versions (IM_{2D} and NAV_{2D}), these shots are simply k_z phase-encoded readout segments. In the 3D versions (IM_{3D} and NAV_{3D}), each shot is a stack of k_x, k_y planes, which is implemented in the pulse sequence by inserting G_z blip gradients to move to the next k_z plane and G_y rewind gradients to move to the k_y starting position. The k_z planes are at contiguous k_z points. For the imaging data (IM_{2D} and IM_{3D}), pre-phasing is used to move these shots around to sample the whole 3D k-space matrix. The navigator always acquires a full k_y resolution readout segment (NAV_{2D}) or stack of low k_y resolution readout segments (NAV_{3D}) at the centre of 3D k-space.

6.2.2 Motion-Induced Phase Simulations

In order to test acquisition and reconstruction strategies in the context of a 3D DW sequence, a simulation environment was implemented in Matlab (Mathworks, Natick, MA) for evaluation of motion-induced diffusion phase artefacts. Translations and rotations of the head are rigid-body movements and have been shown to impart constant and linear phase offsets, respectively [104, 98]. Deformation of

brain tissue during the cardiac cycle is a non-rigid motion that generates a non-linear spatial phase variation across the brain when it occurs during the diffusion-encoding gradients [108]. This motion has been measured with MR techniques [95, 96, 97] and manifests as if the brainstem is pulled downwards in a funnel-shaped motion during systole thus reshaping the medial and inferior parts of the brain.

Wirestam *et al.* [186] described a model for the temporal and spatial variation of the z component of brain tissue velocity, $v_z(x, y, z, t)$, based on experimental data from previous studies [95, 187] and used it to calculate signal loss caused by intravoxel dephasing when diffusion encoding is orientated in the z direction. In this study we use the model presented by Wirestam *et al.* to generate the expected phase maps for diffusion encoding along z so that multi-shot DW acquisitions can be simulated. This diffusion encoding direction has been shown to be the most problematic for diffusion phase artefacts [108, 109] so acquisition and reconstruction strategies can be tested in the worst-case scenario.

The equations for the temporal and spatial variation of velocity used in the simulations were taken from Ref. [186]. Temporal variation in velocity is given by Eq. 6.1:

$$v_z(0, 0, 0, t) = c_1 + c_2 e^{-c_3(t_0-t)^2} \tag{6.1}$$

where c_1 = -0.21 mm/s, c_2 = 2.4 mm/s and c_3 = 295 s^{-2} and the origin is in the centre of an axial slice at the level of the pons. The parameter t_0 is the time from the start of the pulse sequence simulation to the maximum brain tissue velocity. The spatial velocity profile is given by Eq. 6.2:

$$v_z(x, y, z, t) = v_z(0, 0, 0, t) e^{-a_1(x^2+y^2)-a_2 z^{4.31}} \tag{6.2}$$

where a_1 = 0.0022 mm^{-2} a_2 = 2.3$\times$10^{-8} mm$^{-4.31}$. Examples of the temporal and

spatial variation of velocity are plotted in Fig. 1. 6.2. Using this model for $v_z(\mathbf{r}, t)$,

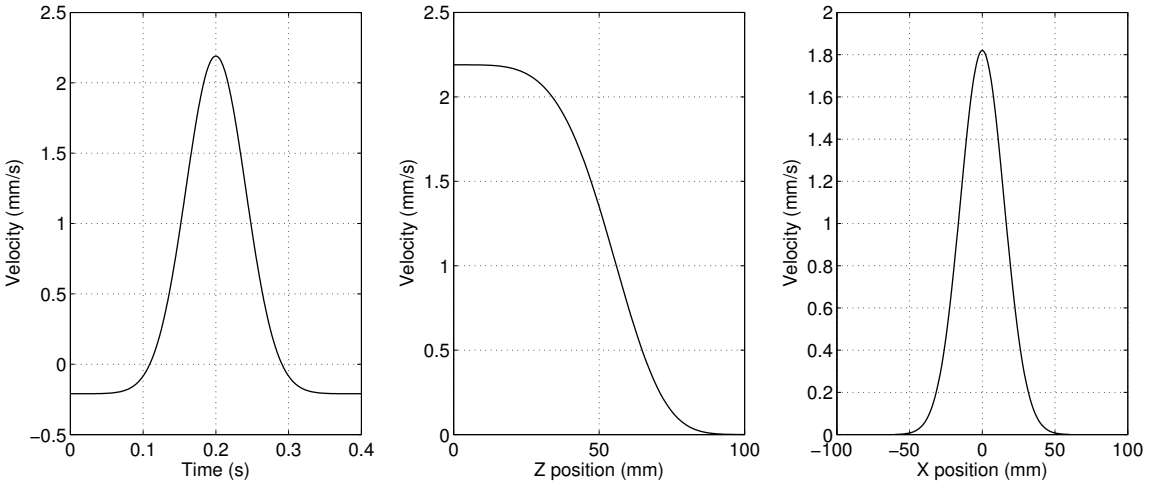

Figure 6.2: Plots of the component of brain tissue velocity in the z direction using the model described by Equations 6.1 and 6.2 from Ref. [186]: temporal velocity variation at x = y = z = 0 with t_0 = 0.2 s; velocity profile in z at x = y = 0 and t = t_0 = 0.2 s; velocity profile in x at y = 0, z = 40 mm and t = t_0 = 0.2 s.

the displacement in z since the start of the pulse sequence, $d_z(\mathbf{r}, t)$, could be calculated at each time step. Hence a map of the motion-induced phase imparted by the Stejskal-Tanner diffusion gradients, $\phi(\mathbf{r})$, was generated with Equations 6.3 and 6.4. The simulations used TE = 70 ms and the diffusion gradient timings were duration δ = 20 ms, diffusion time Δ = 37 ms. A diffusion gradient amplitude of 33.4 mT/m was used for a b-value of 1000 s/mm^2.

$$d_z(\mathbf{r}, t) = \int_0^t v_z(\mathbf{r}, t')\mathrm{d}t' \tag{6.3}$$

$$\phi(\mathbf{r}) = \gamma \int_0^{TE} d_z(\mathbf{r}, t)G(t)\mathrm{d}t \tag{6.4}$$

6.2.3 Cardiac-Reordering of k-space Shots

The main motivation for these motion-phase simulations was to investigate a reordering of the k-space acquisition with respect to the cardiac cycle. Previous

work has synchronised the acquisition to the cardiac and respiratory cycles to reduce ghosting artefacts [188, 189] in structural imaging and temporal fluctuations in 3D segmented-acquisition functional MRI (fMRI) [190, 184]. Cardiac synchronisation has also been used in segmented DW acquisitions to improve the accuracy of the "refocusing" navigator correction by enforcing a smoother variation of k-space phase and amplitude between interleaves in the combined data [108]. As discussed in Ref. [108], locally consistent phase-corruption allows a more accurate k-space deconvolution in the refocusing correction. Here, we apply the cardiac-reordering technique to a multi-shot DW acquisition with two compatible aims: 1) acquiring shots that are close to each other in k-space at similar points in the cardiac cycle (so that they have similar phase corruption) to ensure smooth phase and amplitude modulation, and 2) acquiring clean data (in diastole) close to the centre of k-space. The centre of k-space has the highest signal intensity and therefore it has the largest effect on the point-spread function in image-space.

Cardiac gating and previous reordering schemes affect the acquisition time and introduce a variable TR. Tijssen *et al.* were able to use a fixed TR by implementing an adaptive GRAPPA reconstruction to fill in parts of k-space which could not be acquired at the optimum time [184]. In the following simulations, the closest available shot to the corresponding cardiac phase is thus acquired. This ensures that data can always be acquired, even in systole, with a constant TR. Shots acquired during systole that have severe phase corruption could potentially be re-acquired. On-line re-acquisition was simulated, but was not implemented in this study because it would have required modification of the on-line reconstruction code that was beyond the scope of this work. However, this could be done as future work.

6.2.4 SNR Efficiency

SNR efficiency for a spin echo sequence run at 3 T is plotted as a function of TR in Fig. 3.8 for white matter tissue T_1= 1000 ms, T_2= 75 ms [136, 137] and echo time = 70 ms. From the plot in Fig. 3.8, a TR in the range 1-2 s will result in optimum SNR efficiency, which is compatible with a multi-slab sequence with about 10 slab excitations per TR.

6.3 Methods

6.3.1 Simulations

Point-Spread Functions of Acquisition Schemes

We investigated the effects of time-varying signal phase and amplitude on the trajectory PSF using simulations based around a 0.6×0.6×0.6 mm structural image and a 1×1×2.5 mm B_0 field map, both originally acquired in a representative healthy human volunteer. The field map was upsampled (interpolated by zero padding in the Fourier domain) and the structural scan downsampled so that each voxel of the structural image contained 5×5×5 field map values in order to simulate dephasing of signal due to field gradients. The simulation was performed by starting with a 120×120×10 k-space matrix of the downsampled structural and then selecting regions of this matrix to readout in different shots according to the IM_{2D} or IM_{3D} acquisition scheme.

During each readout, inhomogeneity-induced phase was applied in image space to every k-space point in the trajectory. Each k-space point (from the downsampled structural image) was placed in its corresponding location in a 120×120×10 matrix of zeros and Fourier transformed to image space where it was given the inhomogeneity phase (from the upsampled field map) correspond-

ing to its position in the readout. The timestep resolution between simulated k-space points was 2 μs and a linear accrual of phase was used during the readout such that the phase was zero at the centre of the readout (at the spin-echo centre), as shown in Fig. 6.3. Using the field map, the off-resonance phase could be calculated at each time (or k-space) point in the readout. Multiple readouts of sub-sets of k-space were combined to form the complete k-space. In the right hand column of Fig. 6.3 the phase and simulated amplitude variation during a spin-echo readout has been mapped onto one shot of the $\mathrm{IM_{3D}}$ scheme.

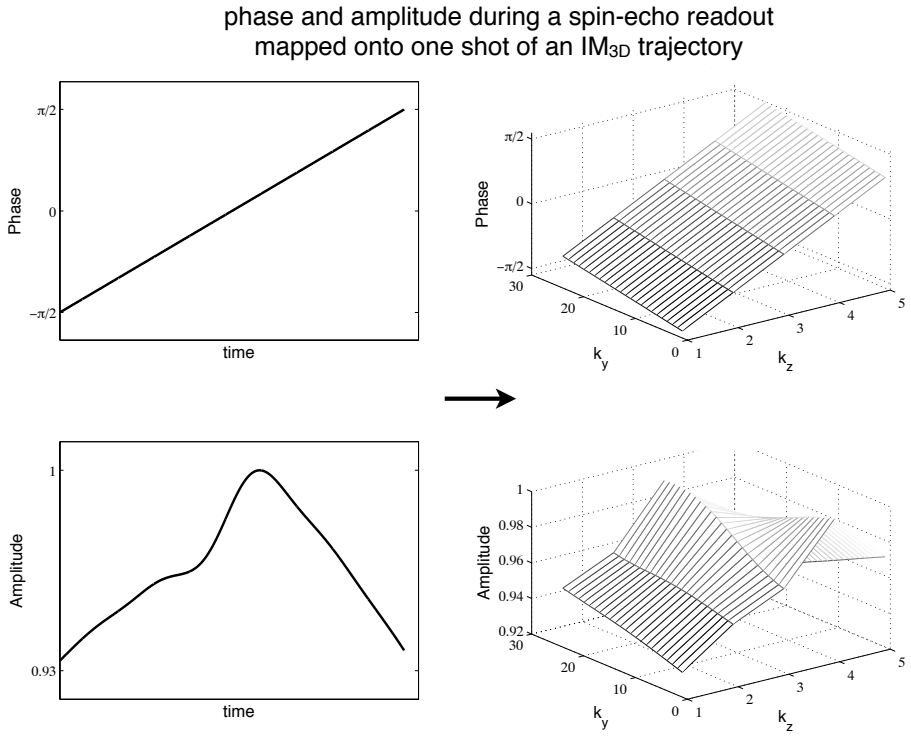

Figure 6.3: Variation of phase and amplitude during one readout (left-hand column) and the resulting phase and amplitude variation in k_y and k_z during one shot of a $\mathrm{IM_{3D}}$ trajectory (right-hand column). In the k_y, k_z plots in the right-hand column, each k_y point represents a k_y phase-encode line of k_x points. The amplitude of the image space mean signal was simulated at the central k-space point in a 120×120×10 matrix using an *in vivo* field map to generate dephasing. The deviation of the shape of the amplitude curve from and ideal spin-echo is due to the dephasing from field inhomogeneity on top of the spin-echo formation. An average value of inhomogeneity was chosen to calculate a maximum phase error of $\pi/2$ in the readout of 24×24×5 points ($k_x \times k_y \times k_z$) using a dwell time of 2 μs.

The phase and the time-evolving image space mean signal (i.e., the k-space signal without imaging gradients during the spin echo) were used to generate estimates of the spatial PSF. The complex signal was mapped onto the corresponding k-space locations for each readout and the multiple readouts were combined before the PSF was calculated by Fourier transformation.

Modified versions of the IM_{3D} scheme were also explored. Rather than sampling contiguous k_z points, the stacks of k_x, k_y planes could be interleaved over the complete k_z range. Also, the k_z modulation could be improved by shifting in time different interleaves with echo-time shifting (ETS) [191, 192].

Motion-Induced Phase

Phase maps at different readout shots were generated for random points in the cardiac cycle by varying the parameter t_0 in Eq. 6.1 between -100 and 100 ms. Simulations were performed using a numerical modified Shepp-Logan phantom and at each shot Rician noise was added to the magnitude image [193] (images are of the phantom are shown in Fig. 6.4 for reference). The phase maps were generated with the model for velocity described above and were saved so that different acquisition and reconstruction strategies could be compared.

A multi-shot DW acquisition was simulated by applying the motion-induced phase error at each shot in image space and then selecting the part of k-space sampled in the readout shot to add to the final multi-shot k-space. The imaging data constituted a readout segment (i.e., limited k_x extent, full k_y extent) from a single k_z partition (trajectory IM_{2D}) and the navigator could either be 1) a 2D fully sampled k_y readout segment at the centre of 3D k-space (trajectory NAV_{2D}) or 2) a 3D stack of k_z planes around the centre of k_z with limited extent in k_x and k_y (trajectory NAV_{3D}). Using the terms defined in Fig. 6.1 and above, the two simulated acquisitions were IM_{2D}+NAV_{2D} and IM_{2D}+NAV_{3D}. The navigator

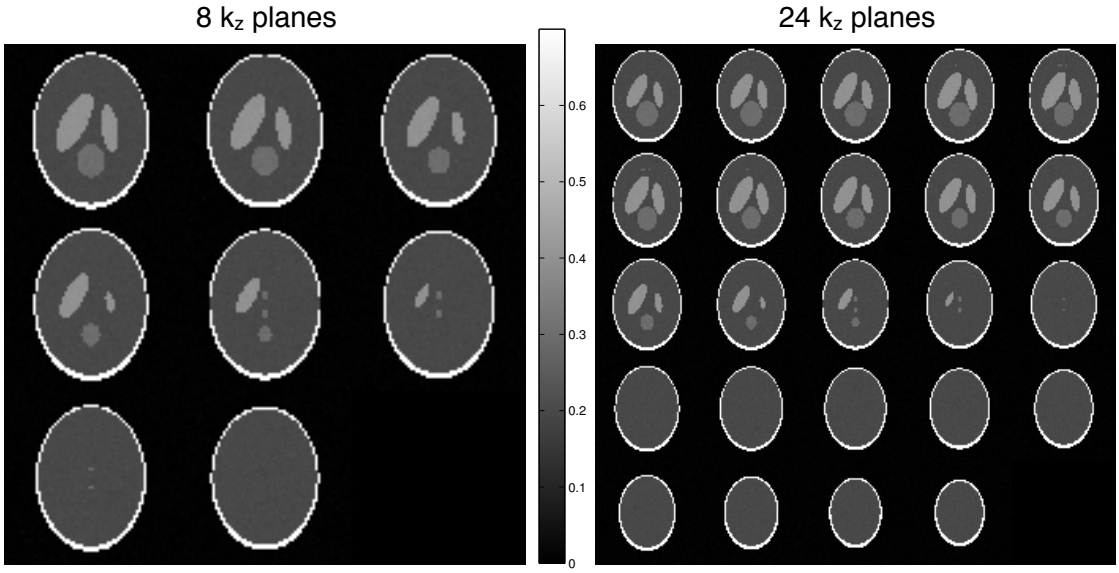

Figure 6.4: Images of the modified Shepp-Logan phantom (with added Rician noise) used in motion-induced phase simulations.

correction was performed in 2D or 3D as appropriate i.e. 2D refocusing correction for the $\mathrm{NAV_{2D}}$ data and 3D refocusing for $\mathrm{NAV_{3D}}$ data. Re-acquisition of motion corrupted segments [131], which is used in the 2D sequence, was simulated with 20% of scan time dedicated to reacquiring the k-space shots with the worst motion corruption. The reacquisition was weighted towards shots at the centre of 3D k-space.

The reconstructions compared in simulations were as follows: no correction, 2D navigation with $\mathrm{NAV_{2D}}$ data and 3D navigation with $\mathrm{NAV_{3D}}$ data. Cardiac-reordering of shots and re-acquisition was simulated with all reconstructions. The voxel size in the simulations was 2×2×2 mm so 8 k_z phase-encodes corresponds to a slab thickness of 16 mm. Artefact performance of the reconstructions was assessed via the quality of the resultant images or via temporal coefficient of variation (CoV) maps, which demonstrate how variable a set of 6 repeated simulations are. The coefficient of variation was computed by dividing the temporal standard deviation by the temporal mean. Quantitative summaries of artefact performance

of each acquisition/reconstruction combination were generated by calculating the mean CoV within the phantom.

6.3.2 In vivo

Based on the simulation results (see below), we propose using a IM_{2D}+NAV_{2D} version of 3D rs-EPI with 2D navigator correction and cardiac-reordering of shots, which we demonstrate by acquiring data with b = 1000 s/mm^2 and with a slab thickness ~16 mm. The 2D navigator correction of 2D planes of k_x, k_y imaging data at 3D phase-encoded k_z locations (IM_{2D}) was compared against theoretical predictions with *in vivo* experiments. The 2D navigator is always the central readout segment at the central k_z partition, i.e. the NAV_{2D} scheme. Additionally, the cardiac-reordering of shots was compared to a standard (sequential) k-space acquisition scheme.

Implementation

A standard 2D rs-EPI sequence [6] was modified to include slice gradient pre- and re-phasing either side of the imaging echo so that readout segments could be k_z phase-encoded (IM_{2D} scheme) and the navigator sampled the central k_x, k_z readout segment (NAV_{2D} scheme). Using the abbreviations introduced above, this sequence was a IM_{2D}+NAV_{2D} acquisition. These changes are shown in Fig 6.5.

Cardiac-reordering of k-space shots was implemented in the sequence so that the decision on which shot to acquire at each excitation was made at sequence run time. The shot choice was determined by the position in the cardiac cycle, which was estimated using pulse oximeter information. Unlike cardiac gating, this reordering scheme does not alter the sequence timing. The cardiac phase, ϕ_c, was calculated by dividing the time since the cardiac trigger, t_{trig}, measured with a pulse oximeter placed on the index finger (a 200 ms delay time between

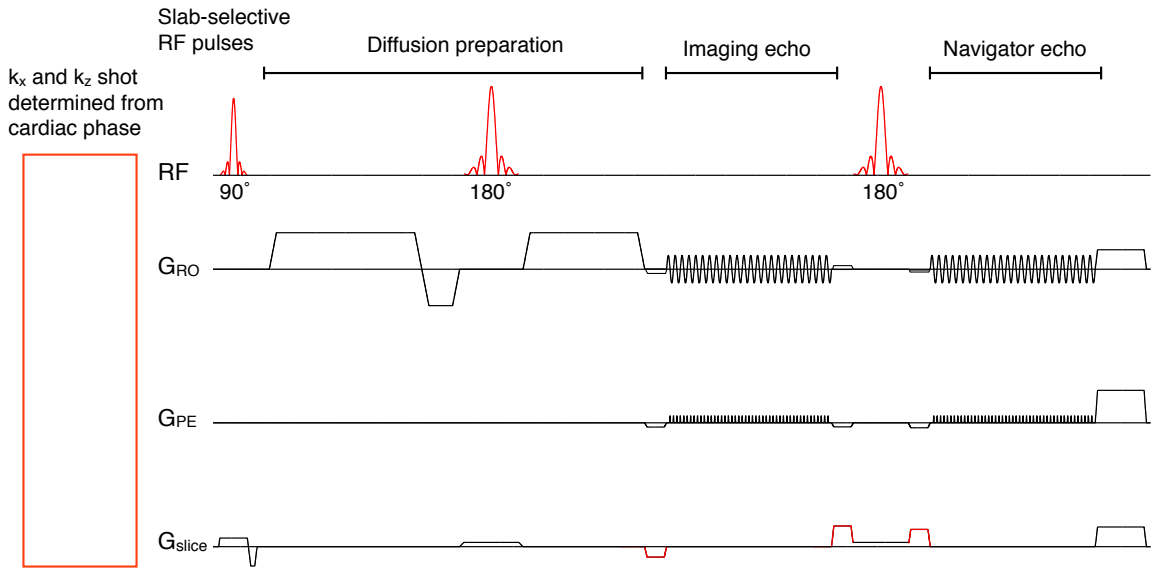

Figure 6.5: 3D rs-EPI pulse sequence with modifications to the original rs-EPI sequence coloured red. These are a module to determine which k-space shot to acquire based on the position in the cardiac cycle, slab-selective pulses and k_z phase-encoding gradients.

the finger and the head was determined empirically), by the mean cardiac cycle duration, T_{RR} (from the previous 10 cardiac cycles), as shown in Eq. B.1.

$$\phi_c = t_{trig}/T_{RR} \tag{6.5}$$

A cardiac phase of 0.4 - 0.6 was assumed to correspond to diastole and therefore to be the optimum time to acquire data. An algorithm (the Matlab version is provided in Appendix B) was used to determine the k_z and k_x shot indices (nz and nx, respectively).

When ϕ_c = 0.5, the algorithm chooses the central k_x, k_z readout segment. For example, with 8 k_z phase-encodes (Nz = 8) and 5 readout segments (Nx = 5) the optimum shot is nz = 4 (or 5), nx = 3. As ϕ_c tends towards 0 or 1 the chosen shot moves toward the edges of 3D k-space. As well as acquiring central k-space during the quiet portion of the cardiac cycle, this scheme also ensures that adjacent segments will have similar values of cardiac phase and therefore that the residual

phase-corruption after navigator correction will change smoothly through k-space [108]. If the chosen shot had already been acquired, the closest unacquired shot was selected. With this method, as the k-space matrix fills up, the selected shots usually become further away from the optimum position. Therefore, a measured cardiac phase between 0.4 and 0.6 is altered to 0.5 with the intention that shots near the centre of k-space are acquired early so that the outer k-space shots are acquired later, when they are more likely to be acquired with a sub-optimal cardiac phase. Logs of cardiac phase and selected shots were created so that this hypothesis could be verified. In a multi-slab acquisition, lists of shots were created for each slab and the shots were selected independently in each slab with the method described above. The objective of the scheme is to acquire cleaner data at the centre of k-space where corrupt data is more likely to result in severe signal loss artefacts.

Although reacquisition was simulated, it was not implemented in the sequence because it required modification of the on-line reconstruction code to accept 3D phase-encoded data rather than 2D multi-slice data, which was beyond the scope of this work. 3D reconstruction code will be the subject of future work so that motion-corrupted shots can be re-acquired to further improve image quality, as predicted by the simulations (see below).

To achieve contiguous slices at the interface between neighbouring slabs in a multi-slab acquisition a combination of slab overlap, k_z oversampling, and separate slab concatenations was used [182, 113]. To reduce saturation between slabs, all readout segments and k_z phase-encodes were acquired in separate concatenations for odd and even slabs. This increased the acquisition time by a factor of two but was necessary to improve artefacts at interfaces between slabs when scanning with TR ~1 s. The approaches used for slab overlap, k_z oversampling, concatenated scanning and a simultaneous multi-slab acquisition are illustrated

in Fig. 6.6.

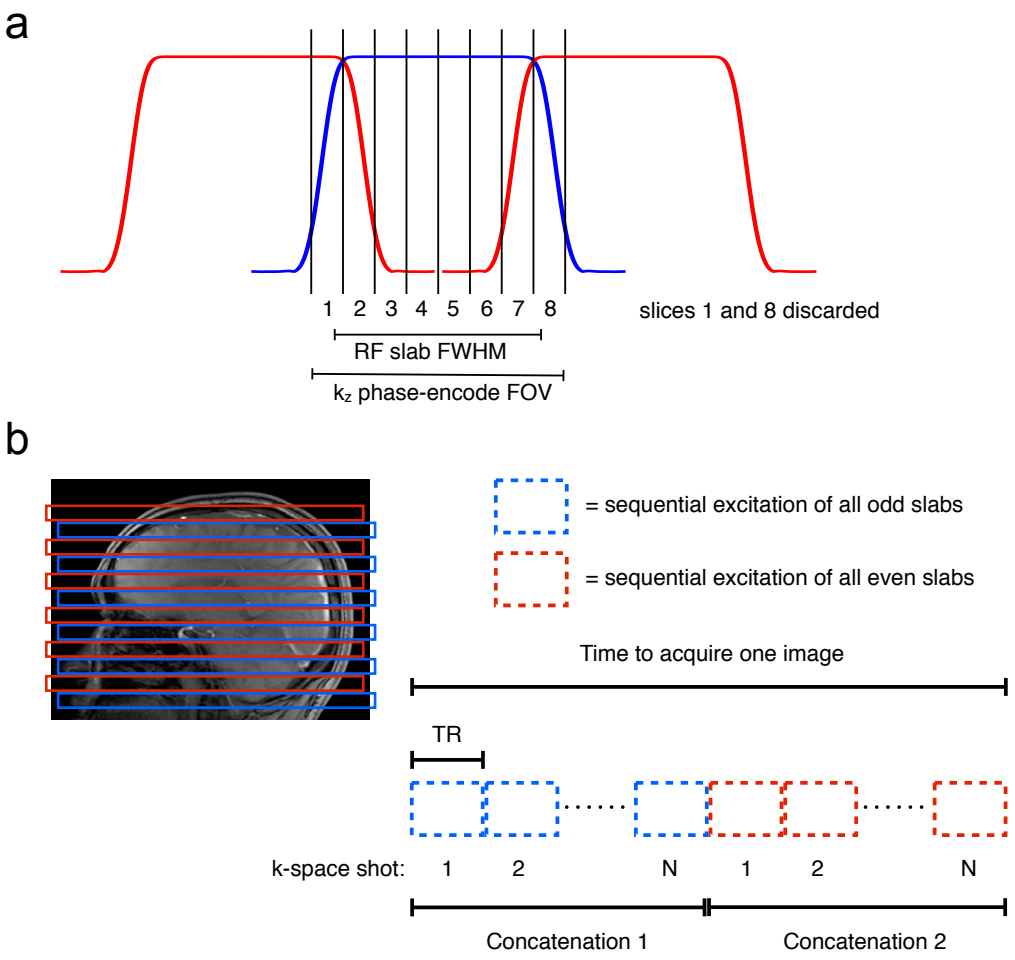

Figure 6.6: (**a**) Illustration with 8 k_z phase-encodes how adjacent slabs were overlapped and slices at either end of the slab were discarded. Also, the k_z phase-encode FOV was oversampled to be larger than the slab profile defined by the full width at half maximum (FWHM) of the RF pulse. (**b**) Illustration of concatenated scanning. The k-space data for odd and even slabs are acquired separately to reduce crosstalk between slabs. The number of k-space shots N = (readout segments) $\times$ (k_z phase-encodes).

Images were reconstructed off-line in Matlab (Mathworks, Natick, MA) from raw scanner data. The rs-EPI reconstruction used was similar to that described in previous work [135] but with modifications for the 3D phase-encoded data. The full pipeline was: phase-correction on a slab-specific basis using a pre-scan at k_y = k_z = 0 and regridding; GRAPPA reconstruction of missing k_y lines [66]; 2D navigator correction of readout segments at all k_x, k_z locations with a navigator

readout segment at the central k_x, k_z location; compilation of a 3D k-space matrix after discarding the extra k_x columns at the edges of each readout segment; 3D Fourier transform and combination of coil data with sum of squares reconstruction. GRAPPA under-sampling of k_z phase-encodes was not possible due to the small z FOV of each slab, which would lead to large g-factor-related losses when trying to unalias voxels a maximum of 8 mm apart (16 mm slab with factor 2 under-sampling).

Data Acquisition and Processing

Data were acquired in three healthy volunteers with a MAGNETOM Verio 3 T scanner (Siemens Healthcare, Erlangen) under an approved technical development ethics protocol. A pulse oximeter was used to send triggers to the scanner at each R-wave so that the cardiac phase could be determined. The diffusion preparation module was a modified Stejskal-Tanner pair [117] to allow a shorter TE. Eddy-current correction was applied using FMRIB's Diffusion Toolbox [150], which consists of an affine registration (12 degrees of freedom) of all volumes to the initial b = 0 s/mm^2 volume. Non-brain tissue was removed from the images using the FSL Brain Extraction Tool [151].

Variability in Repeated Diffusion Acquisitions

Data were acquired with 12 and 18 k_z phase-encodes, with and without cardiac-reordering of shots to verify the improvement in image quality predicted by simulation. Six repeated volumes with diffusion encoding in the z direction (the encoding direction that is most sensitive to motion artefacts) were acquired to assess the variability of the images. Voxel-wise maps of coefficient of variation (CoV) were computed by dividing the temporal standard deviation in signal by the temporal mean signal.

1.83×1.83×2 mm images with 12 or 18 k_z phase-encodes were acquired with TR/TE = 1000/87 ms, flip angle = 68°, FOV = 220×220 mm, matrix = 120×120, 5 in-plane readout-segments and a single slab. One T_2-weighted volume and six DW volumes with b = 1000 s/mm^2 encoding in the z direction were acquired with a 12-channel head coil. The scan times of the 12 and 18 k_z phase-encode acquisitions were 7:05 min and 10:35 min, respectively, and each acquisition was repeated with and without cardiac-reordering.

Multi-Slab Trace-Weighted Images

1.53 mm isotropic trace-weighted images were acquired with TR/TE = 1500/75 ms, flip angle = 77°, FOV = 220×220 mm, matrix = 144×144, 7 readout-segments, 12 slabs with 25% slab overlap, 8 k_z phase-encodes within each slab with 10% k_z oversampling and R = 2 in-plane GRAPPA k_y under-sampling. One T_2-weighted volume and three DW volumes with orthogonal b = 1000 s/mm^2 encoding were acquired in 11:30 min with a 32-channel head coil. Trace-weighted images were calculated by taking the geometric mean of the diffusion-weighted images.

Simultaneous Multi-Slab

1.53 mm isotropic DW images were acquired with TR/TE = 1500/75 ms, flip angle = 77°, FOV = 220×220 mm, matrix = 144×144, 7 readout-segments, 2 simultaneously excited slabs with 72 mm separation, 8 slab excitations (16 slabs in total), 25% slab overlap, and 8 k_z phase-encodes within each slab with 10% k_z oversampling. A 32-channel head coil was used to acquire a T_2-weighted volume and 3 DW volumes with orthogonal b = 1000 s/mm^2 encoding in 11:30 min plus 2:50 min for each of the two reference acquisitions. No in-plane shifting of slices was used in this preliminary demonstration. Multiplexed partition data was unaliased in k-space using the slice-GRAPPA parallel imaging reconstruction [74].

6.4 Results

6.4.1 Simulations

Point-Spread Functions of Acquisition Schemes

The phase and amplitude variation with time of the image space mean signal during the spin-echo are shown in the left-hand column of Fig. 6.3. In the plot of phase variation, the maximum phase error of $\pi/2$ at the extremes of the spin-echo readout is a mean value across the brain. In the right-hand column of Fig. 6.3, these phase and amplitude variations are mapped onto one shot of a IM_{3D} scheme (from Fig. 6.1) which samples 5 k_z planes in one shot. Note that each k_y point represents one phase encode line (along k_x). This demonstrates the amplitude and phase discontinuities in the k_z direction caused by the jump to a new k_z plane at the initial k_y position, in this case after the readout of 24 k_y lines.

When k-space is sampled in multiple shots with this IM_{3D} scheme, the amplitude variation in Fig. 6.3 (lower right panel) will be repeated through k-space. The k_x dimension was split into 5 readout segments and in the IM_{3D} scheme these segments were divided into 5 k_z planes so 5 shots were required to fill the y resolution. The acquisition time of the IM_{2D} and IM_{3D} schemes would then be the same. The results of a IM_{2D} scheme are shown in Fig. 6.7a and the results of a IM_{3D} scheme in Fig. 6.7b.

These results suggest that the PSF in the IM_{3D} scheme introduces aliasing artefacts (indicated by the off-centre peaks along z and y) as well as some blurring. Blurring is exhibited by a broadening of the main PSF peak from a pure delta function. These artefacts are most prominent along z, where signal from alternating slices will alias. There are no PSF problems in the simulations with the IM_{2D} scheme apart from relatively minor blurring along y.

The results of the modified IM_{3D} schemes are shown in Fig. 6.8 and the PSF

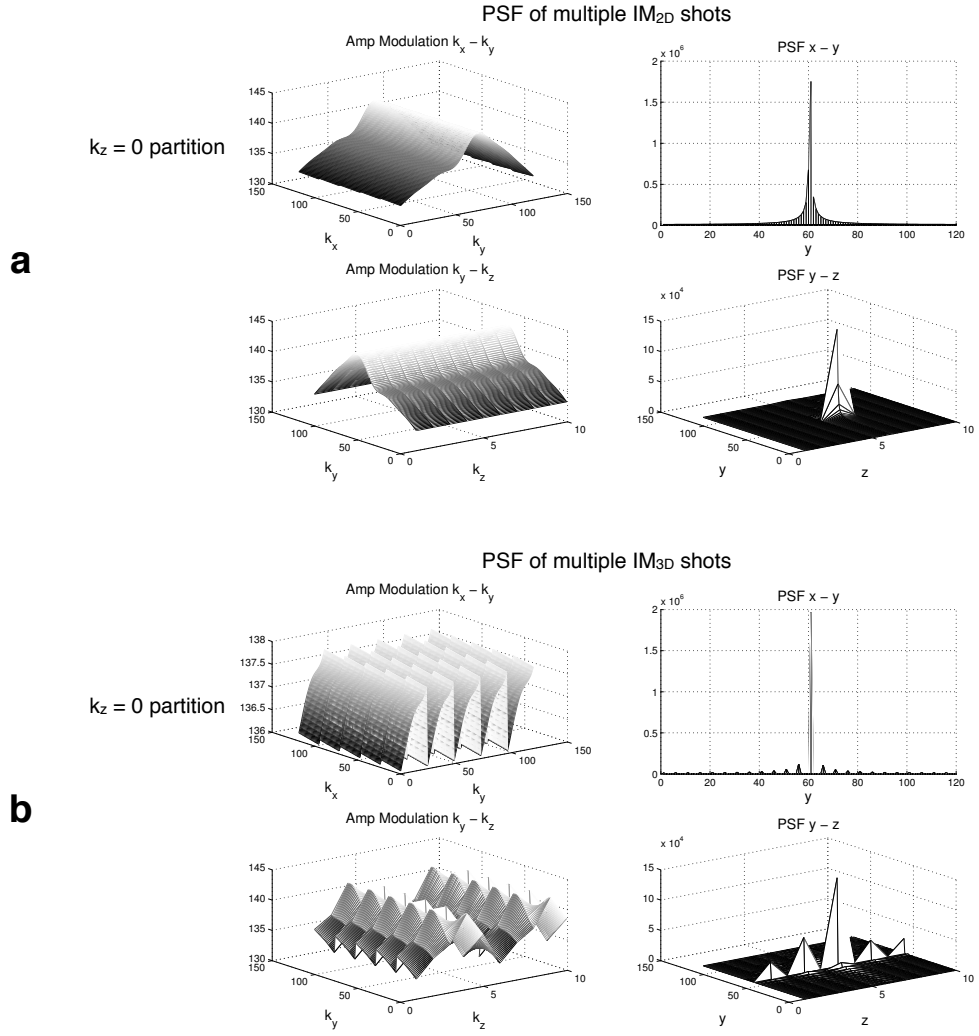

Figure 6.7: (**a**) $\mathrm{IM_{2D}}$ and (**b**) $\mathrm{IM_{3D}}$ scheme simulation results. The modulation in amplitude and the associated point spread functions are shown in the x - y and y - z planes. In the k_y, k_z plot, each k_y point represents a k_y phase-encode line of k_x points. The scales on the amplitude and PSF plots are arbitrary.

results are compared to the original IM_{3D} scheme (which samples contiguous k_z points). The PSF interleaved scheme suggests substantial blurring in the z direction with an interleaved scheme and a slight improvement with ETS. Another factor to consider with the interleaved schemes is that the trajectory would not be as amenable to 3D navigator correction as a contiguous scheme.

In light of these simulations, it was decided that the IM_{2D} trajectory was most promising and that this was the strategy to develop further. Motion phase artefacts were simulated with IM_{2D}+NAV_{2D} navigator correction and a IM_{2D}+NAV_{2D} scheme was implemented in the sequence. A navigator correction of IM_{2D}+NAV_{3D} data was simulated but not implemented, with further work required on this topic.

Motion-Induced Phase

A comparison at b = 1000 s/mm^2 of no correction, 2D navigation, 3D navigation (all without reacquisition) and 2D navigation with reacquisition, with and without cardiac-reordering of shots, is shown in Figures 6.9 and 6.10 for 8 and 24 k_z phase-encodes, respectively. Simulations of a 2D navigator reconstruction with reacquisition of corrupted shots is shown in Fig. 6.11 at b = 1000 s/mm^2 and b = 3000 s/mm^2 for 8, 16 and 24 k_z phase-encodes. Figure 6.12 summarises the artefact performance of the different reconstructions with 8 and 24 k_z phase-encodes (summary of the data shown in Figures 6.9 and 6.10). Figure 6.13 summarises the artefact performance of the 2D navigator reconstruction with reacquisition at b = 1000 s/mm^2 and b = 3000 s/mm^2 for 8, 16 and 24 k_z phase-encodes (summary of the data shown in Fig. 6.11).

The simulation results in Figures 6.9, 6.10 and 6.12 show that the 2D and 3D navigator corrections have similar performance, with slightly better image quality in the 3D correction. However, the 2D correction was chosen to be the

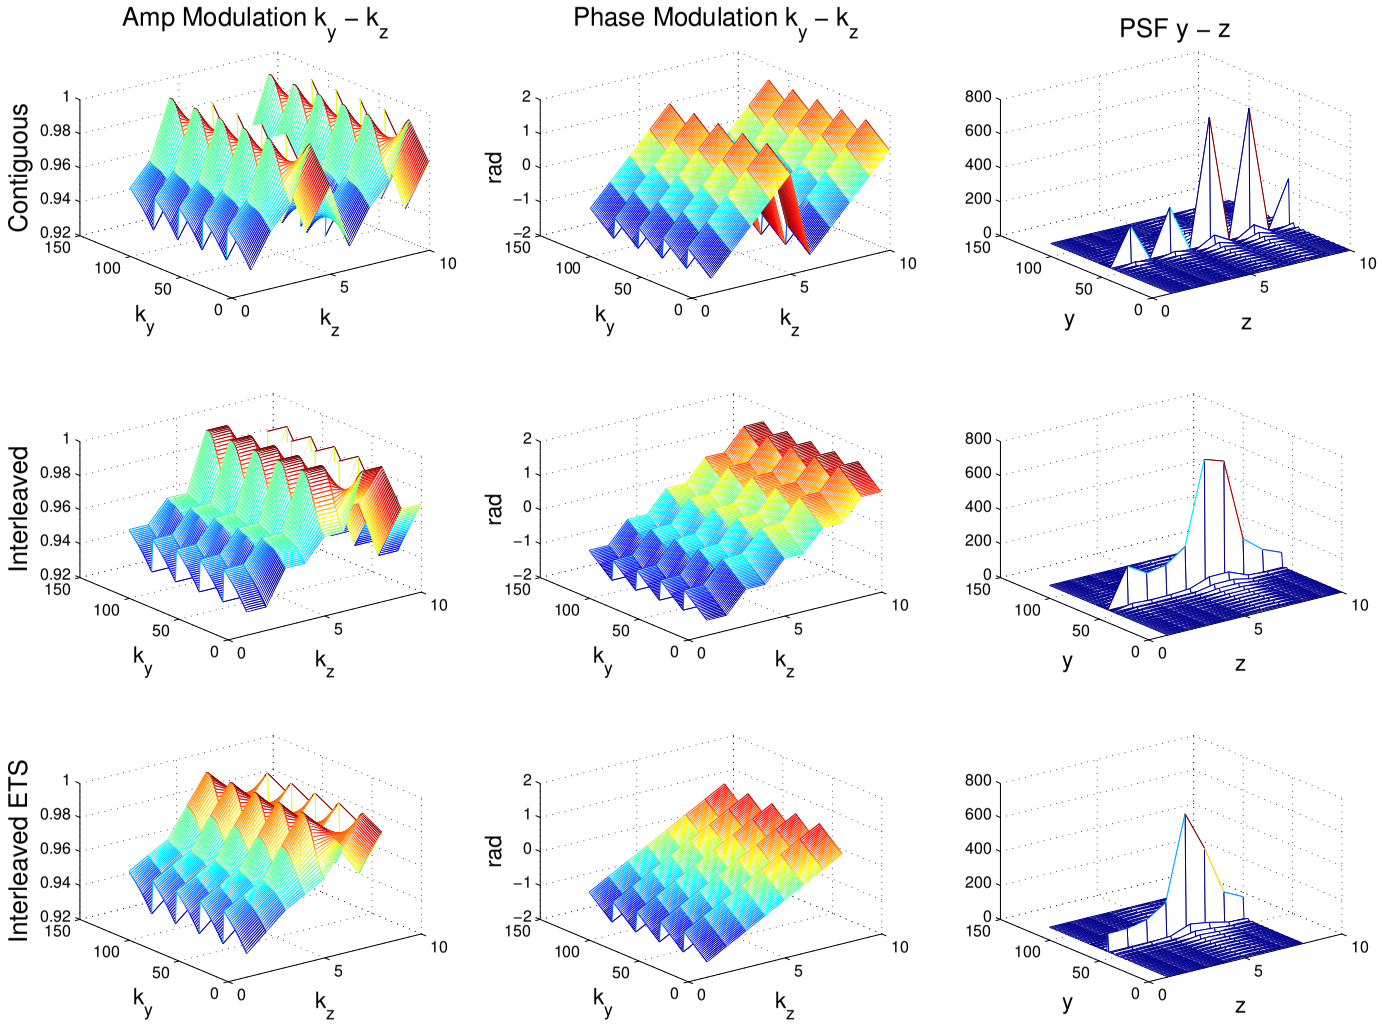

Figure 6.8: Comparison of simulation results of modifications to the IM_{3D} scheme. Amplitude and phase modulation with associated point spread functions (PSF) generated for contiguous (1st row), interleaved (2nd row) and interleaved with echo-time shift (3rd row) are compared. The calculations are based on 5 k_x readout segments, 5 k_y repeats and 2 k_z repeats required to fill a 120×120×10 k-space matrix. The maximum phase error in a readout of 24×24×5 was estimated to have an average value of $\pi/2$. In the k_y, k_z plot, each k_y point represents a k_y phase-encode line of k_x points.

best candidate for implementation (i.e., scheme IM_{2D}+NAV_{2D}) due to the close relationship of the sequence and the reconstruction to the the standard 2D rs-EPI sequence. The simulations also suggest that major improvements in image quality can be realised by reordering the k-space acquisition of shots so that clean data is acquired close to the centre of k-space and phase corruption is locally consistent. This does not lengthen the acquisition time and can be implemented using real-time calculations based on physiological monitoring of the cardiac phase at each shot in a similar fashion to a recent reordering scheme applied to fMRI [184]. The results show that image quality can be improved further by reacquiring the shots with the worst motion corruption [131] in an extension of the current 2D rs-EPI implementation [6]. The simulations suggest (data not shown) that a navigator consisting of either a readout segment at the central k_z partition (NAV_{2D}) or a readout segment split into a stack of k_z partitions with low k_y resolution (NAV_{3D}) at the centre of 3D k-space can be used to judge motion corruption.

Figures 6.11 and 6.13 shows the expected increase in artefacts with stronger diffusion encoding and with increasing slab thickness from 16 to 48 mm for the example case of the 2D navigator reconstruction with reacquisition. Stronger diffusion encoding creates a more complex phase variation across the brain with more phase wraps and in a larger slab there is more scope for phase difference in the slab direction. When there is more complicated phase variation in z the 2D or 3D navigator data becomes less effective at resolving and removing the phase corruption. Therefore the data contain residual phase corruption which causes signal loss and ghosting artefacts due to destructive phase interference as well as discontinuities and periodicity in the k-space signal [108]. The phase variation at each shot is unpredictable so in repeated acquisitions these artefacts will be different, leading to increasing variation in the temporal standard deviation maps.

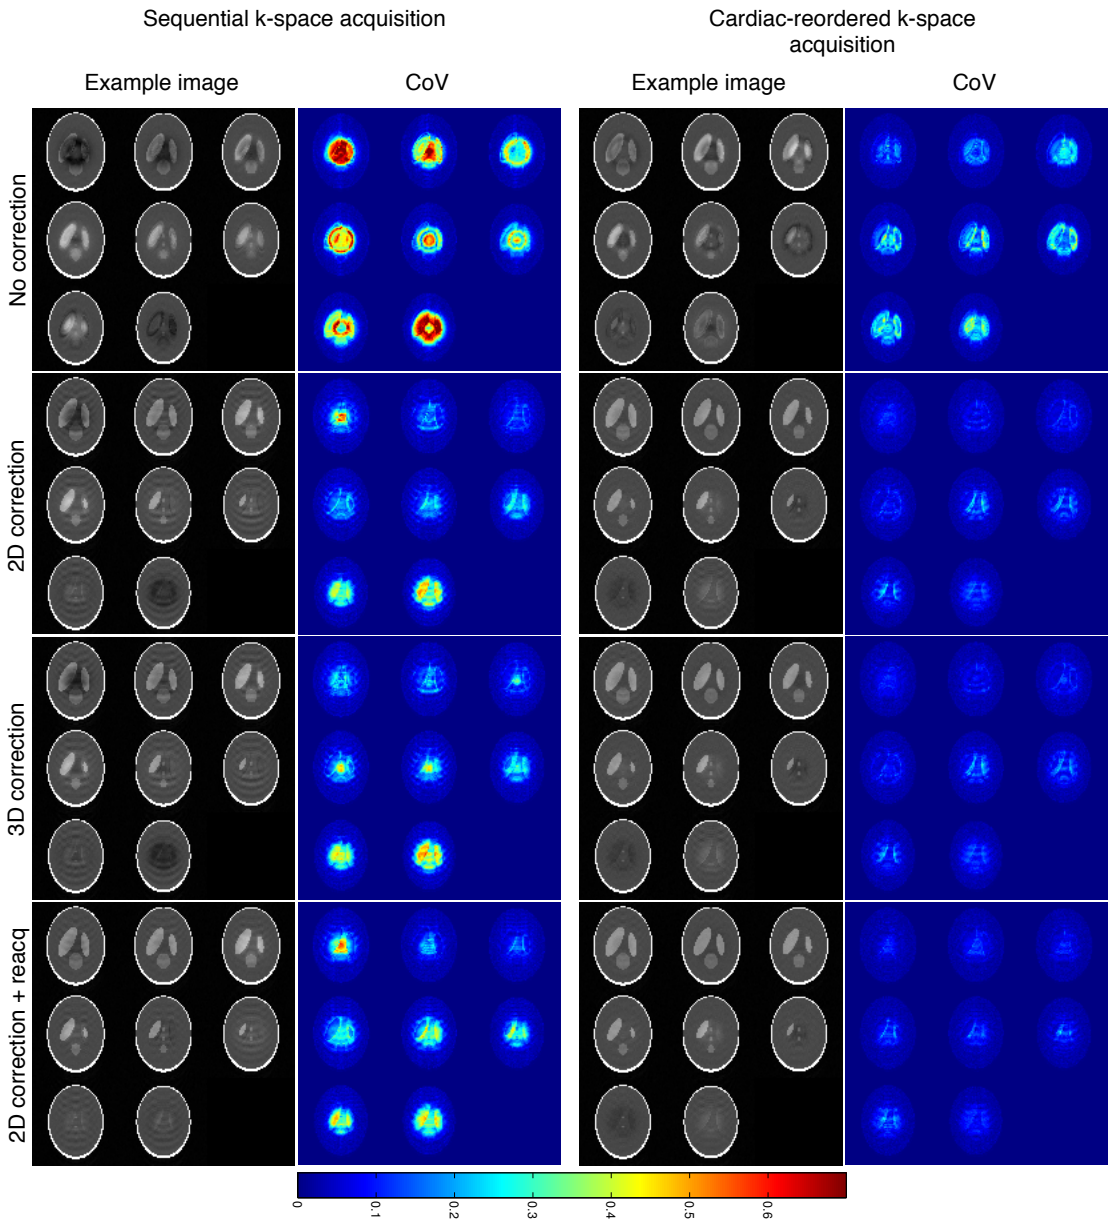

Figure 6.9: Simulation results showing a comparison of acquisition and reconstruction approaches for 8 k_z phase-encodes and 16 mm slab thickness with b = 1000 s/mm^2 diffusion weighting in z. The coefficient of variation (CoV) is over 6 repeated acquisition simulations.

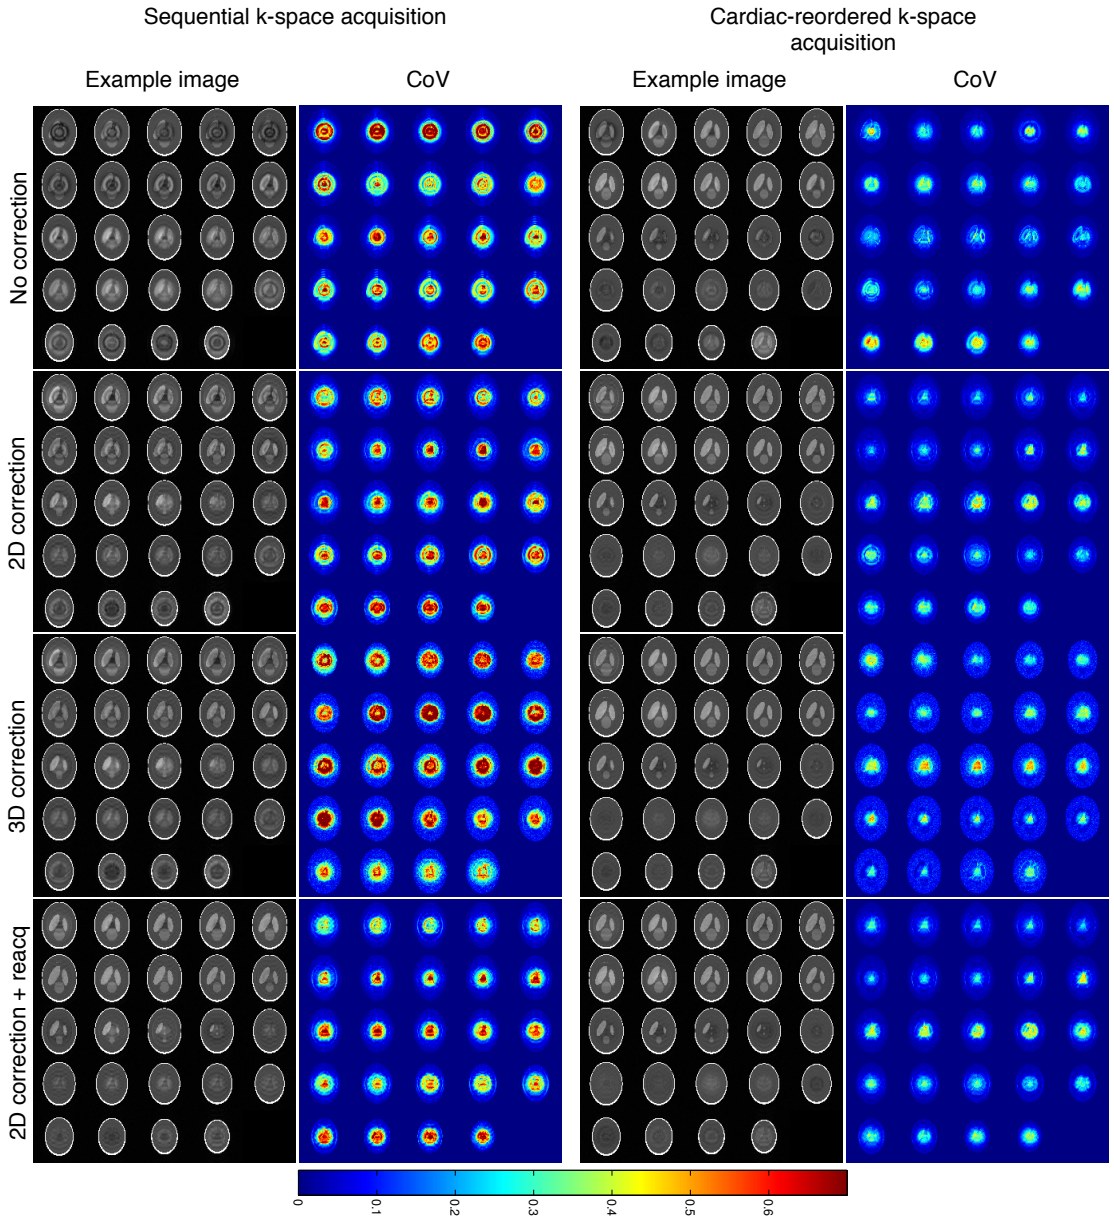

Figure 6.10: Simulation results showing a comparison of acquisition and reconstruction approaches for 24 k_z phase-encodes and 48 mm slab thickness with b = 1000 s/mm^2 diffusion weighting in z. The coefficient of variation (CoV) is over 6 repeated acquisition simulations.

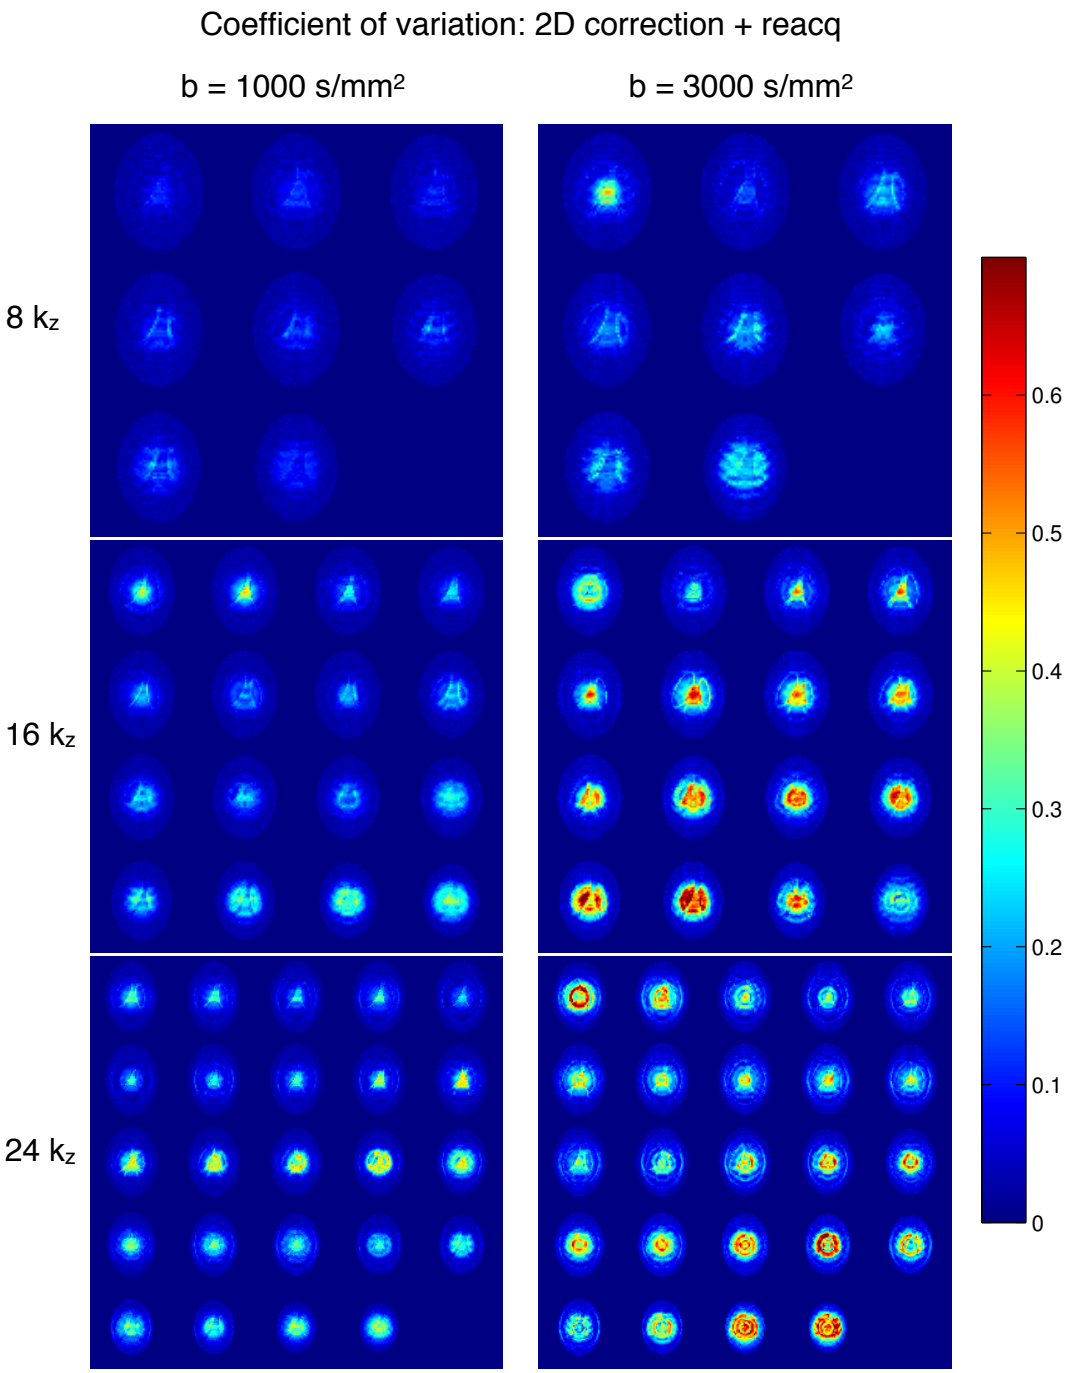

Figure 6.11: Simulation results showing a comparison of b = 1000 s/mm^2 and b = 3000 s/mm^2 diffusion weighting in z for 2D navigator correction with cardiac-reordering of shots and reacquisition of corrupted shots. The coefficient of variation (CoV) is over 6 repeated acquisition simulations.

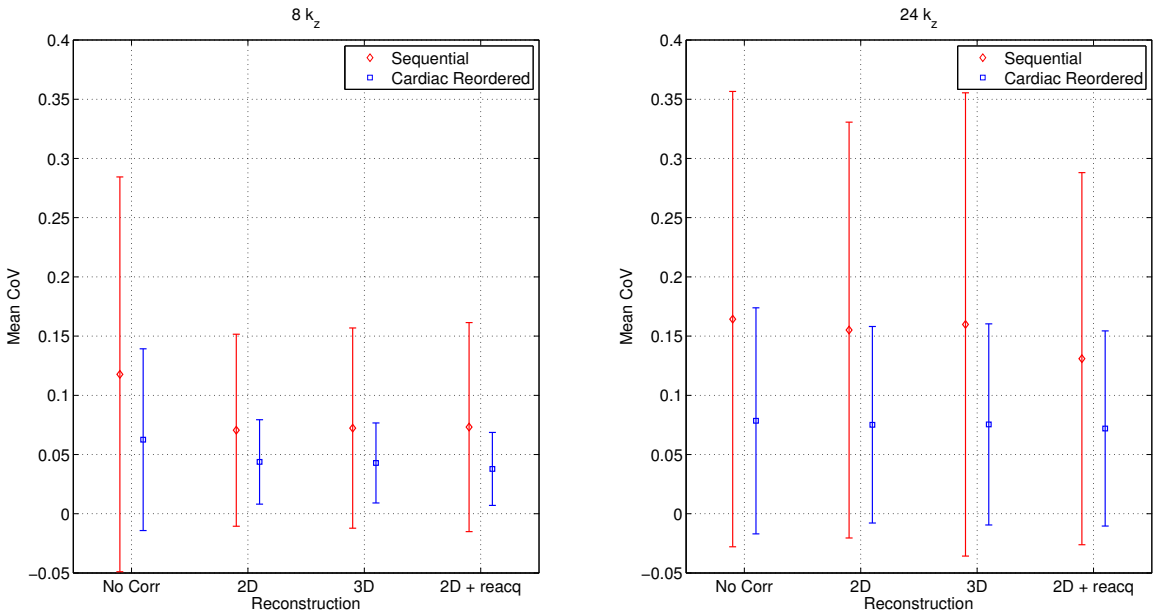

Figure 6.12: Simulation results showing a comparison of the mean coefficient of variation (CoV) for the different reconstruction approaches for 8 and 24 k_z phase-encodes with b = 1000 s/mm^2 diffusion weighting in z. The mean is within a ROI defined by all the voxels in the phantom.

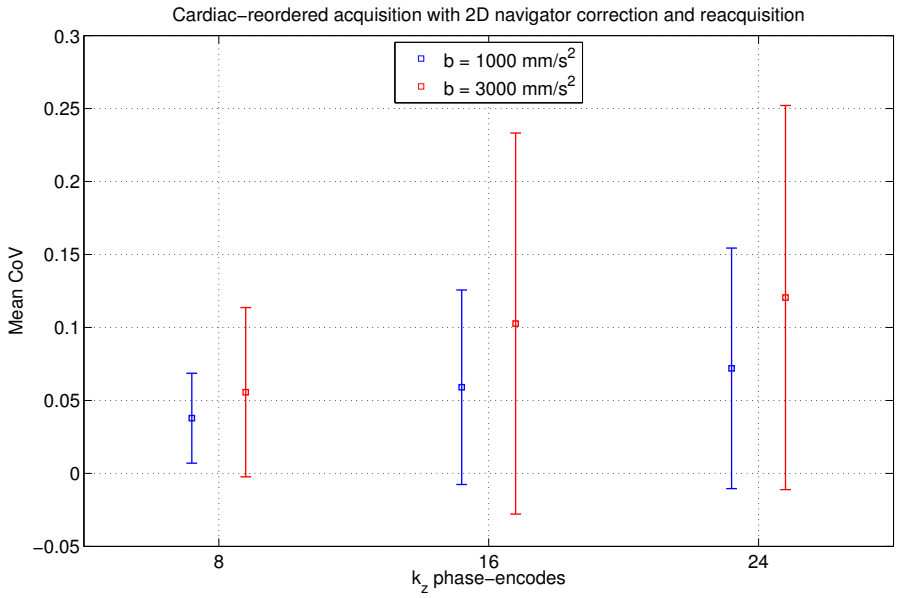

Figure 6.13: Simulation results showing a comparison of the mean coefficient of variation (CoV) between b = 1000 s/mm^2 and b = 3000 s/mm^2 diffusion weighting in z for 8, 16 and 24 k_z phase-encodes. The comparison is for a cardiac-reordered acquisition with 2D navigator correction and reacquistion of the shots with the worst motion corruption. The mean is within a ROI defined by all the voxels in the phantom.

6.4.2 In vivo

Images and variability results are shown in Fig. 6.14 for 12 2 mm slices and Fig. 6.15 for 18 slices. These results demonstrate the successful 2D navigation and the clear reduction of image artefacts with cardiac-reordering of k-space shots. Ignoring edge slices that would be discarded, the largest improvements (high to low CoV) are in the centre of the brain (where the non-rigid deformation is concentrated) and some surrounding white matter structures. In the cardiac-reordered data, the only remaining regions of high CoV are in low signal areas (e.g. CSF). In the 36 mm slab data (Fig. 6.15) the remaining CoV increases compared to the 24 mm slab data (Fig. 6.14). In ss-EPI data, artefacts can be caused by through-plane dephasing and signal loss due to large systolic motions with short timescales. In the 2D navigator corrected DW images, the effect of reordering is clearer in the 36 mm slab data which exhibit artefactual low signal in the central areas of the brain in the sequential k-space acquisition. This signal is restored in the cardiac-reordered data. The mean CoV within the phantom was calculated so that quantitative results could be compared. A summary of simulation and experiment results for b = 1000 s/mm^2 diffusion weighting in z with the 2D navigator correction (without reacquisition) is shown in Table 6.1. A consistent 40-50% improvement with cardiac-reordering was observed in the simulations and experiments, and the improvement increased with more k_z phase-encodes/thicker slabs.

Multi-slab T_2- and trace-weighted images are shown in Fig. 6.18 with evidence of some residual slab joining artefacts. The cardiac cycle durations observed in the volunteers were in the range 600 - 1200 ms. Whole brain diffusion- and trace-weighted images from the simultaneous multi-slab acquisition are shown in Fig. 6.19.

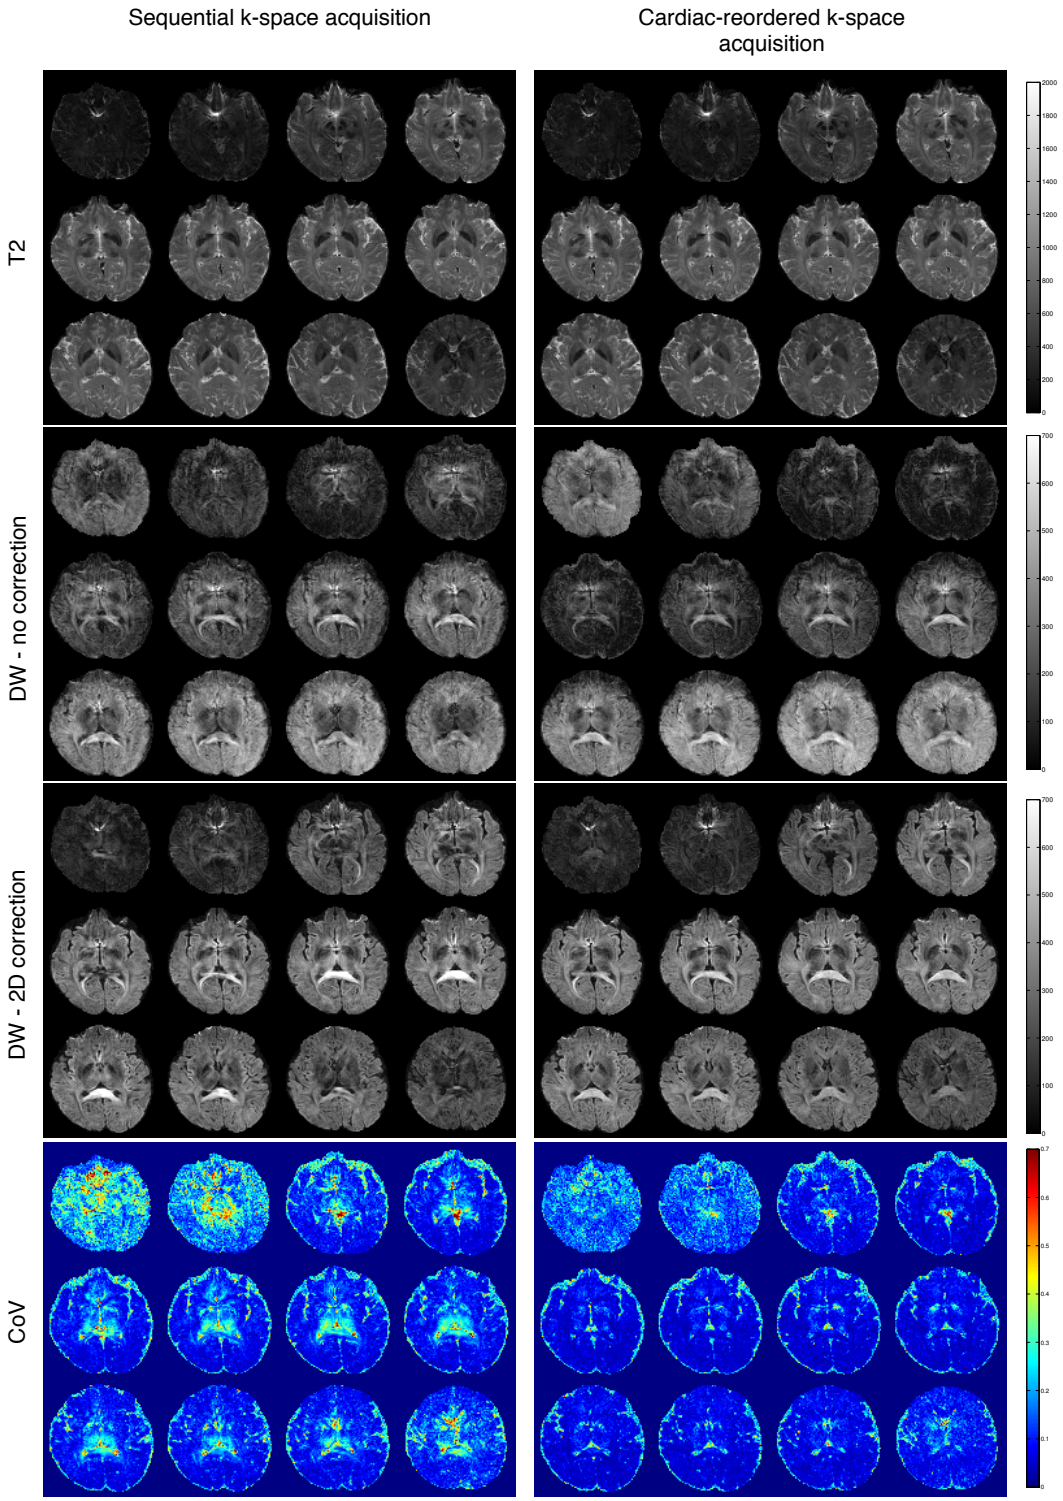

Figure 6.14: Experimental results showing a comparison of images and variability between sequential acquisition of k-space and cardiac-reordered acquisition of k-space in data with 12 k_z phase-encodes and 24 mm slab thickness. Diffusion weighting is in the z direction with b = 1000 s/mm^2. Resolution of the images is 1.8×1.8×2 mm. The coefficient of variation (CoV) is calculated from the temporal standard deviation divided pixel-by-pixel by the temporal mean over 6 repeated acquisitions in the data with 2D correction.

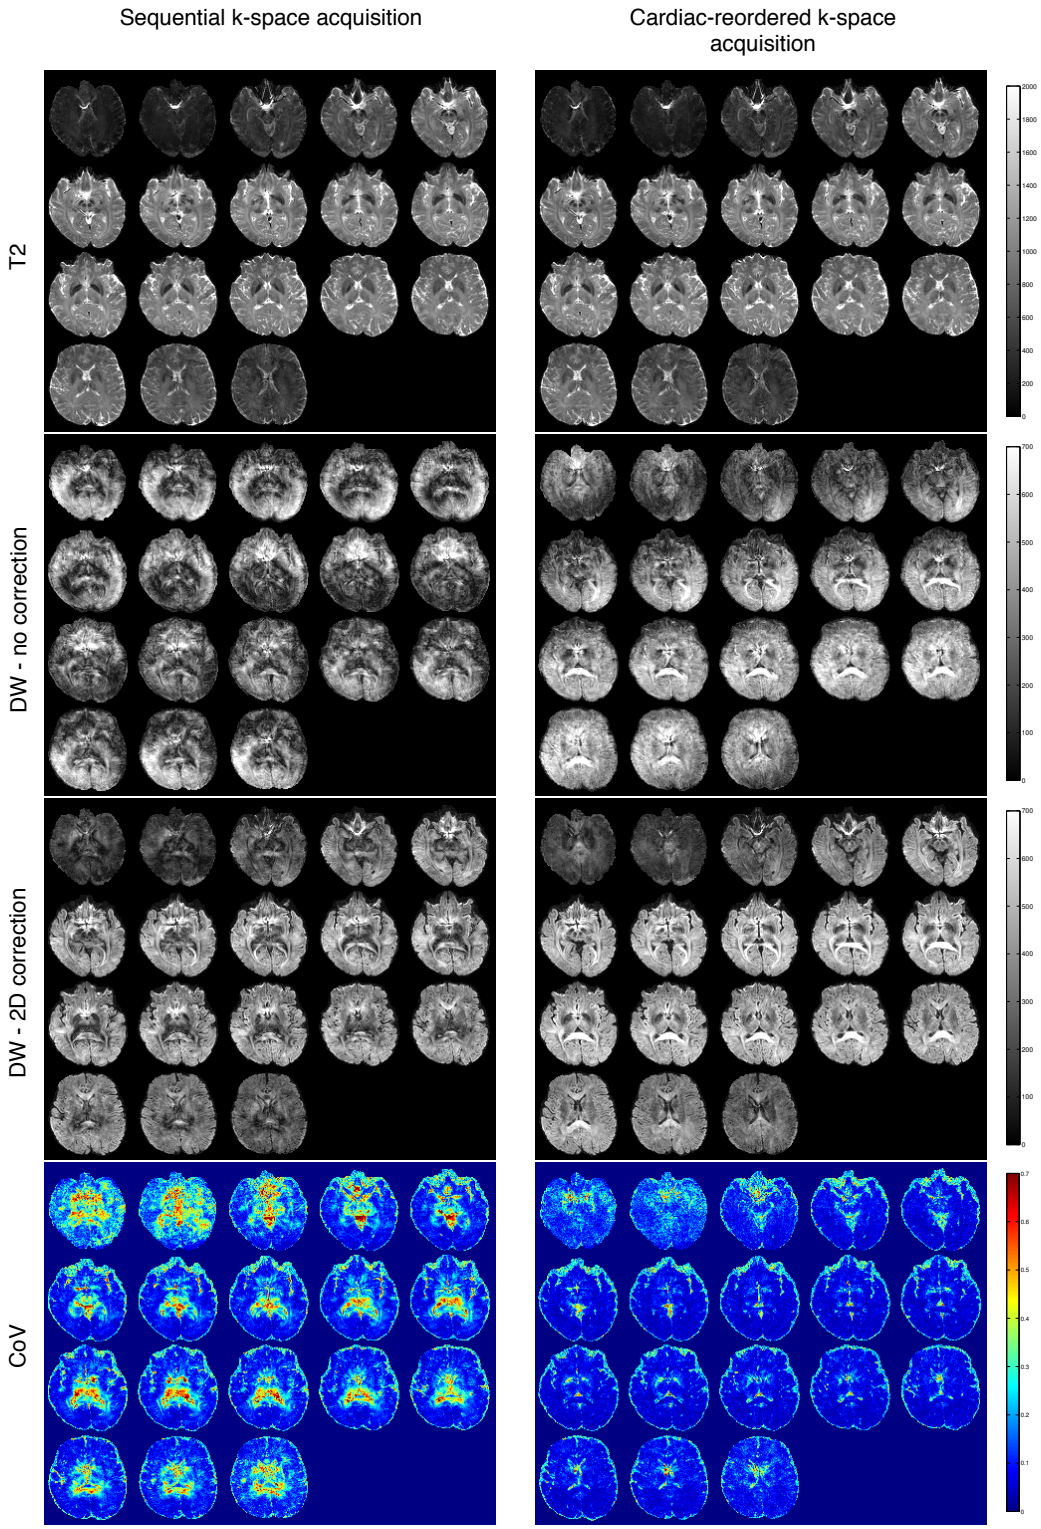

Figure 6.15: Experimental results showing a comparison of images and variability between sequential acquisition of k-space and cardiac-reordered acquisition of k-space in data with 18 k_z phase-encodes and 36 mm slab thickness. Diffusion weighting is in the z direction with b = 1000 s/mm^2. Resolution of the images is 1.8×1.8×2 mm. The coefficient of variation (CoV) is calculated from the temporal standard deviation divided pixel-by-pixel by the temporal mean over 6 repeated acquisitions in the data with 2D correction.

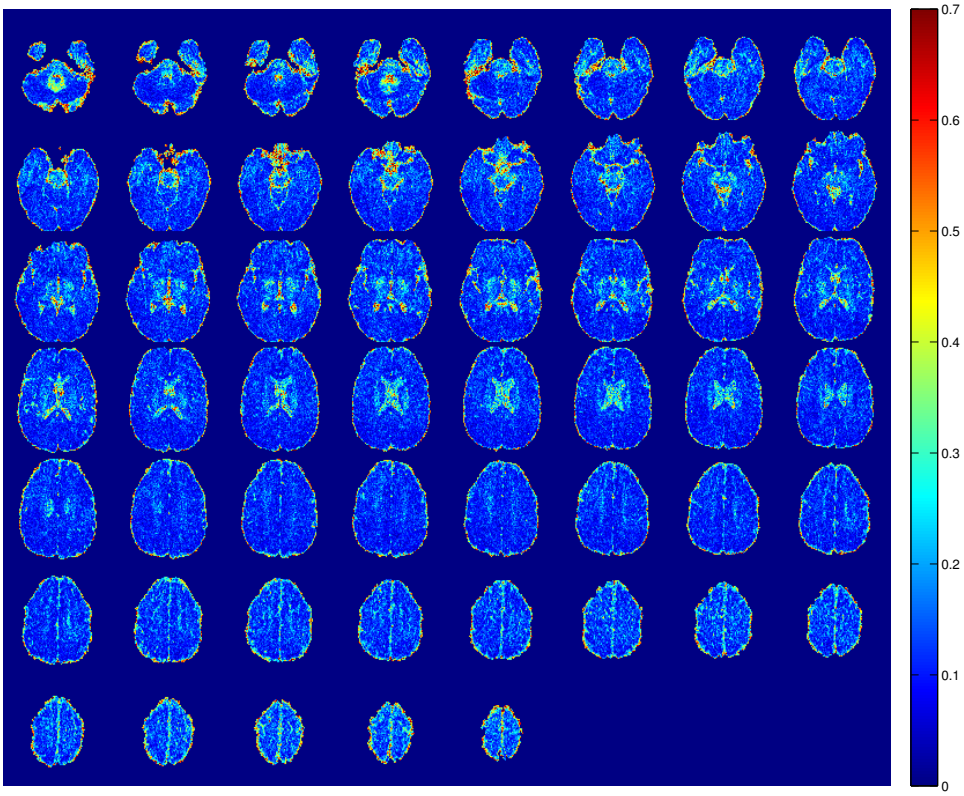

Figure 6.16: Experimental results showing maps of coefficient of variation in a standard ss-EPI protocol. Diffusion weighting is in the z direction with b = 1000 s/mm^2. Resolution of the images is 2×2×2 mm. The coefficient of variation (CoV) is calculated from the temporal standard deviation divided pixel-by-pixel by the temporal mean over 6 repeated acquisitions.

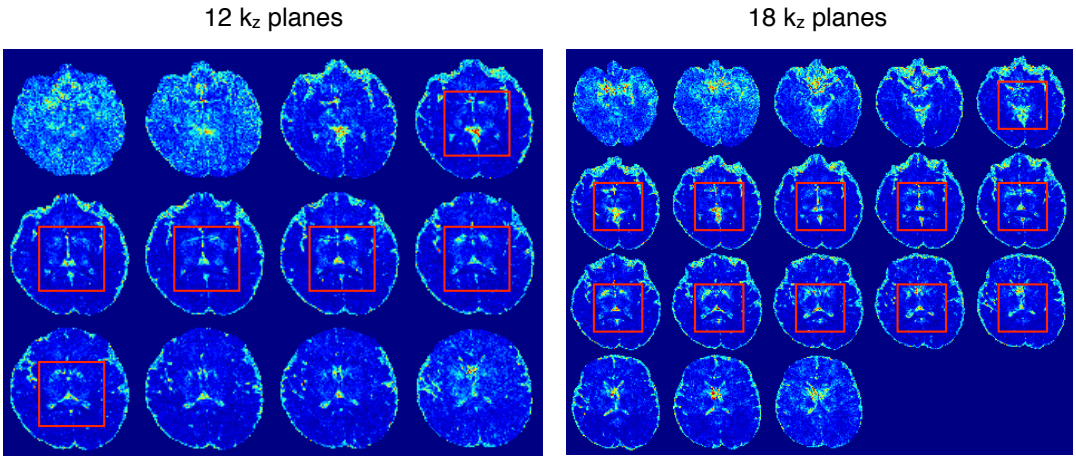

Figure 6.17: Locations of the ROIs used for quantitative comparisons of the coefficient of variation (CoV) in the *in vivo* data with 12 and 18 k_z phase-encodes (2 mm thick slices). The ROIs are shown on the cardiac-reordered CoV maps.

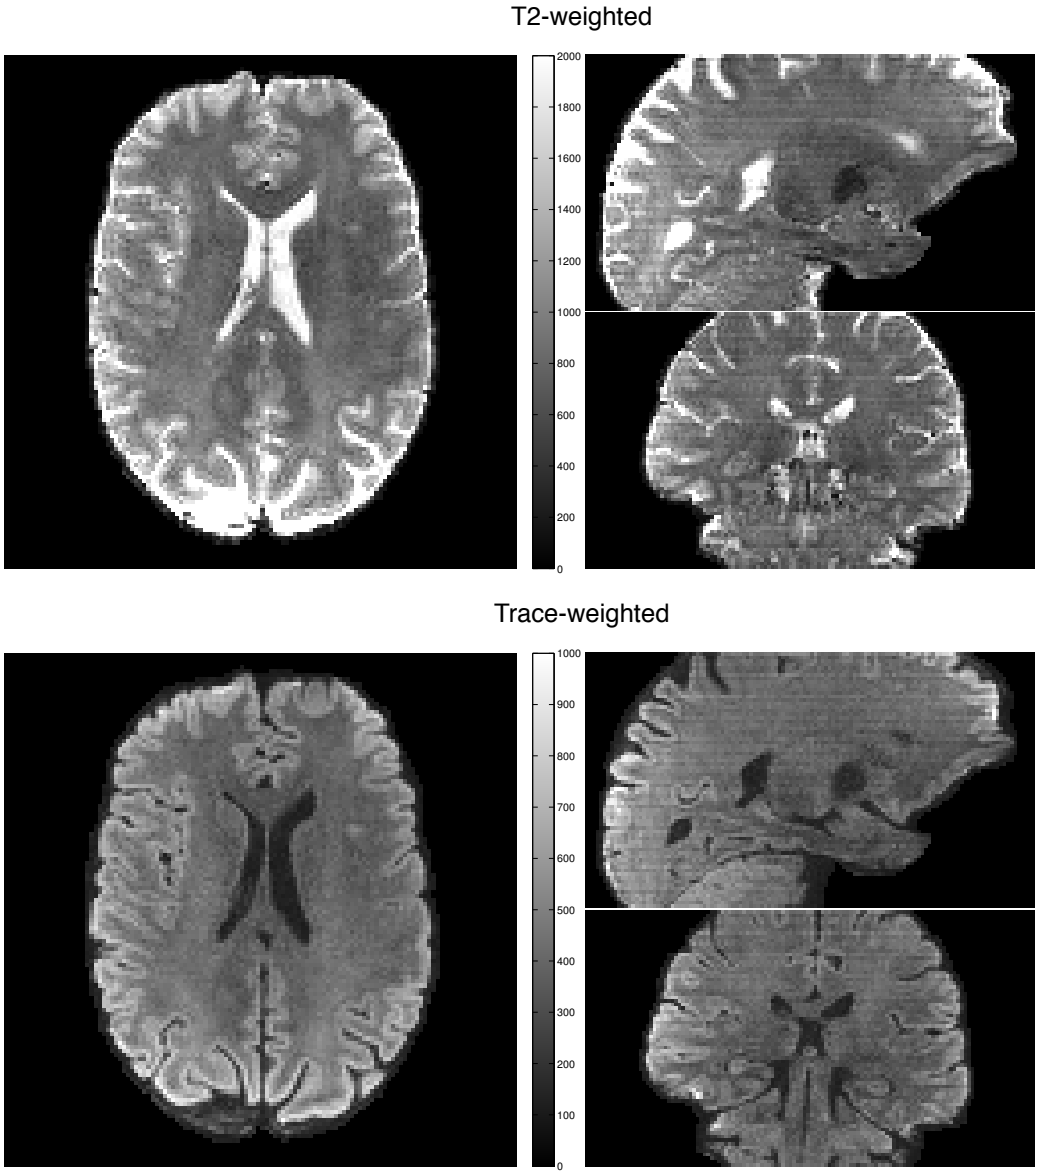

Figure 6.18: Experimental results of a multi-slab acquisition of 1.5 mm isotropic resolution T_2- and trace-weighted images with b = 1000 s/mm^2 diffusion-weighting.

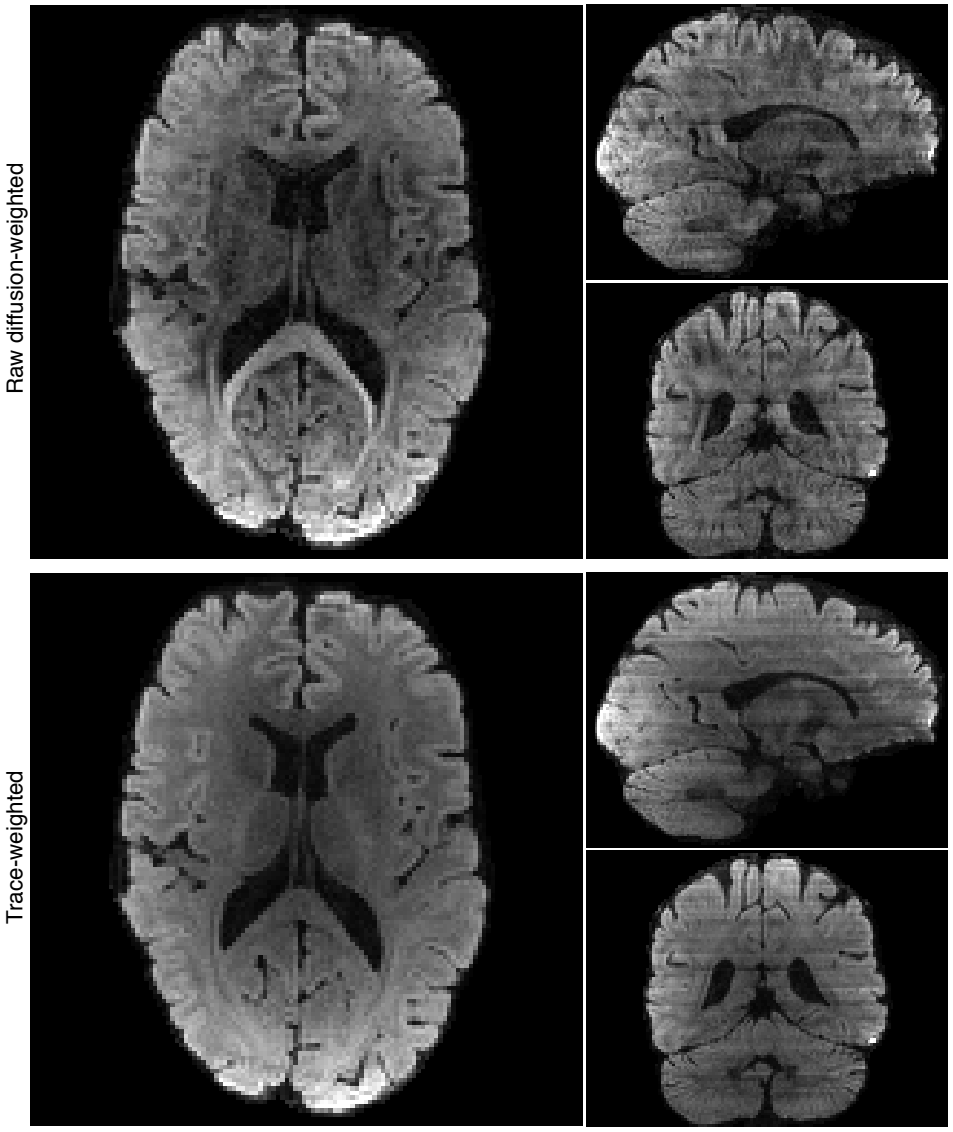

Figure 6.19: Experimental results of a whole brain simultaneous multi-slab acquisition. Raw diffusion- and trace-weighted images are shown.

Simulation			
k_z planes	8	16	24
Slab thickness (mm)	16	32	48
CoV: Sequential	0.071±0.081	0.144±0.199	0.155±0.176
CoV: Cardiac-reordered	0.044±0.036	0.059±0.065	0.075±0.083
% improvement	38	59	52

Experiment		
k_z planes	12	18
Slab thickness (mm)	24	36
CoV: Sequential	0.154±0.103	0.216±0.146
CoV: Cardiac-reordered	0.086±0.068	0.098±0.086
% improvement	44	55

Table 6.1: Comparison of the variability in sequential and cardiac-reordered k-space acquisitions in simulations and experiment. The coefficient of variation (CoV) was compared between the simulations and the experiments. The CoV reported is the mean within the ROIs defined by the voxels in the phantom and in Fig. 6.17 (for the simulation and experimental results, respectively) and the error is the standard deviation within the ROI. Both the simulation and experiment results are when using the 2D navigator correction (without reacquisition) with 6 volumes of b = 1000 s/mm^2 diffusion weighting in the z direction.

6.5 Discussion

This study has demonstrated through simulation and experiment that 2D navigation can be used in thin-slab 3D rs-EPI DW acquisitions and that reordering the k-space acquisition with respect to the cardiac cycle results in clear improvements in image quality. Acquisition of 2D imaging (IM_{2D}) data was chosen based on PSF simulations, which suggested that 3D imaging (IM_{3D}) schemes would suffer from substantial aliasing artefacts. Simulations of motion phase artefacts using a model for brain velocity suggested 2D and 3D corrections (IM_{2D}+NAV_{2D} and IM_{2D}+NAV_{3D}) were possible and that cardiac-reordering of the k-space acquisition and reacquisition of motion corrupted shots were beneficial. Further work on the 3D navigator reconstruction of experimental data with a 3D navigator (scheme IM_{2D}+NAV_{3D}) is required. Modification of the current 2D on-line reconstruction code to accept 3D phase-encoded data would enable re-acquisition, which is expected to improve image quality further.

Predictions made in the simulation environment were confirmed experimentally with successful 2D navigator correction and reduction in variability in repeated acquisitions when diffusion weighting in z (the most problematic encoding direction for DW artefacts). In quantitative analyses of the reduction of image artefacts a 40-50% reduction in variability was observed with cardiac-reordering, in both simulations and experiments. This demonstrates the accuracy of the simulations, which can be used to test further improvements in reconstruction and reordering strategies and evaluation of motion corruption for reacquisition purposes. Ideally, all shots would be acquired as cleanly as possible with no data acquisition during systole (which is the not case in the current reordering scheme). However, as mentioned above, cardiac gating approaches only allow 2-3 slabs per R-R interval and would result in variable and longer acquisition times, which are already long in 3D rs-EPI due to the segmentation of k-space in k_x and k_z. The reordering strategy should also be applicable to other segmented 3D DW sequences. Heart-rate variability could interfere with the cardiac-reordering scheme but in the healthy volunteers scanned to date it has not been a significant problem. Implementing a reacquisition of the shots with worst phase corruption is expected to improve residual levels of phase corruption caused by shots which were not acquired at the optimal time. The adaptive GRAPPA method proposed by Tijssen *et al.* [184] for filling in corrupt or poorly ordered shots is not possible in this multi-slab case due to lack of coil coverage in the slab direction.

Artefacts at the interfaces between slabs are currently evident in the multi-slab data. Combined slab overlap and k_z oversampling and concatenated scanning improved the data but further optimisation of the slab overlap and RF pulses is required to generate clean, isotropic data. Recent work in a 3D DW scan (TR ~5 s) with 2D navigator correction used high time-bandwith product 10 ms refocusing pulses and regridding of slabs with a slab profile weighting to generate data

without visible slab crosstalk effects [112]. Although a TR = 1-2s is optimal for SNR efficiency, the increased T_1 saturation effects that cause slab-joining artefacts may necessitate a longer TR. To mitigate this, scan times can be accelerated by using a simultaneous multi-slab acquisition to acquire multiple slabs at once [71, 74]. A whole brain simultaneous multi-slab acquisition was also demonstrated, which is an extension of work from Chapter 5. The blipped-CAIPI method [72, 74] of shifting slices in-plane to improve SNR in the unaliasing reconstruction was not used in this preliminary demonstration. Also, partial Fourier encoding can be used to reduce the number of readout segments required in the k_x dimension [135].

6.6 Conclusions

A 3D rs-EPI multi-slab sequence for acquisition of true isotropic voxels with an SNR-optimal TR has been demonstrated. An acquisition trajectory for 3D imaging was chosen based on simulations of PSF artefacts. Using this trajectory, a method involving 2D navigator correction and cardiac-reordering of the k-space acquisition was tested in simulation and as implemented *in vivo*. This showed good quality DW data with low levels of motion-phase artefact. A quantitative analysis of simulated and experimental data suggests a 40-50% reduction in diffusion phase-induced variability instabilities when using real time cardiac-reordering of k-space shots compared to a sequential acquisition. Improvements in slab joining artefacts, implementation of 3D navigation and navigator-based reacquisition are the subject of future work. A simultaneous multi-slab acquisition acceleration strategy was also demonstrated.

Chapter 7

Summary and Future Work

7.1 Thesis Summary

Readout-segmented EPI has previously been demonstrated as a successful multi-shot diffusion-weighted sequence that overcomes. The multi-shot motion-induced phase problem is addressed with non-linear navigator correction and reacquisition of severely corrupted shots. The sequence is attractive due to its close relationship to the standard diffusion imaging method single-shot EPI, which is popular due to its high imaging speed and low sensitivity to motion artefacts. The segmentation of k-space in the readout direction in rs-EPI allows acquisition of high-resolution diffusion-weighted images with reduced geometric distortions and blurring, relative to ss-EPI. In this thesis, modifications to the pulse sequence and reconstruction were developed and implemented to extend the sequence to incorporate new encoding approaches. Changes were made to reduce rs-EPI scan times, which are necessarily longer than single-shot approaches due to the multi-shot nature of the sequence. The reduction of rs-EPI scan times to be more comparable to the duration of current ss-EPI diffusion protocols is important for adoption of the rs-EPI technique in clinical stroke and DTI settings.

Furthermore, a 3D multi-slab extension of the original 2D multi-slice sequence was presented. The 3D sequence could allow isotropic high-resolution DW data, which is difficult to attain in 2D multi-slice sequences where RF excitation and SNR considerations limit resolution in the slice direction. A 3D method would also enable imaging with a optimally SNR efficient TR of 1-2 s.

The first of two approaches to reducing the scan time was to acquire fewer readout segments on one side of k-space and to estimate the missing data with a partial Fourier algorithm. This method allows reductions in scan time by the partial Fourier factor (scan time reduction in partial Fourier acquisitions is limited to a factor less than two) but reduces the SNR of the reconstructed images by the square root of the partial Fourier factor. The homodyne and POCS reconstruction algorithms were compared to the gold standard full k-space reconstructions in terms of the accuracy of their estimates of diffusion parameters from 2 mm isotropic DTI acquisitions (30 directions) and signal intensity in $0.9\times0.9\times4$ mm trace-weighted acquisitions. An ROI analysis in commonly-studied white matter fibre tracts of estimates of FA and MD showed that POCS values were closest to the full k-space values. Also, the differences were less than previously reported changes due to inter-subject and inter-session variability and pathology in patient versus control studies. The uncertainties on the estimate of the principal diffusion direction (which dictates the performance of tractography) in partial Fourier encoded data with POCS reconstruction were shown to be comparable to the white matter uncertainties in ss-EPI DTI data (60 directions) and previously reported studies. It was shown that reliable POCS partial Fourier rs-EPI data could be acquired when the scan times were reduced by 3/5 in the 2 mm isotropic acquisition and by 9/13 in the $0.9\times0.9\times4$ mm acquisition.

The second acceleration strategy was a simultaneous acquisition of multiple slices using the recently published blipped-CAIPI method that was applied to ss-

EPI. The simultaneous multi-slice acquisition allowed a lower TR because fewer excitations are required to achieve the same slice coverage when the signal from multiple slices is excited and encoded together. SNR losses in the separation of the slice-multiplexed data were minimised by using a FOV/2 blipped-CAIPI gradient scheme to apply relative in-plane shifts of the slices, thereby reducing the g-factor of the reconstruction. The blipped-CAIPI modification to the rs-EPI sequence was demonstrated with factor two slice acceleration and shown to be compatible with the navigator correction and the navigator-based reacquisition. A comparison of blipped-CAIPI rs-EPI and ss-EPI data in 2 mm and 1.5 mm isotropic DTI acquisitions (60 directions, blipped-CAIPI rs-EPI scan times matched by 3 averages of ss-EPI protocol) demonstrated the reduced blurring in rs-EPI as a consequence of the shorter readout. Tractography using this data showed that rs-EPI clearly delineated fibre tracts and dispersed less than ss-EPI although there were more streamlines reaching cortex in the higher SNR ss-EPI data. The blipped-CAIPI rs-EPI sequence was also used for high-resolution trace-weighted measurements that would be useful for visualisation of small ischemic lesions, such as the one shown in a 1.1×1.1×4 mm blipped-CAIPI rs-EPI trace-weighted image in a TIA patient.

The 3D multi-slab rs-EPI sequence was implemented after considering possible acquisition trajectories and simulating their point-spread functions. Motion artefacts were investigated in the chosen acquisition trajectory through simulations that used a model for the non-linear motion of the brain. The simulation environment informed on the efficacy of reconstruction approaches and the influence of slab thickness and b-value on image artefacts. A scheme to reorder the acquisition of k-space shots at sequence run-time based on information from the cardiac cycle was shown to substantially reduce image artefacts, by approximately 40% in a comparison of the variability in a time-series of repeated images.

The reordering scheme ensures that cleaner data is acquired at the centre of 3D k-space (where motion corruption has a larger effect on image quality) and minimises discontinuities and periodicity in signal amplitude and phase. The 2D multi-slice rs-EPI sequence was modified to excite and encode 3D slabs and the cardiac-reordering scheme was implemented. A decision on which k-space shot to acquire at each excitation was made at sequence run-time based on the position in the cardiac cycle determined from pulse oximeter information. The simulation predictions were confirmed in experiment with approximately 50% reduction in image artefacts compared to a conventional sequential k-space acquisition. The residual image artefacts in the cardiac-reordered acquisition were shown to be similar to those encountered in a standard ss-EPI acquisition. The multi-slab sequence required consideration of slab-joining artefacts, which were mitigated by a combination of overlapping slabs and k_z oversampling. A cardiac-reordered 1.5 mm isotropic trace-weighted acquisition was demonstrated as well as a preliminary simultaneous multi-slab version using the methods from the work on blipped-CAIPI rs-EPI.

7.2 Future Work

On-line Image Reconstruction

All of the image reconstruction of the modifications to the rs-EPI sequence presented in this thesis was performed off-line in Matlab. Image reconstruction on the scanner requires software written in the Siemens ICE environment and translation of Matlab code to ICE was beyond the scope of this work. The ability to view images on the scanner is clearly necessary before the modified rs-EPI sequence can be distributed. Scanner reconstruction would facilitate use of the sequence in clinical stroke settings and for DTI acquisitions for neuroscience re-

search purposes. It would also enable developments to aspects of the sequence related to the reacquisition of motion corrupted segments, which would be of interest in blipped-CAIPI rs-EPI and 3D rs-EPI, as discussed in the relevant sections below.

Blipped-CAIPI rs-EPI

In the blipped-CAIPI rs-EPI experiments the multiband pulses were designed with a simple sum of two phase-modulated waveforms and therefore the number of slices was constrained by limits on B_1 amplitude and SAR. A more sophisticated RF pulse design would allow simultaneous excitation of more than two slices and hence improvements in slice acceleration, potentially to recently demonstrated factors of 6-8 [194]. Two possible options are to use the VERSE algorithm [169] to reduce the B_1 amplitude of the existing pulses or PINS pulses [170] where RF power is independent of the number of slices. Another simpler strategy for summed pulses, would be to shift the peaks of the pulses relative to each other in time, which reduces the peak B_1 of the composite pulse and can therefore be effective for reducing SAR. In highly-accelerated acquisitions the performance of the unaliasing reconstruction is more important because the aliased pixels are closer together and the interaction between aliasing due to separate in-plane and slice acceleration must be considered. Hence the reconstruction should be investigated further, including recently published improvements to the slice-GRAPPA reconstruction [171].

In the case of two simultaneously excited slices, the navigator-based reacquisition was able to detect motion corruption in either of the slices because the distribution width considers signal from all coil elements. In acquisitions with higher slice acceleration it may be necessary to unalias the slice-multiplexed data on-line so that corruption can be judged separately in each of the slices. This would

require implementation of slice-GRAPPA ICE reconstruction. As an alternative to on-line unaliasing, the distribution width could be evaluated on a coil-specific basis so that the coils closest to each slice could be used to judge corruption.

3D rs-EPI

The multi-slab data suffers from residual artefacts at the interfaces between slabs that are exacerbated by the low TR chosen to increase SNR efficiency and reduce scan time. Further optimisation of RF pulses, slab overlap and k_z oversampling are required to improve the data. The possibilty of implementing a 3D navigator correction in the sequence should also be explored.

The 3D experimental data in this thesis did not reacquire the shots with the worst motion corruption as was done in the 2D sequences. Changes to scanner reconstruction code would allow navigator-based reacquisition, which currently is not possible because the ICE code for the 2D sequence does not accept the 3D encoded data. Reacquistion is expected to improve the data quality and robustness to motion artefacts.

The preliminary simultaneous multi-slab acquisition can be extended to include in-plane shifting of data from different slabs with the blipped-CAIPI slice gradient scheme. The effect of through-slab blurring due to the gradient blips and any interaction with the k_z phase-encoding should be considered. This modification would improve the slab unaliasing reconstruction and allow excitation of more closely-spaced slabs.

High-Resolution Data Acquisition at 7 T

Off-resonance artefacts and blurring increase with field strength and therefore rs-EPI has a greater advantage over ss-EPI at 7 T. The reduction in artefacts by using rs-EPI at 7 T has already been demonstrated with the original sequence [130].

Also, the shorter echo times achievable with rs-EPI afford greater SNR or opportunity for stronger diffusion encoding, which is more noticeable at 7 T because T_2 relaxation is faster. An obvious extension to the modified sequence would be to acquire higher resolution data at 7 T. Even higher resolution of localised regions could also be obtained by reducing the FOV as was recently demonstrated with sub-millimetre DTI [195].

Appendix B

Algorithm for Cardiac Reordering of k-space Shots

This Appendix provides Matlab code of the algorithm used in Chapter 6 to choose which k-space shot to acquire based on the value of the cardiac phase, ϕ_c (`phic` in the code). ϕ_c was calculated by dividing the time since the cardiac trigger, t_{trig}, measured with a pulse oximeter placed on the index finger (a 200 ms delay time was determined empirically), by the mean cardiac cycle duration, T_{RR} (from the previous 10 cardiac cycles), as shown in Eq. B.1.

$$\phi_c = t_{trig}/T_{RR} \tag{B.1}$$

A cardiac phase of 0.4 - 0.6 was assumed to correspond to diastole and therefore to be the optimum time to acquire data. The following algorithm was used to determine the k_z and k_x shot indices (`nz` and `nx`, respectively):

```
% phic  = cardiac phase
% Nx    = number of kx shots (readout segments)
% Nx    = number of kz phase-encodes

Ns = Nx*Nz;
nx_middle = ceil(Nx/2); % index of central readout segment

Ns = Nx*Nz; % total number of shots
nx_middle = ceil(Nx/2); % index of central kx readout segment

% shot index
ns = ceil(phic * Ns);

if rem(Nz,2)==0 % when the number of kz slices is even
    % adjust shot index so that central kx readout segment ...
       is the optimum shot (at phic=0.5)
    ns = ns + nx_middle;
    % if ns becomes larger than Ns loop back to ns=1,2.. etc.
    if ns > Ns
        ns = ns - Ns;
```

```
    end
end

% determine kz and kx shot indices from ns
nz = ceil(ns/Nx);
nx = rem((ns-1),Nx) + 1;

% e.g. with Nx=5, Nz=8, phic = 0.5:
%           ns = 23, nz=5 and nx=3
```

References

[1] R. Gilberto González. Clinical MRI of acute ischemic stroke. *Journal of Magnetic Resonance Imaging*, 36(2):259–271, 2012.

[2] Jeffrey L. Saver. Time is brain–quantified. *Stroke*, 37(1):263–266, 2006.

[3] P J Basser, J Mattiello, and D LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J Magn Reson B*, 103(3):247–254, Mar 1994.

[4] Heidi Johansen-Berg and Timothy E J Behrens. Just pretty pictures? What diffusion tractography can add in clinical neuroscience. *Curr Opin Neurol*, 19(4):379–385, Aug 2006.

[5] R Turner, D Le Bihan, J Maier, R Vavrek, L K Hedges, and J Pekar. Echo-planar imaging of intravoxel incoherent motion. *Radiology*, 177(2):407–414, Nov 1990.

[6] David A Porter and Robin M Heidemann. High resolution diffusion-weighted imaging using readout-segmented echo-planar imaging, parallel imaging and a two-dimensional navigator-based reacquisition. *Magn Reson Med*, 62(2):468–475, Aug 2009.

[7] E. Mark Haacke, Robert W. Brown, Michael R. Thompson, and Ramesh Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley-Liss, 1999.

[8] Peter Jezzard, Paul M Matthews, and Stephen M Smith. *Functional MRI: an introduction to methods*. Oxford University Press, 2001.

[9] Matt A Bernstein, K F King, and X J Zhou. *Handbook of MRI Pulse Sequences*. Elsevier Academic Press, 2004.

[10] Dwight G Nishimura. *Principles of Magnetic Resonance Imaging*. Stanford University, 2010.

[11] Walther Gerlach and Otto Stern. Über die richtungsquantelung im magnetfeld. *Annalen der Physik*, 379(16):673–699, 1924.

[12] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch. A New Method of Measuring Nuclear Magnetic Moment. *Phys. Rev.*, 53:318–318, Feb 1938.

[13] F. Bloch. Nuclear Induction. *Phys. Rev.*, 70:460–474, Oct 1946.

[14] E. M. Purcell, H. C. Torrey, and R. V. Pound. Resonance Absorption by Nuclear Magnetic Moments in a Solid. *Phys. Rev.*, 69:37–38, Jan 1946.

[15] P Mansfield and P K Grannell. NMR 'diffraction' in solids? *Journal of Physics C: Solid State Physics*, 6(22):L422, 1973.

[16] Paul C. Lauterbur. Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance. *Nature*, 242(5394):190–191, 03 1973.

[17] Ronald Bracewell. *The Fourier Transform and its Applications*. McGraw-Hill, 1978.

[18] Anil Kumar, Dieter Welti, and Richard R Ernst. NMR Fourier zeugmatography. *Journal of Magnetic Resonance (1969)*, 18(1):69–83, 4 1975.

[19] W A Edelstein, J M Hutchison, G Johnson, and T Redpath. Spin warp NMR imaging and applications to human whole-body imaging. *Phys Med Biol*, 25(4):751–756, Jul 1980.

[20] G N Hounsfield. Computerized transverse axial scanning (tomography). 1. Description of system. *Br J Radiol*, 46(552):1016–1022, Dec 1973.

[21] S Ljunggren. A simple graphical representation of Fourier-based imaging methods. *J Magn Reson*, 54:338–343, 1983.

[22] P Mansfield. Imaging by nuclear magnetic resonance. *J Phys E: Sci Instrum*, 21:18–30, 1988.

[23] P Mansfield. Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10(3), 1977.

[24] P Mansfield and I. L Pykett. Biological and medical imaging by NMR. *Journal of Magnetic Resonance (1969)*, 29(2):355–373, 2 1978.

[25] J Hennig, A Nauerth, and H Friedburg. RARE imaging: a fast imaging method for clinical MR. *Magn Reson Med*, 3(6):823–833, Dec 1986.

[26] G Johnson and J. M. S Hutchison. The limitations of NMR recalled-echo imaging techniques. *Journal of Magnetic Resonance (1969)*, 63(1):14–30, 6 1985.

[27] R J Ordidge and P Mansfield. Nuclear magnetic resonance methods. *United States patent 4509015*, 1985.

[28] I L Pykett and R R Rzedzian. Instant images of the body by magnetic resonance. *Magn Reson Med*, 5(6):563–571, Dec 1987.

[29] A M Howseman, M K Stehling, B Chapman, R Coxon, R Turner, R J Ordidge, M G Cawley, P Glover, P Mansfield, and R E Coupland. Improvements in snap-shot nuclear magnetic resonance imaging. *Br J Radiol*, 61(729):822–828, Sep 1988.

[30] P Margosian, F Schmitt, and D Purdy. Faster MR imaging: imaging with half the data. *Health Care Instrum*, 1:195–197, 1986.

[31] P B Roemer and J S Hickey. Self-shielded gradient coils for nuclear magnetic resonance imaging. *European patent application 87101198*, 1986.

[32] R Turner and R M Bowley. Passive screening of switched magnetic field gradients. *J Phys E: Sci Instrum*, 19:876–879, 1986.

[33] P Mansfield and B Chapman. Active magnetic screening of gradient coils in NMR imaging. *Journal of Magnetic Resonance (1969)*, 66(3):573–576, 2 1986.

[34] Robert Turner. Perfusion studies and fast imaging. In *Cerebral Blood Flow (Rescigno A, Boicelli A)*, pages 245–258, Plenum, New York, 1988.

[35] R Turner, D Le Bihan, and A S Chesnick. Echo-planar imaging of diffusion and perfusion. *Magn Reson Med*, 19(2):247–253, Jun 1991.

[36] J W Belliveau, D N Jr Kennedy, R C McKinstry, B R Buchbinder, R M Weisskoff, M S Cohen, J M Vevea, T J Brady, and B R Rosen. Functional mapping of the human visual cortex by magnetic resonance imaging. *Science*, 254(5032):716–719, Nov 1991.

[37] S Ogawa, D W Tank, R Menon, J M Ellermann, S G Kim, H Merkle, and K Ugurbil. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89(13):5951–5955, 1992.

[38] K K Kwong, J W Belliveau, D A Chesler, I E Goldberg, R M Weisskoff, B P Poncelet, D N Kennedy, B E Hoppel, M S Cohen, and R Turner. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci U S A*, 89(12):5675–5679, Jun 1992.

[39] P A Bandettini, E C Wong, R S Hinks, R S Tikofsky, and J S Hyde. Time course EPI of human brain function during task activation. *Magn Reson Med*, 25(2):390–397, Jun 1992.

[40] Franz Schmitt, Michael K. Stehling, Robert Turner, and Peter Mansfield. *Echo-Planar Imaging: Theory, Technique and Application*. Springer, 1998.

[41] D Q Chen, R B Marr, and P C Lauterbur. Reconstruction from NMR data acquired with imaging gradients having arbitrary time dependence. *IEEE Trans Med Imaging*, 5(3):162–164, 1986.

[42] A Zakhor, R Weisskoff, and R Rzedzian. Optimal sampling and reconstruction of MRI signals resulting from sinusoidal gradients. *IEEE Trans Signal Processing*, 39(9):2056–2065, September 1991.

[43] J I Jackson, C H Meyer, D G Nishimura, and A Macovski. Selection of a convolution function for Fourier inversion using gridding. *IEEE Trans Med Imaging*, 10(3):473–478, 1991.

[44] H Bruder, H Fischer, H E Reinfelder, and F Schmitt. Image reconstruction for echo planar imaging with nonequidistant k-space sampling. *Magn Reson Med*, 23(2):311–323, Feb 1992.

[45] C H Meyer, J M Pauly, A Macovski, and D G Nishimura. Simultaneous spatial and spectral selective excitation. *Magn Reson Med*, 15(2):287–304, Aug 1990.

[46] H W Park, D J Kim, and Z H Cho. Gradient reversal technique and its applications to chemical-shift-related NMR imaging. *Magn Reson Med*, 4(6):526–536, Jun 1987.

[47] J M Gomori, G A Holland, R I Grossman, W B Gefter, and R E Lenkinski. Fat suppression by section-select gradient reversal on spin-echo MR imaging. *Radiology*, 168(2):493–495, 1988.

[48] P Jezzard and R S Balaban. Correction for geometric distortion in echo planar images from B0 field variations. *Magn Reson Med*, 34(1):65–73, Jul 1995.

[49] Jesper L R Andersson, Stefan Skare, and John Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage*, 20(2):870–888, Oct 2003.

[50] Daniel Gallichan, Jesper L R Andersson, Mark Jenkinson, Matthew D Robson, and Karla L Miller. Reducing distortions in diffusion-weighted echo planar imaging with a dual-echo blip-reversed sequence. *Magn Reson Med*, 64(2):382–390, Aug 2010.

[51] F Farzaneh, S J Riederer, and N J Pelc. Analysis of T2 limitations and off-resonance effects on spatial resolution and artifacts in echo-planar imaging. *Magn Reson Med*, 14(1):123–139, Apr 1990.

[52] D C Noll, D G Nishimura, and A Macovski. Homodyne detection in magnetic resonance imaging. *IEEE Trans Med Imaging*, 10(2):154–163, 1991.

[53] E. M. Haacke, E. D. Lindskog, and W. Lin. A fast, iterative, partial-Fourier technique capable of local phase recovery. *Journal of Magnetic Resonance (1969)*, 92(1):126–145, 1991.

[54] R R Rzedzian. High speed, high resolution, spin echo imaging by mosaic scan and MESH. In *Proceedings of the Society of Magnetic Resonance in Medicine*, New York, USA, 1987 (abstract 51).

[55] B Chapman, R Turner, R J Ordidge, M Doyle, M Cawley, R Coxon, P Glover, and P Mansfield. Real-time movie imaging from a single cardiac cycle by NMR. *Magn Reson Med*, 5(3):246–254, Sep 1987.

[56] P Mansfield, A M Howseman, and R J Ordidge. Volumar imaging using NMR spin echoes: echo-volumar imaging (EVI) at 0.1 T. *J Phys E: Sci Instrum*, 22:324–330, 1989.

[57] A W Song, E C Wong, and J S Hyde. Echo-volume imaging. *Magn Reson Med*, 32(5):668–671, Nov 1994.

[58] P Mansfield, R Coxon, and J Hykin. Echo-volumar imaging (EVI) of the brain at 3.0 T: first normal volunteer and functional imaging results. *J Comput Assist Tomogr*, 19(6):847–852, Nov-Dec 1995.

[59] Thomas Witzel, Jonathan Polimeni, Graham C Wiggins, Fa-Hsuan Lin, S Biber, M Hamm, R Seethamraju, and Lawrence L Wald. Single-shot echo-volumar imaging using highly parallel detection. In *Proceedings of the 16th Annual Meeting of ISMRM*, Toronto, Canada, 2008 (abstract 2367).

[60] P B Roemer, W A Edelstein, C E Hayes, S P Souza, and O M Mueller. The NMR phased array. *Magn Reson Med*, 16(2):192–225, Nov 1990.

[61] David J Larkman and Rita G Nunes. Parallel magnetic resonance imaging. *Phys Med Biol*, 52(7):R15–55, Apr 2007.

[62] J. W Carlson. An algorithm for NMR imaging reconstruction based on multiple RF receiver coils. *Journal of Magnetic Resonance (1969)*, 74(2):376–380, 9 1987.

[63] J W Carlson and T Minemura. Imaging time reduction through multiple receiver coil data acquisition and image reconstruction. *Magn Reson Med*, 29(5):681–687, May 1993.

[64] D K Sodickson and W J Manning. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med*, 38(4):591–603, Oct 1997.

[65] K P Pruessmann, M Weiger, M B Scheidegger, and P Boesiger. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*, 42(5):952–962, Nov 1999.

[66] Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med*, 47(6):1202–1210, Jun 2002.

[67] P M Jakob, M A Griswold, R R Edelman, and D K Sodickson. AUTO-SMASH: a self-calibrating technique for SMASH imaging. SiMultaneous Acquisition of Spatial Harmonics. *MAGMA*, 7(1):42–54, Nov 1998.

[68] Felix A. Breuer, Stephan A. R. Kannengiesser, Martin Blaimer, Nicole Seiberlich, Peter M. Jakob, and Mark A. Griswold. General formulation for quantitative G-factor calculation in GRAPPA reconstructions. *Magnetic Resonance in Medicine*, 62(3):739–746, 2009.

[69] Philip M Robson, Aaron K Grant, Ananth J Madhuranthakam, Riccardo Lattanzi, Daniel K Sodickson, and Charles A McKenzie. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magn Reson Med*, 60(4):895–907, Oct 2008.

[70] Markus Weiger, Klaas P Pruessmann, and Peter Boesiger. 2d sense for faster 3d mri. *MAGMA*, 14(1):10–19, Mar 2002.

[71] D J Larkman, J V Hajnal, A H Herlihy, G A Coutts, I R Young, and G Ehnholm. Use of multicoil arrays for separation of signal from multiple slices simultaneously excited. *J Magn Reson Imaging*, 13(2):313–317, Feb 2001.

[72] Felix A Breuer, Martin Blaimer, Robin M Heidemann, Matthias F Mueller, Mark A Griswold, and Peter M Jakob. Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging. *Magn Reson Med*, 53(3):684–691, Mar 2005.

[73] Rita G Nunes, Joseph V Hajnal, Xavier Golay, and David J Larkman. Simultaneous slice excitation and reconstruction for single shot EPI. In *Proceedings of the 14th Annual Meeting of ISMRM*, Seattle, Washington, USA, 2006 (abstract 293).

[74] Kawin Setsompop, Borjan A Gagoski, Jonathan R Polimeni, Thomas Witzel, Van J Wedeen, and Lawrence L Wald. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn Reson Med*, 67(5):1210–1224, May 2012.

[75] Martin Blaimer, Felix A Breuer, Nicole Seiberlich, Matthias F Mueller, Robin M Heidemann, Vladimir Jellus, Graham Wiggins, Lawrence L Wald, Mark A Griswold, and Peter M Jakob. Accelerated volumetric MRI with a SENSE/GRAPPA combination. *J Magn Reson Imaging*, 24(2):444–450, Aug 2006.

[76] Steen Moeller, Essa Yacoub, Cheryl A Olman, Edward Auerbach, John Strupp, Noam Harel, and Kamil Ugurbil. Multiband multislice GE-EPI at 7 Tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magn Reson Med*, 63(5):1144–1153, May 2010.

[77] Robert Brown. A brief account of microscopical observations made in the months of June, July and August, 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies. *Phil. Mag.*, 4:161–173, 1828.

[78] A. Einstein. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik*, 322(8):549–560, 1905.

[79] A. Einstein. *Investigations on the theory of the Brownian movement*. Dover Publications, 1926.

[80] E. L. Hahn. Spin echoes. *Physical Review*, 80(4), 11 1950.

[81] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Physical Review*, 94(3), 05 1954.

[82] E. O. Stejskal and J. E. Tanner. Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient. *The Journal of Chemical Physics*, 42(1):288–292, 1965.

[83] Christian Beaulieu. The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR Biomed*, 15(7-8):435–455, Nov-Dec 2002.

[84] M E Moseley, Y Cohen, J Mintorovitch, L Chileuitt, H Shimizu, J Kucharczyk, M F Wendland, and P R Weinstein. Early detection of regional cerebral ischemia in cats: comparison of diffusion- and T2-weighted MRI and spectroscopy. *Magn Reson Med*, 14(2):330–346, May 1990.

[85] S Warach, D Chien, W Li, M Ronthal, and R R Edelman. Fast magnetic resonance diffusion-weighted imaging of acute human stroke. *Neurology*, 42(9):1717–1723, Sep 1992.

[86] S Mori and P C van Zijl. Diffusion weighting by the trace of the diffusion tensor within a single scan. *Magn Reson Med*, 33(1):41–52, Jan 1995.

[87] E C Wong, R W Cox, and A W Song. Optimized isotropic diffusion weighting. *Magn Reson Med*, 34(2):139–143, Aug 1995.

[88] Derek K Jones. The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study. *Magn Reson Med*, 51(4):807–815, Apr 2004.

[89] P J Basser. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR Biomed*, 8(7-8):333–344, Nov-Dec 1995.

[90] C Pierpaoli and P J Basser. Toward a quantitative assessment of diffusion anisotropy. *Magn Reson Med*, 36(6):893–906, Dec 1996.

[91] S Pajevic and C Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. *Magn Reson Med*, 42(3):526–540, Sep 1999.

[92] Thomas E. Conturo, Nicolas F. Lori, Thomas S. Cull, Erbil Akbudak, Abraham Z. Snyder, Joshua S. Shimony, Robert C. McKinstry, Harold Burton, and Marcus E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences*, 96(18):10422–10427, 1999.

[93] S Mori, B J Crain, V P Chacko, and P C van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol*, 45(2):265–269, Feb 1999.

[94] P J Basser, S Pajevic, C Pierpaoli, J Duda, and A Aldroubi. In vivo fiber tractography using DT-MRI data. *Magn Reson Med*, 44(4):625–632, Oct 2000.

[95] D Greitz, R Wirestam, A Franck, B Nordell, C Thomsen, and F Stahlberg. Pulsatile brain movement and associated hydrodynamics studied by magnetic resonance phase imaging. The Monro-Kellie doctrine revisited. *Neuroradiology*, 34(5):370–380, 1992.

[96] D R Enzmann and N J Pelc. Brain motion: measurement with phase-contrast MR imaging. *Radiology*, 185(3):653–660, Dec 1992.

[97] B P Poncelet, V J Wedeen, R M Weisskoff, and M S Cohen. Brain parenchyma motion: measurement with cine echo-planar MR imaging. *Radiology*, 185(3):645–651, Dec 1992.

[98] A W Anderson and J C Gore. Analysis and correction of motion artifacts in diffusion weighted imaging. *Magn Reson Med*, 32(3):379–387, Sep 1994.

[99] K Butts, A de Crespigny, J M Pauly, and M Moseley. Diffusion-weighted interleaved echo-planar imaging with a pair of orthogonal navigator echoes. *Magn Reson Med*, 35(5):763–770, May 1996.

[100] M A Griswold, P M Jakob, Q Chen, J W Goldfarb, W J Manning, R R Edelman, and D K Sodickson. Resolution enhancement in single-shot imaging using simultaneous acquisition of spatial harmonics (SMASH). *Magn Reson Med*, 41(6):1236–1245, Jun 1999.

[101] Robin M Heidemann, Mark A Griswold, David A Porter, Berthold Kiefer, Mathias Nittka, J Wang, Axel Haase, and Peter M Jakob. Minimizing distortions and blurring in diffusion weighted single shot EPI using high performance gradients in combination with parallel imaging. In *Proceedings of the 9th Annual Meeting of ISMRM*, Glasgow, Scotland, 2001 (abstract 169).

[102] R Bammer, S L Keeling, M Augustin, K P Pruessmann, R Wolf, R Stollberger, H P Hartung, and F Fazekas. Improved diffusion-weighted single-shot echo-planar imaging (EPI) in stroke using sensitivity encoding (SENSE). *Magn Reson Med*, 46(3):548–554, Sep 2001.

[103] James G Pipe, Victoria G Farthing, and Kirsten P Forbes. Multishot diffusion-weighted FSE using PROPELLER MRI. *Magn Reson Med*, 47(1):42–52, Jan 2002.

[104] R J Ordidge, J A Helpern, Z X Qing, R A Knight, and V Nagesh. Correction of motional artifacts in diffusion-weighted MR images using navigator echoes. *Magn Reson Imaging*, 12(3):455–460, 1994.

[105] A J de Crespigny, M P Marks, D R Enzmann, and M E Moseley. Navigated diffusion imaging of normal and ischemic human brain. *Magn Reson Med*, 33(5):720–728, May 1995.

[106] K Butts, J Pauly, A de Crespigny, and M Moseley. Isotropic diffusion-weighted and spiral-navigated interleaved EPI for routine imaging of acute stroke. *Magn Reson Med*, 38(5):741–749, Nov 1997.

[107] D Atkinson, D A Porter, D L Hill, F Calamante, and A Connelly. Sampling and reconstruction effects due to motion in diffusion-weighted interleaved echo planar imaging. *Magn Reson Med*, 44(1):101–109, Jul 2000.

[108] Karla L Miller and John M Pauly. Nonlinear phase correction for navigated diffusion imaging. *Magn Reson Med*, 50(2):343–353, Aug 2003.

[109] Jennifer A McNab, Daniel Gallichan, and Karla L Miller. 3D steady-state diffusion-weighted imaging with trajectory using radially batched internal navigator echoes (TURBINE). *Magn Reson Med*, 63(1):235–242, Jan 2010.

[110] Lawrence R. Frank, Youngkyoo Jung, Souheil Inati, J. Michael Tyszka, and Eric C. Wong. High efficiency, low distortion 3D diffusion tensor imaging with variable density spiral fast spin echoes (3D DW VDS RARE). *NeuroImage*, 49(2):1510–1523, Jan 2010.

[111] M Engström, A Nordell, M Mårtensson, B Nordell, Roland Bammer, and Stefan Skare. Isotropic resolution in Diffusion Weighted Imaging using 3D multi-slab, multi-echo Echo Planar Imaging. In *Proceedings of the 18th Annual Meeting of ISMRM*, Stockholm, Sweden, 2010 (abstract 1619).

[112] M Engström, R Bammer, and S Skare. 2D Navigated 3D multi-slab DWI at 1.3 mm isotropic resolution. In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 117).

[113] Anh T Van, Diego Hernando, and Bradley P Sutton. Motion-induced phase error estimation and correction in 3D diffusion tensor imaging. *IEEE Trans Med Imaging*, 30(11):1933–1940, Nov 2011.

[114] T G Reese, O Heid, R M Weisskoff, and V J Wedeen. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magn Reson Med*, 49(1):177–182, Jan 2003.

[115] P Jezzard, A S Barnett, and C Pierpaoli. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magn Reson Med*, 39(5):801–812, May 1998.

[116] Jesper L R Andersson and Stefan Skare. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *Neuroimage*, 16(1):177–199, May 2002.

[117] John N Morelli, Val M Runge, Thorsten Feiweier, John E Kirsch, Kenneth W Williams, and Ulrike I Attenberger. Evaluation of a modified Stejskal-Tanner diffusion encoding scheme, permitting a marked reduction in TE, in diffusion-weighted imaging of stroke patients at 3 T. *Invest Radiol*, 45(1):29–35, Jan 2010.

[118] J E Tanner. Use of Stimulated Echo in NMR Diffusion Studies. *J. Chem. Phys.*, 52(5):2523–6, 1970.

[119] Klaus-Dietmar Merboldt, Wolfgang Hanicke, and Jens Frahm. Self-diffusion NMR imaging using stimulated echoes. *Journal of Magnetic Resonance (1969)*, 64(3):479 – 486, 1985.

[120] T P Trouard, R J Theilmann, M I Altbach, and A F Gmitro. High-resolution diffusion imaging with DIFRAD-FSE (diffusion-weighted radial acquisition of data with fast spin-echo) MRI. *Magn Reson Med*, 42(1):11–18, Jul 1999.

[121] H Gudbjartsson, S E Maier, R V Mulkern, I A Morocz, S Patz, and F A Jolesz. Line scan diffusion imaging. *Magn Reson Med*, 36(4):509–519, Oct 1996.

[122] D G Norris. Implications of bulk motion for diffusion-weighted imaging experiments: effects, mechanisms, and solutions. *J Magn Reson Imaging*, 13(4):486–495, Apr 2001.

[123] Robin M Heidemann, Alfred Anwander, Thorsten Feiweier, Thomas R Knosche, and Robert Turner. k-space and q-space: combining ultra-high spatial and angular resolution in diffusion imaging using ZOOPPA at 7 T. *Neuroimage*, 60(2):967–978, Apr 2012.

[124] Chunlei Liu, Roland Bammer, Dong-Hyun Kim, and Michael E Moseley. Self-navigated interleaved spiral (SNAILS): application to high-resolution diffusion tensor imaging. *Magn Reson Med*, 52(6):1388–1396, Dec 2004.

[125] Rita G Nunes, Peter Jezzard, Timothy E J Behrens, and Stuart Clare. Self-navigated multishot echo-planar pulse sequence for high-resolution diffusion-weighted imaging. *Magn Reson Med*, 53(6):1474–1478, Jun 2005.

[126] M D Robson, A W Anderson, and J C Gore. Diffusion-weighted multiple shot echo planar imaging of humans without navigation. *Magn Reson Med*, 38(1):82–88, Jul 1997.

[127] D A Porter and E Mueller. Multi-shot diffusion-weighted EPI with readout mosaic segmentation and 2D navigator correction. In *Proceedings of the 12th Annual Meeting of ISMRM*, Kyoto, Japan, 2004 (abstract 442).

[128] Q Nguyen, M Clemence, J Thornton, and R J Ordidge. Minimizing distortions and blurring in diffusion isotropic diffusion-weighted multishot imaging using automatic reacquisition. In *Proceedings of the 7th Annual Meeting of ISMRM*, Philadelphia, Pennsylvania, USA, 1999 (abstract 559).

[129] David A Porter. 2D-navigator-based re-acquisition for motion artefact suppression in multi-shot, diffusion-weighted imaging. In *Proceedings of the 14th Annual Meeting of ISMRM*, Seattle, Washington, USA, 2006 (abstract 1047).

[130] Robin M Heidemann, David A Porter, Alfred Anwander, Thorsten Feiweier, Keith Heberlein, Thomas R Knosche, and Robert Turner. Diffusion imaging in humans at 7T using readout-segmented EPI and GRAPPA. *Magn Reson Med*, 64(1):9–14, Jul 2010.

[131] Q Nguyen, M Clemence, and R J Ordidge. The use of intelligent reacquisition to reduce scan time in MRI degraded by motion. In *Proceedings of the 6th Annual Meeting of ISMRM*, Sydney, Australia, 1998 (abstract 134).

[132] Samantha J Holdsworth, Stefan Skare, Rexford D Newbould, and Roland Bammer. Robust GRAPPA-accelerated diffusion-weighted readout-segmented (RS)-EPI. *Magn Reson Med*, 62(6):1629–1640, Dec 2009.

[133] S J Holdsworth, K Yeom, S Skare, A J Gentles, P D Barnes, and R Bammer. Clinical application of readout-segmented- echo-planar imaging for diffusion-weighted imaging in pediatric brain. *AJNR Am J Neuroradiol*, 32(7):1274–1279, Aug 2011.

[134] LL Hayes, RA Jones, S Palasis, D Aguilera, and DA Porter. Drop metastases to the pediatric spine revealed with diffusion-weighted MR imaging. *Pediatr Radiol*, Nov 2011.

[135] Robert Frost, David A Porter, Karla L Miller, and Peter Jezzard. Implementation and assessment of diffusion-weighted partial Fourier readout-segmented echo-planar imaging. *Magn Reson Med*, 68(2):441–451, Aug 2012.

[136] J P Wansapura, S K Holland, R S Dunn, and W S Jr Ball. NMR relaxation times in the human brain at 3.0 tesla. *J Magn Reson Imaging*, 9(4):531–538, Apr 1999.

[137] Greg J Stanisz, Ewa E Odrobina, Joseph Pun, Michael Escaravage, Simon J Graham, Michael J Bronskill, and R Mark Henkelman. T1, T2 relaxation and magnetization transfer in tissue at 3T. *Magn Reson Med*, 54(3):507–512, Sep 2005.

[138] Robert Frost, David A Porter, Gwenaelle Douaud, Peter Jezzard, and Karla L Miller. Reduction of diffusion-weighted readout-segmented EPI scan time using a blipped-CAIPI modification. In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 116).

[139] M E Moseley, J Kucharczyk, J Mintorovitch, Y Cohen, J Kurhanewicz, N Derugin, H Asgari, and D Norman. Diffusion-weighted MR imaging of acute stroke: correlation with T2-weighted and magnetic susceptibility-enhanced MR imaging in cats. *AJNR Am J Neuroradiol*, 11(3):423–429, May 1990.

[140] K J van Everdingen, J van der Grond, L J Kappelle, L M Ramos, and W P Mali. Diffusion-weighted magnetic resonance imaging in acute stroke. *Stroke*, 29(9):1783–1790, Sep 1998.

[141] O Wu, W J Koroshetz, L Ostergaard, F S Buonanno, W A Copen, R G Gonzalez, G Rordorf, B R Rosen, L H Schwamm, R M Weisskoff, and A G Sorensen. Predicting tissue outcome in acute human cerebral ischemia using combined diffusion- and perfusion-weighted MR imaging. *Stroke*, 32(4):933–942, Apr 2001.

[142] Jessica N E Redgrave, Shelagh B Coutts, Ursula G Schulz, Dennis Briley, and Peter M Rothwell. Systematic review of associations between the presence of acute ischemic lesions on diffusion-weighted imaging and clinical predictors of early stroke risk after transient ischemic attack. *Stroke*, 38(5):1482–1488, May 2007.

[143] Ahmet Baki Yagci, Nurgul Ozari, Zafer Aybek, and Ender Duzcan. The value of diffusion-weighted MRI for prostate cancer detection and localization. *Diagn Interv Radiol*, 17(2):130–134, Jun 2011.

[144] Robert Frost, David A Porter, Thorsten Feiweier, and Peter Jezzard. Homodyne reconstruction of partial Fourier readout-segmented EPI for diffusion imaging. In *Proceedings of the 18th Annual Meeting of ISMRM*, Stockholm, Sweden, 2010 (abstract 1625).

[145] Samantha J Holdsworth, Stefan Skare, Murat Aksoy, R L O'Halloran, and Roland Bammer. Comparison of two alternative approaches for diffusion-weighted readout-segmented (RS)-EPI. In *Proceedings of the 19th Annual Meeting of ISMRM*, Montreal, Canada, 2011 (abstract 3932).

[146] Bregman L M. Determination of a common point of convex sets by the method of successive projection (Theorems for determining common point of convex sets by method of successive projection). *Akademiia Nauk SSR, Doklady*, 162:487–490, 1965.

[147] Stephen M Smith, Heidi Johansen-Berg, Mark Jenkinson, Daniel Rueckert, Thomas E Nichols, Karla L Miller, Matthew D Robson, Derek K Jones, Johannes C Klein, Andreas J Bartsch, and Timothy E J Behrens. Acquisition and voxelwise analysis of multi-subject diffusion data with tract-based spatial statistics. *Nat Protoc*, 2(3):499–503, 2007.

[148] D K Jones, M A Horsfield, and A Simmons. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn Reson Med*, 42(3):515–525, Sep 1999.

[149] G McGibney, M R Smith, S T Nichols, and A Crawley. Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI. *Magn Reson Med*, 30(1):51–59, Jul 1993.

[150] T E J Behrens, M W Woolrich, M Jenkinson, H Johansen-Berg, R G Nunes, S Clare, P M Matthews, J M Brady, and S M Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med*, 50(5):1077–1088, Nov 2003.

[151] Stephen M Smith. Fast robust automated brain extraction. *Hum Brain Mapp*, 17(3):143–155, Nov 2002.

[152] T E J Behrens, H Johansen Berg, S Jbabdi, M F S Rushworth, and M W Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34(1):144–155, Jan 2007.

[153] 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*, 20(1):45–57, Jan 2001.

[154] R Bammer, M Augustin, S Strasser-Fuchs, T Seifert, P Kapeller, R Stollberger, F Ebner, H P Hartung, and F Fazekas. Magnetic resonance diffusion tensor imaging for characterizing diffuse and focal white matter abnormalities in multiple sclerosis. *Magn Reson Med*, 44(4):583–591, Oct 2000.

[155] O Ciccarelli, D J Werring, C A Wheeler-Kingshott, G J Barker, G J Parker, A J Thompson, and D H Miller. Investigation of MS normal-appearing brain using diffusion tensor MRI with clinical correlations. *Neurology*, 56(7):926–933, Apr 2001.

[156] E. Heiervang, T. E. J. Behrens, C. E. Mackay, M. D. Robson, and H. Johansen-Berg. Between session reproducibility and between subject variability of diffusion MR and tractography measures. *NeuroImage*, 33(3):867–877, 11 2006.

[157] Stephen M Smith, Mark Jenkinson, Heidi Johansen-Berg, Daniel Rueckert, Thomas E Nichols, Clare E Mackay, Kate E Watkins, Olga Ciccarelli, M Zaheer Cader, Paul M Matthews, and Timothy E J Behrens. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage*, 31(4):1487–1505, Jul 2006.

[158] Derek K Jones. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magn Reson Med*, 49(1):7–12, Jan 2003.

[159] Fu-Nien Wang, Teng-Yi Huang, Fa-Hsuan Lin, Hsiao-Wen Chung, David S Tuch, Ming-Chung Chou, Cheng-Yu Chen, and Kenneth K Kwong. Propeller EPI: application to diffusion tensor imaging. In *Proceedings of the 12th Annual Meeting of ISMRM*, Kyoto, Japan, 2004 (abstract 2462).

[160] Fu-Nien Wang, Teng-Yi Huang, Fa-Hsuan Lin, Tzu-Chao Chuang, Nan-Kuei Chen, Hsiao-Wen Chung, Cheng-Yu Chen, and Kenneth K Kwong. PROPELLER EPI: an MRI technique suitable for diffusion tensor imaging at high field strength with reduced geometric distortions. *Magn Reson Med*, 54(5):1232–1240, Nov 2005.

[161] Stefan Skare, Rexford D Newbould, Dave B Clayton, and Roland Bammer. Propeller EPI in the other direction. *Magn Reson Med*, 55(6):1298–1307, Jun 2006.

[162] Robert Frost, David A. Porter, Karla Miller, and Peter Jezzard. Implementation and assessment of diffusion-weighted partial Fourier readout-segmented echo-planar imaging. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2011.

[163] Martyn N J Paley, Kuan J Lee, James M Wild, Paul D Griffiths, and Elspeth H Whitby. Simultaneous parallel inclined readout image technique. *Magn Reson Imaging*, 24(5):557–562, Jun 2006.

[164] Kawin Setsompop, Borjan Gagoski, Jonathan Polimeni, Thomas Witzel, V J Wedeen, and Lawrence L Wald. Blipped CAIPIRHINA for simultaneous multi-slice EPI with reduced g-factor penalty. In *Proceedings of the 18th Annual Meeting of ISMRM*. Stockholm, Sweden, 2010 (abstract 551).

[165] Kawin Setsompop, J. A. Cohen-Adad, Jennifer A McNab, Borjan Gagoski, V J Wedeen, and Lawrence L Wald. Improving SNR per unit time in diffusion Imaging using a blipped-CAIPIRINHA simultaneous multi-slice EPI

acquisition. In *Proceedings of the 18th Annual Meeting of ISMRM*. Stockholm, Sweden, 2010 (abstract 187).

[166] David A Feinberg, Steen Moeller, Stephen M Smith, Edward Auerbach, Sudhir Ramanna, Matt F Glasser, Karla L Miller, Kamil Ugurbil, and Essa Yacoub. Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging. *PLoS One*, 5(12):e15710, 2010.

[167] Kawin Setsompop, Berkin Bilgic, Julien Cohen-Adad, M. Dylan Tisdall, Boris Keil, Thomas Witzel, Yogesh Rathi, Van J Wedeen, Elfar Adalsteinsson, and Lawrence L Wald. Whole-brain DSI in 4 minutes: sparse sampling in q-space with simultaneous multi-slice acquisition. In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 693).

[168] J Pauly, P Le Roux, D Nishimura, and A Macovski. Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm [NMR imaging]. *IEEE Trans Med Imaging*, 10(1):53–65, 1991.

[169] Steven Conolly, Dwight Nishimura, Albert Macovski, and Gary Glover. Variable-rate selective excitation. *Journal of Magnetic Resonance (1969)*, 78(3):440–458, 7 1988.

[170] David G Norris, Peter J Koopmans, Rasim Boyacioglu, and Markus Barth. Power Independent of Number of Slices (PINS) radiofrequency pulses for low-power simultaneous multislice excitation. *Magn Reson Med*, 66(5):1234–1240, Nov 2011.

[171] Kawin Setsompop and Lawrence L Wald. A multi-kernel approach for reducing inter-slice image ghosting in simultaneous multi-slice EPI. In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 3349).

[172] P J Basser, J Mattiello, and D LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophys J*, 66(1):259–267, Jan 1994.

[173] L R Frank. Anisotropy in high angular resolution diffusion-weighted MRI. *Magn Reson Med*, 45(6):935–939, Jun 2001.

[174] Lawrence R Frank. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magn Reson Med*, 47(6):1083–1099, Jun 2002.

[175] Eun-Kee Jeong, Seong-Eun Kim, and Dennis L Parker. High-resolution diffusion-weighted 3D MRI, using diffusion-weighted driven-equilibrium (DW-DE) and multishot segmented 3D-SSFP without navigator echoes. *Magn Reson Med*, 50(4):821–829, Oct 2003.

[176] Xavier Golay, Hangyi Jiang, Peter C M van Zijl, and Susumu Mori. High-resolution isotropic 3D diffusion tensor imaging of the human brain. *Magn Reson Med*, 47(5):837–843, May 2002.

[177] Jiun-Jie Wang, Ralf Deichmann, Robert Turner, and Roger Ordidge. 3D DT-MRI using a reduced-FOV approach and saturation pulses. *Magn Reson Med*, 51(4):853–857, Apr 2004.

[178] Eun-Kee Jeong, Seong-Eun Kim, Eugene G Kholmovski, and Dennis L Parker. High-resolution DTI of a localized volume using 3D single-shot diffusion-weighted STimulated echo-planar imaging (3D ss-DWSTEPI). *Magn Reson Med*, 56(6):1173–1181, Dec 2006.

[179] Karla L Miller, Brian A Hargreaves, Garry E Gold, and John M Pauly. Steady-state diffusion-weighted imaging of in vivo knee cartilage. *Magn Reson Med*, 51(2):394–398, Feb 2004.

[180] Youngkyoo Jung, Alexey A Samsonov, Walter F Block, Mariana Lazar, Aiming Lu, Jing Liu, and Andrew L Alexander. 3D diffusion tensor MRI with isotropic resolution using a steady-state radial acquisition. *J Magn Reson Imaging*, 29(5):1175–1184, May 2009.

[181] R. L. O'Halloran, M. Aksoy, A. T. Van, and R. Bammer. 3D isotropic high-resolution diffusion-weighted MRI of the whole brain with a motion-corrected steady-state free precession sequence. *Magn Reson Med*, doi: 10.1002/mrm.24489, 2012.

[182] K Oshio, F A Jolesz, P S Melki, and R V Mulkern. T2-weighted thin-section imaging with the multislab three-dimensional RARE technique. *J Magn Reson Imaging*, 1(6):695–700, Nov-Dec 1991.

[183] M Engström, R Bammer, and S Skare. Sequence modifications and reconstruction strategies for multi-slab multi-echo DWI. In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 2301).

[184] Rob H N Tijssen, Thomas W Okell, and Karla L Miller. Real-time cardiac synchronization with fixed volume frame rate for reducing physiological instabilities in 3D FMRI. *Neuroimage*, 57(4):1364–1375, Aug 2011.

[185] P. Mansfield, P. R. Harvey, and M. K. Stehling. Echo-volumar imaging. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 2(3):291–294, Oct 1994.

[186] R Wirestam, D Greitz, C Thomsen, S Brockstedt, M B Olsson, and F Stahlberg. Theoretical and experimental evaluation of phase-dispersion effects caused by brain motion in diffusion and perfusion MR imaging. *J Magn Reson Imaging*, 6(2):348–355, Mar-Apr 1996.

[187] A Franck, D Greitz, B Nordell, and F Stahlberg. A theoretical study of amplitude modulation and time shifting in quantitative MR measurements of motion in brain tissue. *Magn Reson Imaging*, 11(5):739–747, 1993.

[188] 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*, 9(4):835–838, Jul-Aug 1985.

[189] M H Cho, W S Kim, and Z H Cho. CSF flow artifact reduction using cardiac cycle ordered phase-encoding method. *Magn Reson Imaging*, 8(4):395–405, 1990.

[190] 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*, 41(5):983–991, May 1999.

[191] David A. Feinberg and Koichi Oshio. Gradient-echo shifting in fast MRI techniques (GRASE imaging) for correction of field inhomogeneity errors and chemical shift. *Journal of Magnetic Resonance (1969)*, 97(1):177 – 183, 1992.

[192] D A Feinberg and K Oshio. Phase errors in multi-shot echo planar imaging. *Magn Reson Med*, 32(4):535–539, Oct 1994.

[193] H Gudbjartsson and S Patz. The Rician distribution of noisy MRI data. *Magn Reson Med*, 34(6):910–914, Dec 1995.

[194] David A Feinberg, Steen Moeller, S Smith, Edward Auerbach, Kamil Ugurbil, and Essa Yacoub. Multiplexed Echo Planar Imaging with Sub-second Whole Brain FMRI and Fast Diffusion Imaging. In *Proceedings of the 19th Annual Meeting of ISMRM*, Montreal, Canada, 2011 (abstract 635).

[195] R M Heidemann, A Anwander, Thorsten Feiweier, Cornelius Eichner, Ralf Lutzjendorf, Johannes Bernading, Thomas R Knosche, and Robert Turner. Sub-millimeter diffusion MRI at 7T: Does resolution matter? In *Proceedings of the 20th Annual Meeting of ISMRM*, Melbourne, Australia, 2012 (abstract 1877).