Motion Correction for MRI using External Tracking Devices

by

James A. Smith

Thesis submitted to

The University of Nottingham

for the degree of

Doctor of Philosophy

November 2018

Acknowledgements

The work in this thesis was only possible because of the help and support of many others.

I would like to thank my supervisors Penny Gowland and Paul Glover who have been exceedingly patient and very supportive. Olivier has been an excellent teacher and very generous in giving up his own time to help with the project, either calibrating the camera for me or teaching me how to pulse program. I want to also acknowledge the rest of the motion team: Laura Bischoff for the discussions we had when we were first figuring out the camera, Pueng for saving countless scanning sessions when things went wrong, Laura Bortolotti for always keeping a smile on my face, as well as everyone else who has been involved either directly or indirectly with the project, I'm sorry I can't name you all here.

I have called the Sir Peter Mansfield Imaging Centre my home for $>$ 4 years and I would like to thank all the students, academics and supporting staff who have made me feel welcome throughout my time there.

I would not have got this far in life had I not had constant support and kindness of my parents, Jill and Michael, even if my research is still a mystery to you. And I could not have survived 26 years of my parents without my sister, Charlotte, to back me up.

I also want to thank my friends for the moral support throughout my PhD. A special thank you to Ellie for 8 years of friendship and psychic conversations, Michael for his ability to set new MATLAB challenges for me and always keeping me on my toes and Ian for keeping me grounded and reminding me there is a life outside of work.

Abstract

Subject motion is a major problem for Magnetic Resonance Imaging (MRI), both clinically and in research. Typical MRI scans can last several minutes, during which substantial subject movements can occur. These movements negatively affect image quality resulting in loss of data or the need for repeated scans.

There are currently a large number of potential solutions to subject motion. Some of these solutions involve altering the scan to make it more resistant to subject motion, others involve tracking the subject and either retrospectively correcting data or prospectively updating the acquisition sequence to follow the subject as they move. The method by which the subject is tracked varies depending on the users needs; MR based methods such as navigators are often the simplest but may increase scan times, hardware based methods such as NMR probes or optical tracking are the most accurate methods but require additional set up times and are more expensive.

In this thesis the use of a Moiré phase tracking (MPT) camera is investigated for its use in prospective motion correction.

When using an external tracking camera, data is acquired in the camera's frame of reference and must be converted to the magnet's frame of reference (cross calibration) before it can be used for motion correction. This causes an additional source of error when determining the subject position. The impact of this error was investigated to determine the acceptable accuracy of the calibration transform. A novel method of measuring calibration stability is demonstrated to increase confidence in the calibration transform.

Using the calibration transform, a software patch (Geocam) was written to enable prospective motion correction on a Philips 7 T MRI scanner. This patch

was compared to existing motion correction techniques available on the scanner which correct for inter-volume motion. The new functionality (intra-volume motion correction) introduced by the Geocam patch was demonstrated with phantom and *in-vivo* data by performing large movements during a volume imaging sequence.

The Geocam patch was then demonstrated in combination with a novel method of determining B_0 maps to correct geometric distortions in EPI data. The proposed method was able to improve the registration of EPI data to an anatomical images compared to standard methods.

Despite the advantages of optical tracking cameras, they still suffer from significant limitations due to the requirement of line of sight between the marker and camera. Commercial optical tracking cameras are also relatively expensive. A small electronic sensor is proposed that could potentially address these limitations. The proposed sensor performed comparably to existing methodologies in terms of overall accuracy. Future work to enable to the sensor to be used for prospective motion correction is proposed.

This thesis addresses the practical implications of employing optical tracking camera technology for prospective motion correction at 7 T and proposes alternative devices which could be used.

Contents

1 Thesis Overview **8**

2 Principles of MRI **11**

2.1 Principles of NMR 12

2.1.1 Spin and Energy Levels 12

2.1.2 Excitation 16

2.1.3 Relaxation 18

2.1.4 Free Induction Decay 22

2.1.5 Echoes 23

2.1.6 Chemical Shift 26

2.1.7 Susceptibility 27

2.2 MR Imaging 27

2.2.1 K-space 28

2.2.2 Signal Localization 30

2.2.3 Imaging Sequences 34

2.2.4 Parallel Imaging 37

2.3 MRI System 38

2.3.1 Static Field 39

2.3.2 Shim Coils 40

2.3.3 Gradient Coils 41

2.3.4 RF Coils 42

3 Review of Motion Correction in MRI **43**

3.1 Motion and MRI 44

3.2 Motion by Anatomy 46

3.2.1 Movement in Body Imaging 46

3.2.2 Movement in Neuroimaging 49

3.3 Image Registration 50

3.4 Minimising Motion Effects 51
3.4.1 Restraint and Sedation 51
3.4.2 2D vs 3D Imaging 52
3.4.3 Gradient Motion Nulling 53
3.5 Correcting Motion 54
3.5.1 Retrospective Correction 54
3.5.2 Prospective Correction 55
3.5.3 Prospective vs Retrospective Correction 56
3.6 MR Tracking Methods 58
3.6.1 Navigators 58
3.6.2 PROPELLER 59
3.7 External Tracking Devices 60
3.7.1 RF Coils 61
3.7.2 Field Probes 61
3.7.3 Optical Tracking 62
3.7.4 Moiré Phase Tracking 64

4 Strategies for Cross Calibration 67
4.1 Introduction 68
4.2 Theory 73
4.3 Method 76
4.3.1 Determining Camera Calibration 76
4.3.2 Simulation of Camera Calibration Error 78
4.3.3 Calibration Error Propagation 80
4.3.4 Assessing Calibration Reliability 82
4.4 Results 83
4.4.1 Simulation of Camera Calibration Error 83
4.4.2 Calibration Error Propagation 86
4.4.3 Assessing Calibration Reliability 90
4.5 Discussion 93
4.5.1 Simulation of Camera Calibration Error 93
4.5.2 Calibration Error Propagation 95
4.5.3 Assessing Calibration Reliability 96
4.5.4 Practical considerations 97
4.6 Conclusion 99

5 Geocam Software Patch 100
5.1 Introduction . . . 101
5.2 Implementation . . . 103
5.2.1 Camera output . . . 103
5.2.2 Determining Image Volume Location . . . 104
5.2.3 Updating Scanner Geometry . . . 108
5.3 Method . . . 109
5.3.1 Demonstration of Geocam and Prospective Image Registration . . . 109
5.3.2 Determining Geocam Accuracy . . . 110
5.3.3 Intra-volume Motion Correction . . . 113
5.4 Results . . . 115
5.4.1 Demonstration of Geocam and Prospective Image Registration . . . 115
5.4.2 Determining Geocam Accuracy . . . 116
5.4.3 Intra-volume Motion Correction . . . 119
5.5 Discussion . . . 125
5.5.1 Demonstration of Geocam and Prospective Image Registration . . . 125
5.5.2 Determining Geocam Accuracy . . . 127
5.5.3 Intra-volume Motion Correction . . . 128
5.5.4 Geocam patch Limitations . . . 130
5.6 Conclusion . . . 132

6 EPI Geometric Distortion Correction 133
6.1 Introduction . . . 134
6.2 Theory . . . 137
6.2.1 Magnetic Susceptibility . . . 137
6.2.2 Geometric Distortion . . . 137
6.3 Method . . . 139
6.3.1 Modelling B_0 . . . 139
6.3.2 Geometric Distortion Correction . . . 140
6.4 Results . . . 144
6.4.1 Modelling B_0 . . . 144
6.4.2 Distortion Correction . . . 145
6.5 Discussion . . . 152
6.5.1 Modelling B_0 . . . 152

6.5.2 Distortion Correction . . . 153
6.5.3 Optimisation . . . 154
6.5.4 EPI Phase Artefact . . . 154
6.5.5 Alternative Correction Methods . . . 155
6.6 Conclusion . . . 156

7 Electromagnetic Position Sensor for MRI 157
7.1 Introduction . . . 158
7.2 Theory . . . 159
7.2.1 Determining Orientation . . . 159
7.2.2 Determining Position . . . 162
7.2.3 Determining Sensor Error . . . 165
7.3 Method . . . 166
7.3.1 Sensor Construction . . . 166
7.3.2 Controlling Movement . . . 168
7.3.3 Calibrating the Sensor Coil . . . 169
7.3.4 Sensor vs Optical Tracking . . . 170
7.4 Results . . . 172
7.4.1 Calibrating the Sensor Coil . . . 173
7.4.2 Sensor Vs Optical Tracking . . . 174
7.5 Discussion . . . 176
7.5.1 Calibrating the Sensor Coil . . . 176
7.5.2 Sensor vs Optical Tracking . . . 176
7.5.3 Further Work . . . 177
7.6 Conclusion . . . 180

8 Summary 181
8.1 Concluding Remarks . . . 182
8.2 Future Work . . . 184
8.2.1 Geocam Patch Improvements . . . 184
8.2.2 Marker Fixation . . . 184
8.2.3 Alternative Calibration Methods . . . 185
8.2.4 EPI Distortion Correction Applications . . . 186
8.2.5 Sensor Development . . . 186

A Sensor Circuit Schematics 201

Chapter 1

Thesis Overview

Magnetic Resonance imaging (MRI) has revolutionised modern day clinical imaging. However, it can take several minutes to acquire an image of diagnostic quality. It is required that during this time the subject remain stationary, however, this is often not possible, especially for clinical patients. Subject movement during acquisition can cause substantial image artefacts, limiting the diagnostic use of an image. Even with compliant subjects, the effective resolution of an image is limited by small subject movements. It is therefore crucial that methods are developed that can correct for subject movement during an MRI scan.

This thesis focuses on methods of motion correction for magnetic resonance imaging by tracking the subject movement. The primary methodology discussed in this thesis is a Moiré phase tracking (MPT) camera used to track subject motion.

Chapter 2 describes the physics of magnetic resonance imaging starting from the principles of nuclear magnetic resonance (NMR). This chapter goes on to discuss how NMR translates into MRI, methods of acquiring and reconstructing images and the hardware involved in generating the required magnetic fields.

Chapter 3 contains a literature review of existing motion correction methodologies and outlines some of the advantages of each technique. Both prospective and retrospective methods as well as general methods for reducing the impact of subject motion during imaging are discussed. There is a focus in this chapter on external tracking devices for neuroimaging although non-rigid body imaging is also discussed.

Chapter 4 considers the calibration and set up of an optical tracking camera for use in prospective motion correction. The focus of this chapter is the determination of the calibration transform (the transform that describes the position and orientation of the camera relative to iso-centre). An analysis of the calibration error is performed in simulation, both the occurrence of errors in the calibration and how those errors propagate. This chapter concludes by giving approximate values for the maximum recommended calibration error and proposes a novel method of identifying bad calibration transforms.

Chapter 5 contains details regarding a prospective motion correction patch

(Geocam) that was created for use on a 7 T Philips MRI scanner. It begins by describing the transformations required to use an optical camera to provide geometry updates for prospective motion correction. The use of the Geocam patch is compared to existing motion correction methods available on the scanner. The accuracy of the Geocam patch was determined experimentally by scanning a custom made phantom in the presense of large movements. The new functionality of the patch to correct for intra-volume motion is then demonstrated with phantom and in-vivo experiments.

Chapter 6 uses the Geocam patch to prospectively correction for subject motion during a typical EPI acquisition. The geometric distortions that arise due to B_0 inhomogeneity are corrected for using a novel method of generating predicted B_0 maps using linear regression. The corrected EPI images are then registered to other EPI and anatomical images to measure the degree to which the images were corrected.

Chapter 7 describes a new method of measuring subject motion through the use of a small electronic sensor. This sensor comprises of an accelerometer and an anisotropic magneto resistive sensor to determine the direction of known field vectors. These vectors are used to determine the sensor's orientation within the MRI scanner. A small pick-up coil is used to detect gradient pre-pulses generated by the scanner, which are then used to determine the sensor's position within the scanner. The sensor is then compared to an optical tracking camera and evaluated for accuracy.

Chapter 8 contains a summary of the results in the thesis and a discussion on potential future work.

Chapter 2

Principles of MRI

2.1 Principles of NMR

2.1.1 Spin and Energy Levels

Magnetic resonance imaging (MRI) uses the principles of Nuclear Magnetic Resonance (NMR) as demonstrated by Felix Bloch and Edward Purcell in separate experiments [1, 2] which described the interaction between nuclear spin and external magnetic fields. While NMR spectroscopy uses the NMR signal to probe the atomic envirnoment of molecules, MRI spatially encodes this same signal to generate images. The use of this phenonomen for medical imaging was first shown by Peter Mansfield [3] and seperately by Paul Lauterbur [4].

NMR relies on the atomic property of nuclei known as spin ($\mathbf{I}$ for nuclear spin or s for electron spin) which is a manifestation of angular momentum ($\mathbf{L}$) at the quantum level. Spin is a fundamental property of sub-atomic particles much in the same way charge and mass are. Angular momentum gives rise to the magnetic moment ($\boldsymbol{\mu}$) as a result of a spinning charge. The relationship between angular momentum and magnetic moment is given by,

$$\boldsymbol{\mu} = \mathbf{I}\gamma, \tag{2.1}$$

where γ is the gyromagnetic ratio of the nuclei. The magnetic moment is a description of the magnetic field and orientation generated by a spinning charge. The net magnetisation of a sample can be described as the sum of all magnetic moments in the sample.

The most common nuclei used in MRI is hydrogen (proton MRI), this is because of its high gyromagnetic ratio (as seen in Table 2.1) and the natural abundance of hydrogen (in the form of water) within the body, which means proton MRI gives the highest signal-to-noise ratio compared to other nuclei present within the body.

The spin number of a nucleus is determined by the sub-atomic ratio of hadrons that make it up. Nuclei containing an even number of protons and neutrons have a spin of zero ($I = 0$). Nuclei containing an odd number of protons and neutrons

Nucleus	$\gamma/2\pi$ (MHz T^{-1})
^{1}H	42.577
^{3}He	-32.434
^{13}C	10.705
^{15}N	-4.316
^{17}O	-5.772
^{19}F	40.052
^{23}Na	11.262
^{31}P	17.235

Table 2.1: A table of gyromagnetic ratios of different nuclei [5]

have integer spins ($I = 1, 2, 3$... etc) and most other nuclei have half integer spins ($I = \frac{1}{2}, \frac{3}{2}, \frac{5}{2}$... etc).

Because the spin of a nuclei is quantised, so too is the angular momentum ($\mathbf{p}$). Angular momentum is given as,

$$\mathbf{p} = \hbar\sqrt{I(I+1)}, \tag{2.2}$$

where $\hbar$ is the reduced ($\frac{h}{2\pi}$) . The z-component of $\mathbf{p}$ is described as,

$$p_z = \hbar m_I, \tag{2.3}$$

where m_I can take values $-I, -I+1, ..., I-1, I$. For protons, m_I has two possible values $m_I = \frac{1}{2}$ or $-\frac{1}{2}$.

In the absence of an external magnetic field, the possible spin states are degenerate (have the same energy). However, in the presence of an external magnetic field the energy levels diverge for each of the spin states; this is known as the Zeeman effect [6]. For protons, which have two possible spin states, the spins will either align with the magnetic field (parallel) or against the magnetic field (anti - parallel). By convention the magnetic field is described as being in the z direction, and so the energy levels of each spin state can be described using:

$$E = -\mu_z B_0, \tag{2.4}$$

where μ_z is the z-component of the magnetic moment which is described as,

$$\mu_z = \gamma \hbar m_z, \tag{2.5}$$

with energy levels corresponding to nuclear spins that are aligned with (low E) or against (high E) the field. The energy difference between these spin states is related to the gyromagnetic ratio and the applied magnetic field, given the relationship,

$$\Delta E = \gamma \hbar B_0. \tag{2.6}$$

Since ΔE is small, the population of each spin state within an external magnetic field is nearly identical at room temperature. The ratio between the spin states can be described using the Boltzmann distribution:

$$\frac{N_\uparrow}{N_\downarrow} = e^{-\frac{\Delta E}{k_B T}}, \tag{2.7}$$

where k_B is the Boltzmann constant, T is the temperature (in Kelvin) and $N_\uparrow$ and $N_\downarrow$ are the populations in the spin up and spin down spin states respectively. Equation 2.7 shows that the population difference can be increased by decreasing the temperature (not recommended for *in-vivo* experiments), increasing B_0 (an expensive option which involves changing the large static magnet) or scanning nuclei with higher gyromagnetic ratios (the increase in ΔE may be outweighed by reduced abundance of the nuclei in-vivo however). Because the population difference between spin-up and spin-down nuclei is so small in MRI we consider a relatively large sample of spins in a volume and measure the net magnetisation within, which is described as,

$$M_0 = \frac{\gamma^2 \hbar^2 B_0 N_s}{4 k_B T}, \tag{2.8}$$

where N_s is the total number of spins in the system. For biological tissue there are approximately 10^{22} protons per cm^3.

A spin in an external magnetic field will experience a torque due to its magnetic

moment. Classically torque arises from the conservation of angular momentum as described by Newton's second law of motion:

$$\frac{d\mathbf{I}}{dt} = \mathbf{T}, \tag{2.9}$$

where $\mathbf{T}$ is torque and $\frac{d\mathbf{I}}{dt}$ is the change in angular moment over time. The torque experienced by a spin in an external field can be described as,

$$\mathbf{T} = \boldsymbol{\mu} \times \mathbf{B_0}, \tag{2.10}$$

For a large collection of spins the magnetic moment can be substituted for magnetisation $\mathbf{M}$. Equation 2.10 can then be written in terms of the change of angular momentum and hence the change of magnetisation:

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B_0}, \tag{2.11}$$

The magnetisation can then be decoupled into its vector components ($\mathbf{M} = \hat{\mathbf{x}}M_x + \hat{\mathbf{y}}M_y + \hat{\mathbf{z}}M_z$) where the applied field is along the z axis,

$$\begin{aligned} M_x(t) &= M_{0x}\cos(\omega t) + M_{0y}\sin(\omega t), \\ M_y(t) &= M_{0y}\cos(\omega t) + M_{0x}\sin(\omega t), \\ M_z(t) &= M_0, \end{aligned} \tag{2.12}$$

Equation 2.12 shows that the magnetisation is not solely aligned with $\mathbf{B_0}$, that is to say, the magnetic moment of a spin can not be entirely described by its z-component alone. Instead the nuclear magnetic moment *precesses* around the z axis as seen in Figure 2.1.

The rate of precession is dictated by the energy difference between the spin states.

A transition between high and low energy states can be induced if electromagnetic (EM) radiation with this energy difference is applied (at the resonant

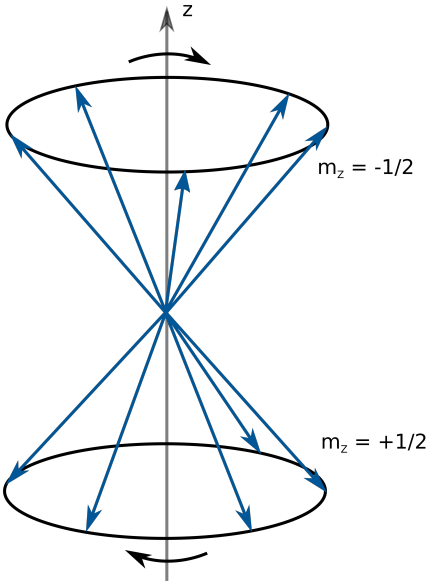


Figure 2.1: An diagram showing how the magnetic moment of an ensemble of spins (blue) precesses about the z axis.

frequency). The energy of a photon is related to its frequency (ω) by,

$$E = \hbar\omega = hf \tag{2.13}$$

where h is Planck's constant. Setting E to energy difference we get the expression:

$$\omega_L = \gamma B_0 \tag{2.14}$$

where $\gamma_{(Hz)}$ is expressed in HzT^{-1} and $\gamma_{(rads^{-1})}$ is expressed as $rads^{-1}T^{-1}$. This frequency is known as the Larmor frequency. At typical MRI field strengths the Larmor frequency is in the radio frequency part of the spectrum.

2.1.2 Excitation

In section 2.1.1 we introduced how an applied magnetic field oscillating at the Larmor frequency induces a transition between energy states. This radio-frequency wave is a time varying electromagnetic wave. In this sense there are two magnetic fields associated with MRI, the static field (B_0) and the oscillating magnetic

field (B_1), also known as the RF field because of its frequency band. This RF field has two effects on the collection of spins. Firstly because the RF field is at the resonant frequency of the spin states, all the spins come into phase. Secondly, the RF field effectively tips the bulk magnetisation vector ($\mathbf{M}$) from pointing along the z-axis, to oscillating in the x-y plane, creating a spiral-like motion (Figure 2.2A). To simplify this description, it is often beneficial to describe a new frame of reference that rotates about the x-y plane at the Larmor frequency, referred to as the rotating frame (as opposed to the stationary "laboratory frame"). If co-ordinates in the laboratory frame are denoted as x, y, z then the corresponding co-ordinates in the rotating frame are denoted at x', y', z'. A diagram showing the net magnetisation for a system of precessing spins in the presence of an RF pulse is shown in Figure 2.2. In the rotating frame, B_1 appears stationary and B_0 acts along the z direction. The z' component of the magnetisation is known as the longitudinal magnetisation while the x' y' component of the magnetisation is known as the transverse magnetisation.

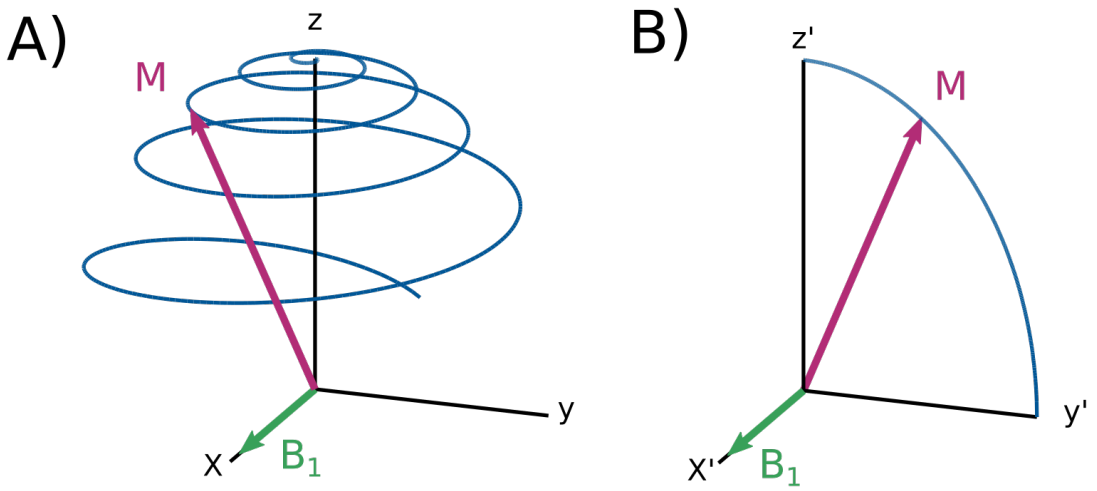


Figure 2.2: An diagram of the net magnetisation in the laboratory frame (A) and rotating frame (B). The net magnetisation ($\mathbf{M}$) is shown in pink, and the B_1 shown in green

The angle by which an RF pulse tips the net magnetisation is known as the

flip angle (α) and is described as,

$$\alpha = \int_0^\tau \gamma B_1(t)dt, \tag{2.15}$$

τ is the duration of the RF field. Because τ is relatively short the application of the RF field is often referred to as the RF pulse. In MRI this process is referred to as excitation. Additional RF pulses can be used in some imaging sequences to refocus the net magnetisation (spin echo, Section 2.1.5) or invert it.

After the application of an RF pulse the spins are coherent (all in the same phase state) creating a single net magnetisation vector of the spin system which is rotating about the transverse plane.

2.1.3 Relaxation

After excitation, spins begin to decay back to their thermal equilibrium position. There are two decay processes through which relaxation occurs; longitudinal relaxation is the process by which the net magnetisation (M_0) recovers back along the z axis. During transverse relaxation, the individual spins that make up the net magnetisation dephase (lose coherence) . These relaxation mechanisms form the basis of image contrast.

Longitudinal Relaxation

The relaxation of M_z towards the z-axis is caused by the exchange of energy (gained from the RF pulse) to the surrounding molecules. This energy transfer occurs because of the interactions with the surrounding molecules' magnetic field. This magnetic field experienced by a spin is randomly fluctuating due to Brownian motion. The interaction between spins is called dipole-dipole interaction, and is dependent on how well coupled a spin is to its surrounding molecules or "lattice". The time constant of this energy exchange is defined as T_1. The recovery of the net magnetisation due to longitudinal relaxation can be described in the Bloch

Equation:

$$\frac{dM_z(t)}{dt} = \frac{M_0 - M_z(t)}{T_1}, \tag{2.16}$$

and its solution:

$$M_z = M_0(1 - (1 - M_z(0)) \exp^{-\frac{t}{T_1}}). \tag{2.17}$$

Typical T_1 values for white and grey matter at 7 T are 1130 ms and 1940 ms respectively [7]. The longitudinal relaxation for two different tissue types can be seen in Figure 2.3.

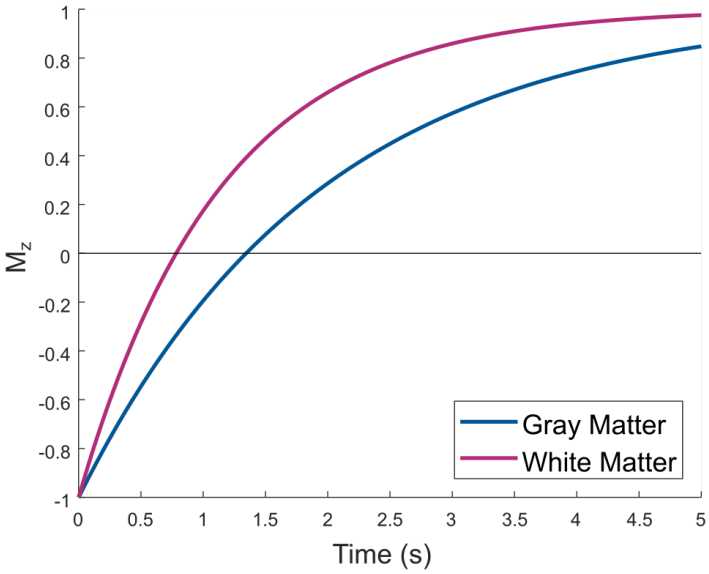


Figure 2.3: The longitudinal magnetisation recovery due to T_1 relaxation for white and gray matter at 7 T following a 180° pulse.

Because the longitudinal relaxation is dependent on the coupling of spins to the surrounding molecules, T_1 varies between tissues. By choosing appropriate times to sample the NMR signal, image contrast based on the T_1 value of the tissue can be obtained.

Transverse Relaxation

Transverse relaxation is caused by the dephasing of spins that make up the net magnetisation. Each spin that has been excited by a 90° RF pulse will precess about the xy plane at the Larmor frequency. Each spin experiences a varying magnetic fields arising from its surrounding neighbours, since the surrounding molecules each generate small magnetic fields. These magnetic fields change the field experienced by the spin which in turn changes its resonant frequency slightly. The small difference in resonant frequency between each spin causes phase differences over time.

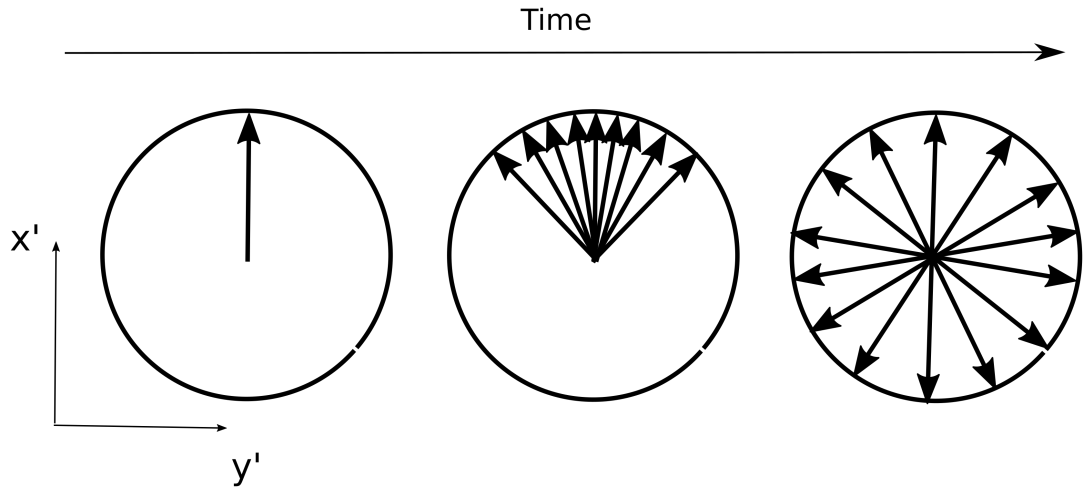


Figure 2.4: A diagram showing the dephasing of spins in the rotating frame. Immediately after the RF pulse all the spins that make up the net magnetisation point in the same direction. At some time later, due to small fluctuations in frequency, spins lose phase coherence

The rate at which the spins dephase is described by T_2. The resulting loss of magnetisation because of the transverse relaxation can be described in the Bloch equation :

$$\frac{dM_{xy}(t)}{dt} = -\frac{M_{xy}(t)}{T_2}, \tag{2.18}$$

and its solution:

$$M_{xy} = M_0 \exp^{-\frac{t}{T_2}} . \tag{2.19}$$

T_2 is always shorter than T_1 relaxation. Typical T_2 values for white and grey matter at 7 T are 46 ms and 55 ms respectively [8]. The value of T2 however can vary greatly due to factors such as diffusion, which causes water (and hence protons) to move during imaging. The transverse relaxation for two different tissue types can be seen in Figure 2.5.

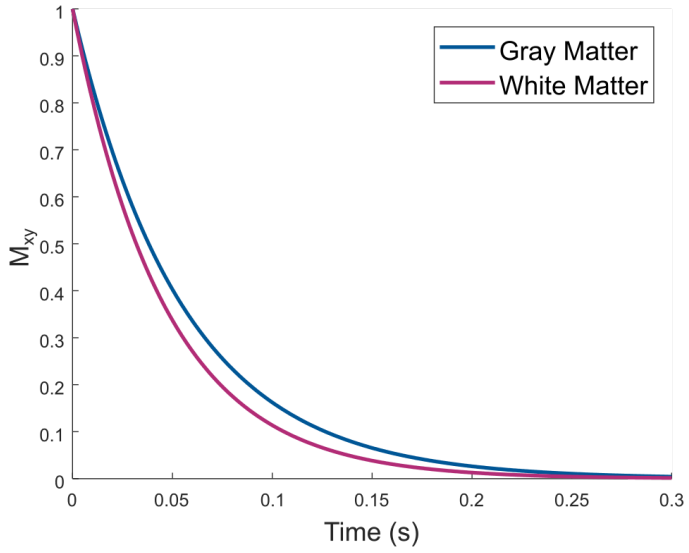


Figure 2.5: The transverse magnetisation loss due to T_2 relaxation for white and gray matter at 7 T

Experimentally, transverse relaxation occurs at a faster rate than can be explained solely by T_2. This is because in addition to the random field fluctuations from neighbouring molecules there are also larger field fluctuations from inhomogeneities in the local field strength. These field fluctuations are usually a result of the susceptibility of tissue to become magnetised in response to an external magnetic field. This magnetisation causes the field to be perturbed locally. These field fluctuations increase the transverse magnetisation relaxation rate.

The rate of relaxation which considers the contribution due to tissue suscepti-

bility is T_2* is described as,

$$\frac{1}{T_2*} = \frac{1}{T_2} + \frac{1}{T_2'}, \tag{2.20}$$

where T_2' is a variable describing the rate of decay due to perturbations in B_0.

T_2* is always shorter than or equal to T_2. The contribution from T_2' is higher at higher field strengths due to the increased inhomogeneity that is present at higher fields. Typical T_2* values for white matter and gray matter at 7 T are 27 ms and 33 ms respectively [9]. Images which are weighted by T_2 rather than T_2* can be obtained by using a spin echo sequence (Section 2.1.5) which reverses the phase loss from tissue susceptibility.

The contribution from all the relaxation parameters can be combined to give the phenomenological Bloch equation.

$$\frac{d\mathbf{M}}{dt} = \gamma(\mathbf{M} \times \mathbf{B}) - \frac{(M_z - M_0)\hat{\mathbf{z}}}{T_1} - \frac{M_x\hat{\mathbf{x}} + M_y\hat{\mathbf{y}}}{T_2} \tag{2.21}$$

where $\mathbf{B}$ is the combined external magnetic field as a result of $\mathbf{B_0}$ and $\mathbf{B_1}$.

2.1.4 Free Induction Decay

A 90° RF pulse converts the longitudinal magnetisation to transverse magnetisation. Faraday's law of induction states that a rotating magnet will produce an electromotive force (EMF) in nearby conductive coils. Because of this the transverse magnetisation, which is rotating about the z axis in the laboratory frame, will cause currents to be induced in nearby receiver coils aligned transversally. These signals will have the same frequency as the rate of precession of the spins (in this case the Larmor frequency). However, due to the T_2* relaxation mechanisms described in Section 2.1.3, the signal will also decay exponentially. This measurable signal is called the Free Induction Decay (FID) an example of which is shown in Figure 2.6.

The FID decays primarily due to the effects of T_2*

Measuring the FID is the most basic NMR experiment but can yield valuable data about T_2*. Despite being the simplest process it is often difficult to measure

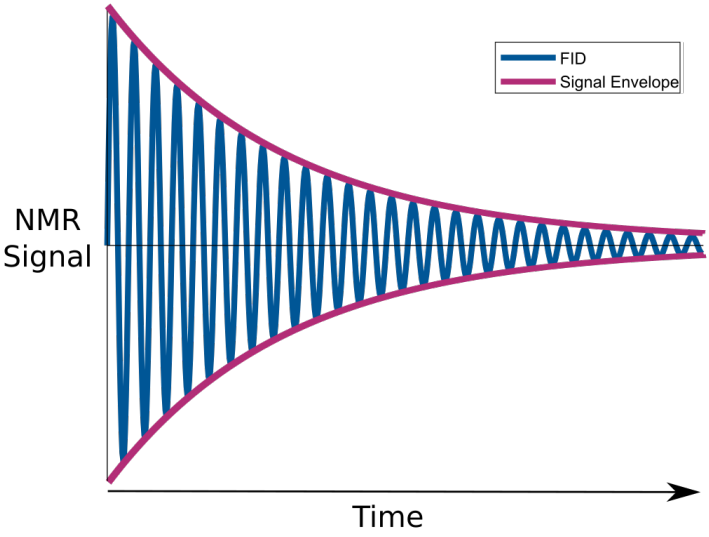


Figure 2.6: The initial NMR signal after an RF pulse, known as the free induction decay, decays with T_2 assuming no field inhomogeneities.

an FID. The RF received channels are gated to avoid being saturated by the transmitted RF pulse. Because of this the early FID is often difficult to measure, this is known as the "deadtime" of the system (typically $<$1 ms).

2.1.5 Echoes

In most MR imaging the FID is not used to sample the NMR signal, instead the NMR signal is measured from an Echo. An echo is a manifestation of the FID that is caused by the rephasing of NMR spins [10]. The FID rapidly decays primarily due to T_2* from the combined effect of random local field variations, B_0 field inhomogeneities and spatially varying gradient fields. Echoes are generated by applying some mechanism to refocus the phase e.g gradient or spin echo.

Spin Echo

For a spin echo, an additional 180° RF pulse is applied at time τ after the excitation pulse to flip the magnetisation in the x-y plane. The phase gained by spins precessing faster than the Larmor frequency (due to spatial fluctuations in B_0) is now undone as the precession is now reversed and the increased frequency acts

in the opposite direction reducing the phase difference between spins. The same applies for spins that previously had lower precession rates. Just as phase coherence is lost exponentially after the excitation pulse, phase coherence is now gained after the inversion pulse, exponentially increasing the amplitude of the NMR signal. The peak signal of the echo occurs at time 2τ also known as the echo time (TE). A diagram showing the dephasing and rephasing of spins can be seen in Figure 2.7.

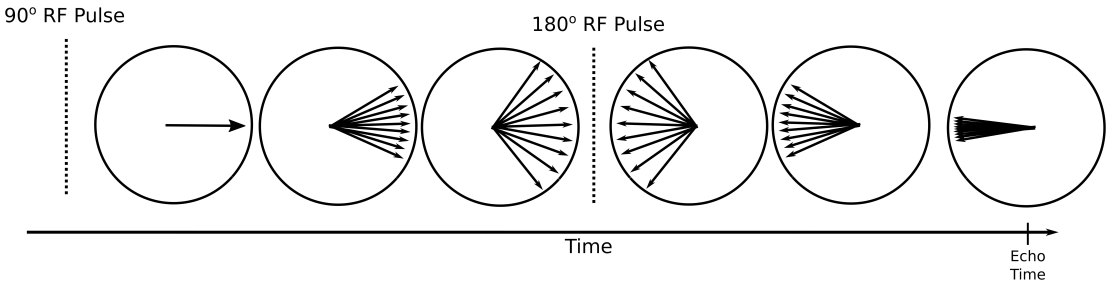


Figure 2.7: A diagram showing a collection of spins dephasing after a 90° RF pulse. The spins begin to rephase after a 180° RF pulse, with the echo measured at the echo time. The spins do not entirely rephase due to irrevesible phase loss from T_2

The amplitude of the NMR signal is limited by the transverse decay that occurs due to T_2. The phase lost as a result of T_2 can not be regained because it was lost by random, time varying, unpredictable interactions with neighbouring molecules. A simple spin echo sequence can be seen in Figure 2.8.

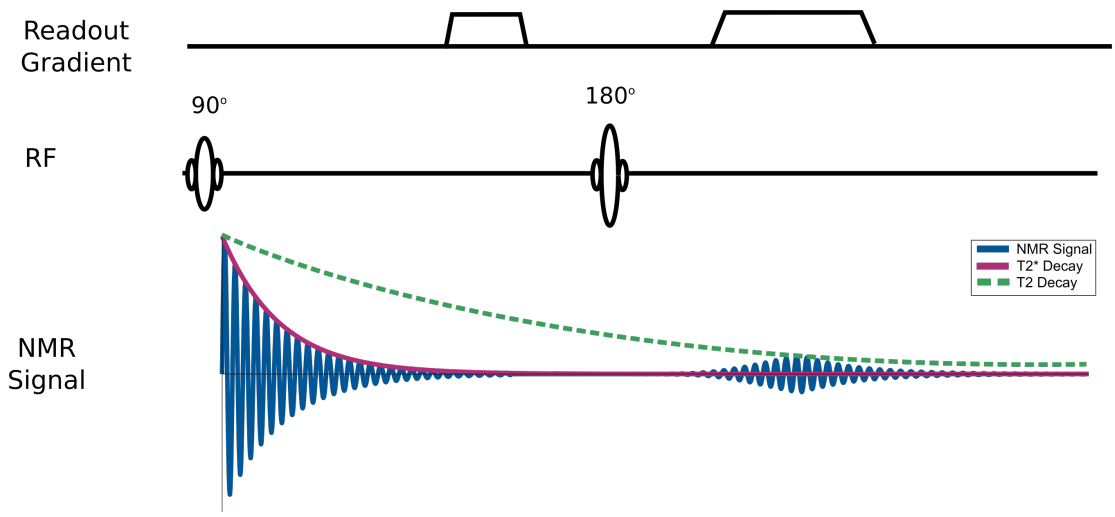


Figure 2.8: A simple spin echo sequence where the second 180 degree RF pulse acts to refocus the phase that is lost due to field inhomogeneities.

Gradient Echo

For a gradient echo, after the excitation pulse, a negative gradient is applied across the sample [11]. This dephasing gradient causes spins to precess at different speeds, and because of this dephase faster. At time $TE/2$, a positive rephasing gradient is applied across the sample. This causes the spins to rephase and hence the signal to return with the peak amplitude of the signal occurring at time TE. Unlike spin echo imaging, this refocusing will only affect the spins that have dephased due to the initial dephasing gradient; The NMR signal will still decay with T_2* . Because only one RF pulse is needed, gradient echoes can be acquired with much shorter echo times. The rephasing gradient is typically twice the length of the dephasing gradient, to ensure that the peak of the echo occurs at the the centre of k-space (the dephasing gradient moves it through k-space by $-x$ then after the first half of the rephasing gradient it is shifted back by x). A typical gradient echo sequence is shown in Figure 2.9 and is discussed further in Section 2.2.3. Gradient echoes are also known as gradient recalled echoes and field echoes.

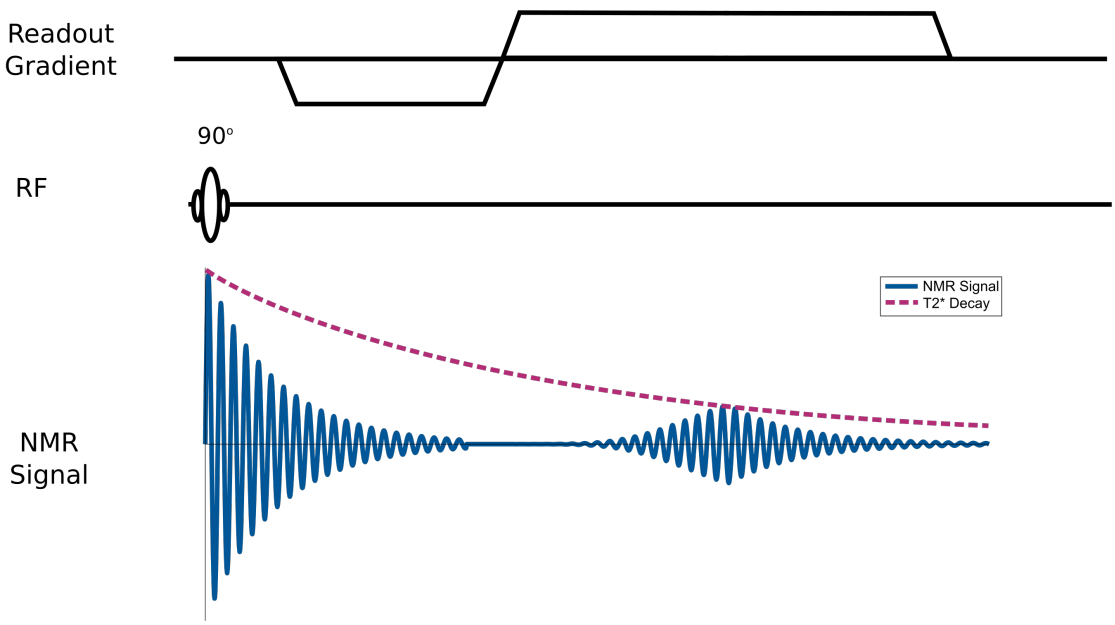


Figure 2.9: A simple gradient echo sequence where the gradients are used to deliberately dephase and rephase the signal to form a gradient echo.

2.1.6 Chemical Shift

It has been assumed thus far that all protons resonate at the same frequency with the exception of subtle variations due to inhomogeneity in the local field. The resonant frequency of protons can also vary due to the electronic shielding of its local chemical environment, which changes the effective B_0 field experienced by the proton. This is known as chemical shift. The effective field experienced by a spin is given by Equation 2.22.

$$B_{eff} = (1 - \sigma)B_0 \tag{2.22}$$

where σ is the chemical shift shielding constant.

This is the principle behind NMR spectroscopy whereby the free induction decay (Section 2.1.4) is analysed. The slight difference in resonance frequency caused by the chemical shift can be probed to give information about a spin's chemical environment.

2.1.7 Susceptibility

Susceptibility (χ) is a measure of a material's response to an external magnetic field. Materials placed into a field may become polarised and develop magnetisation. The induced magnetisation (M_χ) can be described in terms of the susceptibility of the material in Equation 2.23.

$$M_\chi = \frac{\chi}{\mu_0} B_0 \tag{2.23}$$

where μ_0 is the permeability of free space. The resulting field perturbation due to this induced magnetisation is described in Equation 2.24

$$\Delta B_0 \propto \chi B_0 \tag{2.24}$$

Equation 2.24 shows that at higher field strengths the inhomogeneity of the static field due to susceptibility increases. This change in magnetic field due to susceptibility contributes to the transverse relaxation which causes T_2* to be shorter than T_2.

Materials can be split into two categories depending on the sign of χ. Materials with $\chi > 0$ are known as paramagnetic while materials with $\chi < 0$ are referred to as diamagnetic. The effect of magnetic susceptibility on echo planar imaging is discussed in Chapter 6.

2.2 MR Imaging

As mentioned at the begining of this chapter MRI uses the principles of NMR to create images. The earliest NMR images used a form of back projection to create a 2D image. *Lauterbur et al.* [4] used a linear gradient field applied at various angles and sampled the signal generated at each angle. Each signal was then back projected to generate a 2D image. Shortly after, selective excitation was proposed by *Garroway et al*. [12], and then later *Ernst et al.* demonstrated the first Fourier imaging method [13] referred by *Bottomley* as spin warp imaging [14].

The principle behind turning an NMR signal into an image involves two stages: encoding the signal so that it is spatially dependent, and reconstructing the signal from spatial frequency to image space.

2.2.1 K-space

Images in MRI are acquired by measuring the frequency and phase of NMR signals. Signals acquired in this space are described in k-space [15–17]. k-space (k_x, k_y, k_z) is a multi dimensional manifestation of the spatial frequencies that make up an image (x, y, z), in an identical way that a 1D signal can be decomposed into its frequency components as shown by Joseph Fourier in 1822. Because of this k-space is often referred to as Fourier space.

NMR signals are detected as spatially encoded (see section 2.2.2) signals that correspond to k-space. k-space is then transformed into image space using a discrete Fourier transform. This transform can be either 2D or 3D depending on the encoding of the NMR signal.

The route by which k-space is sampled is referred to as the k-space trajectory. Each point in k-space is responsible for a sinusoidal wave across the image in image space. The k-space data is complex so that it contains both a magnitude (the frequency of the sine wave) and a phase (the phase of the sine wave). The image is given as the integral of all the frequencies defined in k-space. Because of this, every point in k-space contributes to the entire image in image space. Figure 2.10 shows how a single point (and its complex conjugate pair) generates an wave pattern across the entire image in image space.

Since images are effectively acquired in k-space, it is important to understand how properties of k-space affect image space. The resolution of an image is defined by the highest frequency components of k-space (the highest k_x, k_y, k_z sampled). While the k-space trajectory may be continuous, the k-space data is sampled discretely along the trajectory; the spacing between the sampling points defines the field of view (FOV) of the image by Nyquist theorem. In this sense the field of view and resolution between image space and Fourier space are symmet-

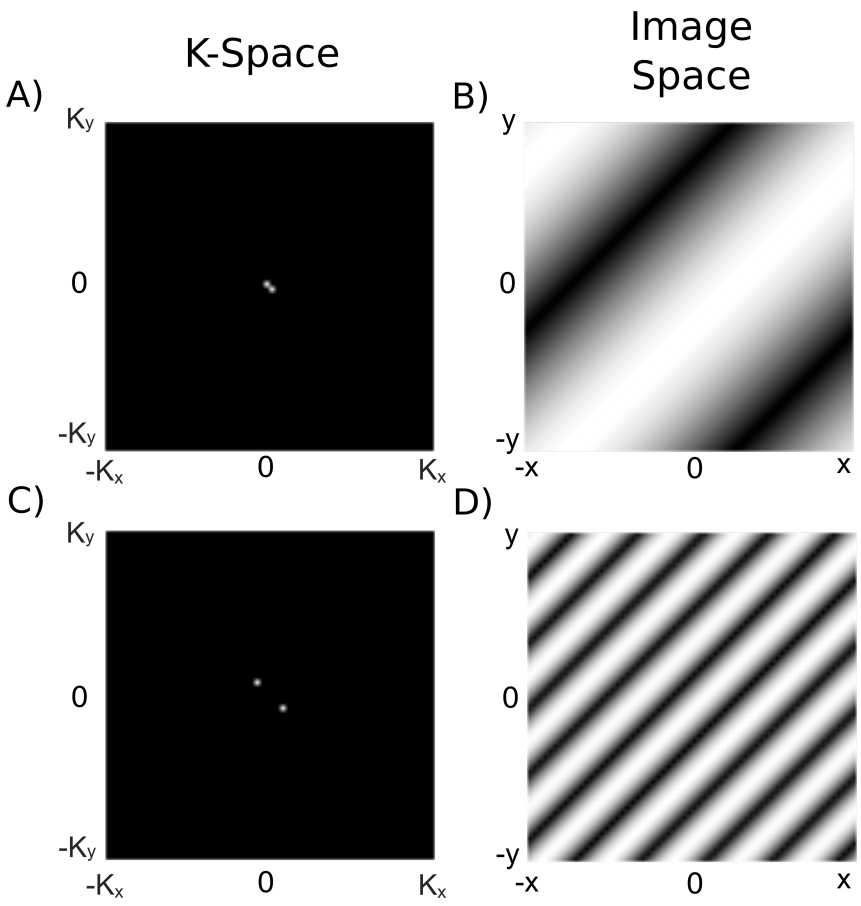


Figure 2.10: An example of the conversion from Fourier space to image space. B) is the Fourier transform of A) and D) is the Fourier transform of C). In A) a single data point (and its conjugate pair) are located close to the centre, generating low spatial frequencies in image space. In C) the data point is located further out in k-space generates higher spatial frequencies in image space.

rically linked: the resolution of k-space defines the field of view of the image, and the resolution of the image is defined by the field of view of k-space.

k-space contains low spatial frequencies in the centre and higher spatial frequencies at the extremities, i.e. the coarse image data is contained in the centre of k-space, which is also where the bulk of the information is contained. The high spatial details are towards the edges of k-space as seen in figure 2.11

In theory, only half of k-space needs to be acquired because of the inherent symmetry in k-space where one half is the complex conjugate of the other. On Philips scanners this option is known as “half scan”, where the user can elect to

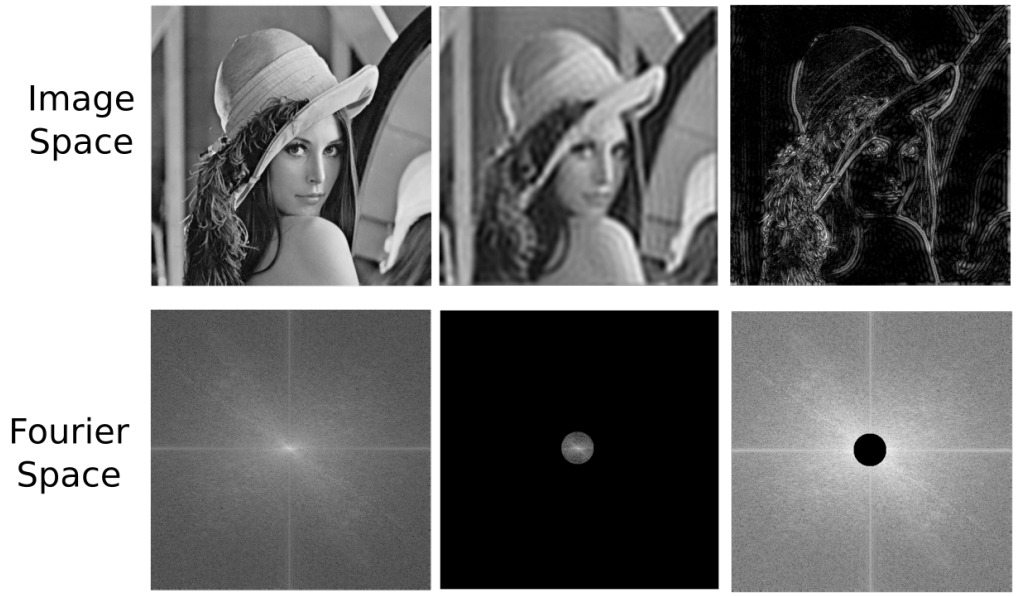


Figure 2.11: An example of image space and the corresponding k-space. Filters have been applied to k-space that either include or exclude low frequency components, images are then produced by performing a 2D fast Fourier transform.

acquire only half of the total k-space (with the rest determined by determining the complex conjugate). However, more than 50% of the k-space is acquired in order to determine the phase at the centre of k-space.

2.2.2 Signal Localization

Three directions need to be considered during the acquisition of a 2D MRI image; the readout direction (conventionally denoted as x/k_x) the phase encoding direction (conventionally denoted as y/k_y) and the slice direction (z). Spatial encoding is achieved though the use of gradients which are applied throughout an imaging sequence. All gradients generate a change in the z component of the magnetic field. However, the strength of the gradient varies depending on the direction in which the the gradient is applied (x, y, or z),

$$G_x = \frac{dB_z}{dx}, G_y = \frac{dB_z}{dy}, G_z = \frac{dB_z}{dz}, \tag{2.25}$$

the local magnetic field experienced by a spin during a particular gradient is then:

$$B_z(r) = B_0 + G_r r, \tag{2.26}$$

where r is some arbitrary position within the gradient field. This change in the magnetic field causes a change in the frequency of precession of a spin (Equation 2.14). By applying gradients at given periods and different directions during the sequence, the change in precession rate can be used to encode spatial information.

Slice Selection

Earlier it was stated that an RF pulse at the Larmor frequency is needed to excite a spin; it is this principle that is used to spatially encode a slice. In the absence of any gradients, an RF pulse excites the entire volume, in MRI this is known as a non-selective RF pulse. During a slice selective pulse, a gradient is applied across the volume at the same time. The gradient changes the Larmor frequency of the spins along this direction and so only those spins that are at a frequency contained within the RF pulse will be excited. The position along the gradient of the excited slice can be changed by altering the central frequency of the RF pulse (the carrier frequency). By convention, the slice encoding direction is referred to as the z direction.

In reality, an RF pulse does not contain a single frequency, but a range of frequencies specified by its bandwidth. This is one method by which the width of the excited slice can be altered (increasing the bandwidth increases the number of frequencies within the excitation range and hence gives a thicker slice). Alternatively if a steeper gradient is used, the region of spins containing frequencies within the RF bandwidth is smaller and hence makes the slice thinner. The width of a slice (Δz) excited using a slice selective RF pulse of bandwidth $\Delta\omega$ is described as,

$$\Delta z = \frac{\Delta\omega}{\gamma G_z}. \tag{2.27}$$

In the low flip angle approximation slice profile can be described as the Fourier transform of the RF envelope. To create a perfect rectangular profile of excitation corresponding to perfect excitation of all spins within with bandwidth, an infinitely long sinc pulse would be required.

The gradient applied during the slice excitation RF pulse results in a change of frequency across the excitation band, however, over time this difference in frequency results in a change in phase, so that following the RF excitation there is a difference in phase among the excited spins. To counteract this a negative gradient is applied after the pulse which acts to reverse the phase accumulated by the spins during the initial gradient. The initial gradient acts to scan on k-space from the lowest frequency in the bandwidth, to the highest frequency within the bandwidth. To recentre k-space the refocusing gradient is half the length of the initial gradient to move back to k_0 .

Frequency Encoding

Much like slice selection, Frequency encoding also depends on changing frequencies due to the presence of magnetic field gradient, however in this case this gradient is applied during the readout stage of the MRI sequence to spatially vary the frequency of the signal.

In the absence of relaxation the signal from an isochromat (collection of spins which resonate at the same frequency) is given by:

$$S(\mathbf{r},t) \propto \rho(\mathbf{r}) \exp^{i\int_0^t \omega_L(\mathbf{r},t')dt'}, \tag{2.28}$$

where ρ is the proton density, $\mathbf{r}$ is a position within the magnet, t is a point in time at which the signal is sampled, ω_L is the Larmor frequency and t' represents the time of the evolution of the Larmor frequency for $t' = 0$ to $t' = t$.

The total signal measured will be contributed to by all isochromats within the excited volume,

$$S(t) \propto \int_V \rho(\mathbf{r}) e^{i\int_0^t \omega_L(\mathbf{r},t')dt'} dV, \tag{2.29}$$

where V is the excited volume and $dV = dxdydz$. Substituting in Equations 2.26 and 2.14 we can describe the signal in response to a gradient field as

$$S(t) \propto \int_V \rho(\mathbf{r}) e^{i \int_0^t G_r r dt'} dV, \tag{2.30}$$

The phase evolution $\mathbf{k}(t)$ in k-space due to the gradient is defined as ,

$$\mathbf{k}(t) = \gamma \int_0^t \mathbf{G}(t') dt', \tag{2.31}$$

substituting Equation 2.31 into Equation 2.30 gives:

$$S(t) \propto \int_V \rho(\mathbf{r}) \exp^{i\mathbf{r}\cdot\mathbf{k}(t)} dV, \tag{2.32}$$

A gradient applied in the x direction then gives a signal that is proportional to:

$$S(t) \propto \int_V \rho(\mathbf{r}) \exp^{ixk_x(t)} dV, \tag{2.33}$$

Assuming only a thin slice is excited during slice selection. Equation 2.32 can then be described as a Fourier transform:

$$S(t) \propto \int \mathcal{F}_x[\rho(r)] dy, \tag{2.34}$$

Hence by performing an inverse Fourier transform of the signal information about the spin density along the direction of the gradient (x) can be obtained. However because this is a 1D Fourier transform there is no information about the spin density along the y direction.

Phase Encoding

Phase encoding allows one to encode the 2nd dimension of a slice by creating a signal whose phase varies linearly across the sample. By convention this direction is known as the y direction.

A gradient is applied after the RF pulse but before the readout gradient. This gradient acts to change the frequencies of the spins in a similar way to the slice

selection gradient, however in this case no rephasing gradient is applied, and thus the spins acquire a phase dependent on the strength and time they are exposed to the phase encoding gradient. The amount of phase acquired by an isochromat can be described as,

$$\Phi_y = \gamma y G_y t_\phi, \tag{2.35}$$

where t_ϕ is the time the isochromat is exposed to the G_y gradient. The signal acquired from between $t' = t_\phi$ and $t' = t_\psi$ after exposure to a phase encoding gradient and during a readout gradient is:

$$S(t) \propto \int\int \rho(\mathbf{r}) \exp^{i\Phi_y - ix\mathbf{k}_x(t)} \, dxdy, \tag{2.36}$$

or :

$$S(t) \propto \int \mathcal{F}_x y[\rho(r)], \tag{2.37}$$

where $\mathcal{F}_x y$ is the 2D Fourier transform along x and y.

It may be necessary to rephase the spins after readout if the same slice is expected to be imaged again shortly after (in say an fMRI experiment). In this case the phase encoding may be reversed by applying a negative phase encoding gradient for the same duration as the forward phase encoding gradient, after readout.

2.2.3 Imaging Sequences

Many different imaging sequences can be used to image the brain. Most imaging sequences use the same basic principles presented in this chapter but with slight variations depending on the desired properties of the image.

Spin Warp Imaging

Spin warp imaging [14] is the most general form of MR imaging. In this sequence k-space is acquired in a series of rows. Each line is acquired by sequential fre-

quency and phase encoding steps as described in Section 2.2.2. The sequence is then repeated by varying the amplitude of the phase encoding gradient each time to fully sample k-space.

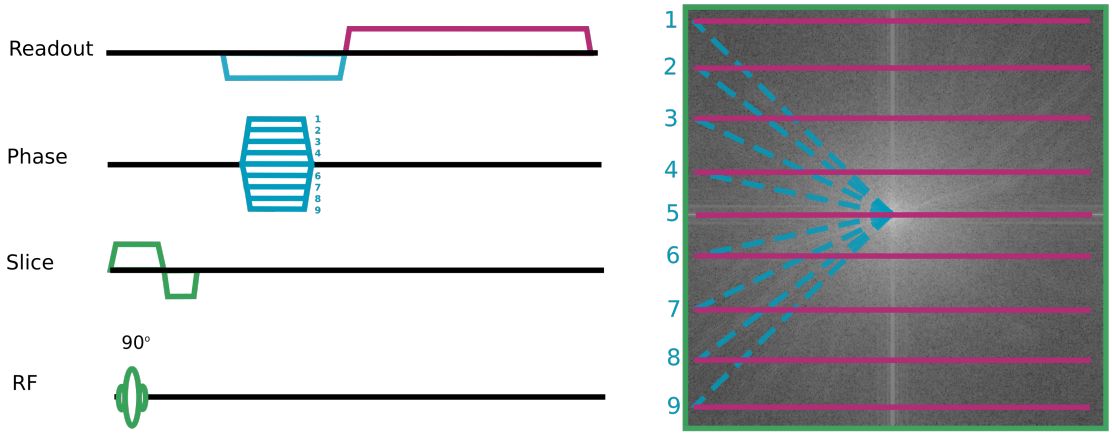


Figure 2.12: A pulse sequence diagram for spin warp imaging sequence using gradient echoes. The slice selection (green) is performed first to excite a slice of spins. Phase encoding (blue) and frequency encoding (pink) are then used to traverse k-space. This sequence is then repeated (in this case 9 times) to fully sample k-space.

Echo Planar Imaging

Echo planar imaging (EPI) is a fast imaging sequence where the whole of k-space is sampled after a single RF-pulse. k-space is sampled in a “raster-like” fashion.

For gradient echo multislice EPI (where k-space is acquired in a series of 2D profiles), this is achieved by first applying a 90° RF pulse at the same time as the slice select gradient. Shortly after, a large phase and readout gradient are applied to ensure that k-space is first sampled at the extremities. Readout gradients are then alternately applied to sample back and forth across k_x. Between the alternating readout gradients, short phase gradient “blips” are applied to advance the trajectory in the phase encoding direction. A pulse sequence diagram for EPI and its associated k-space trajectory can be seen in Figure 2.13.

The repeated reversal of the readout gradients cause refocusing of spins (gradient echoes) which means there is a peak in NMR signal at the centre of each

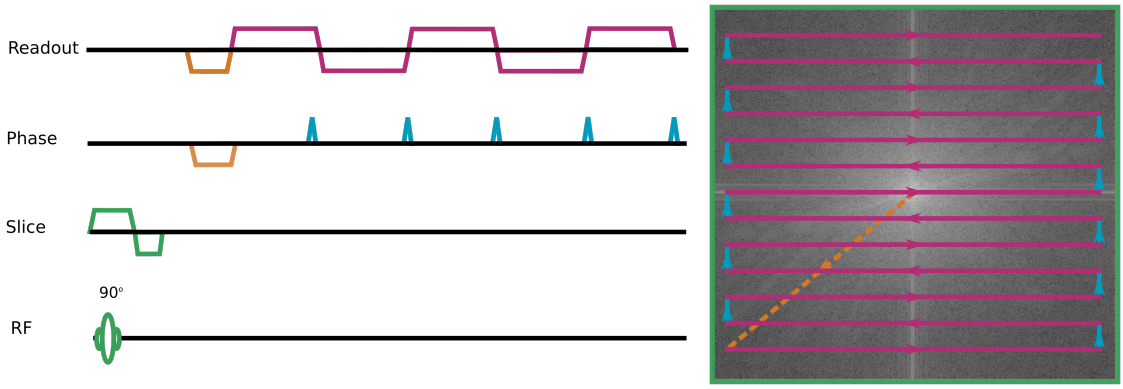


Figure 2.13: The pulse sequence diagram for a 2D gradient echo EPI sequence and its k-space trajectory.

readout gradient (corresponding to $k_x = 0$). Over the course the the acquisition the overall NMR signal decays with T_2*.

Because of the speed in which EPI images are acquired and its T_2* sensitivity, it forms the cornerstone of functional magnetic resonance imaging (fMRI) where many images must be acquired over an extended period of time.

Spoiled gradient echo Imaging

Spoiled gradient echo imaging describes a diverse collection of sequences which sample k-space through multiple RF excitations or "repetitions" which use a similar pulse sequence to that shown in Figure 2.12. These excitations occur frequently meaning the repetition time (TR) is short (TR $< T_2*$). The short repetition time means that the longitudinal relaxation does not fully recover, instead it reaches a new equilibrium magnetisation (M_{SS}), this situation is known as steady-state free precession (SSFP).

The short TR also means that transverse magnetisation is often also not fully dephased before each excitation. However, in spoiled gradient echo SSFP sequences the transverse magnetisation is spoiled between repetitions.

Spoiling is the act of disrupting the transverse magnetisation. The transverse magnetisation of a sample decays with T_2*, for sequences with long TRs the tranvserse magnetisation decays away naturally over time. If the sequence requires shorter TRs however, the transverse magnetisation must be destroyed by other

means. For gradient spoiling, additional gradients lobes are applied at the end of each repetition. The strength of these gradients are varied between excitations which causes spins to rapidly lose phase coherence. For RF-spoiling the phase of the RF pulse is varied between repetitions [18]. The phase variation between RF pulses is described as ,

$$\Phi_n = \frac{1}{2}\Phi_0(n^2 + n + 2), \tag{2.38}$$

where Φ_n is the phase of the nth RF pulse and Φ_0 is either 117° or 123°.

The signal of a spoiled gradient echo sequence can be described as,

$$S \propto \frac{\sin\alpha(1 - \exp^{-TR/T_1})}{1 - (\cos\alpha)\exp^{-TR/T_1}} \exp^{TE/T_2*}. \tag{2.39}$$

So the signal depends upon 3 variable parameters, TR, TE and flip angle(α). By varying these parameters different image contrast can be obtained. Long TE's can be used to increase the image's T_2* weighting while larger flip angles and shorter TR's increase the image's T_1 weighting.

2.2.4 Parallel Imaging

Parallel imaging is a technique that reduces scan times by using multiple receive coil arrays to measure the MR signal. Each coil in the array is located in a spatially different location within the MR environment and hence has a different spatial sensitivity.

Sensitivity encoding (SENSE) is a common parallel imaging technique in which at the start of a scan session the coil sensitivity profiles are determined [19]. Subsequent SENSE enabled scans then acquire every R k-space lines where R is the SENSE factor. For SENSE factor 2 or every other line in the phase encoding direction is acquired, SENSE factor 3 every 3rd line is acquired etc. This results in an image with a reduced matrix size that is wrapped in the phase encoding direction by a factor of R.

Each pixel in the reduced matrix is the sum of R pixels in the true image. But each pixel contribution is weighted by the sensitivity of a particular coil. By the

spatial sensitivity of each coil it is possible to resolve the original image.

Using this method one can in theory produce acceleration factors of up to 32 (for a 32 channel receive coil) however this would cause a large drop in signal to noise so in practice much smaller factors are used.

A crucial stage in the SENSE process is the acquisition of coil sensitivity maps at the beginning of a scan session, these maps are obtained by scanning a large field-of-view (FOV) with each coil. Each image is then normalised and smoothed to obtain a sensitivity profile, however sensitivity can only be obtained where signal is generated, i.e where there is an object to scan.

An alternative to SENSE is GeneRalized Autocalibrating Partial Parallel Acquisition (GRAPPA) [20]. Similar to SENSE, in GRAPPA k-space lines in the phase encoding direction are skipped, however the centre region of k-space is fully sampled (autocalibration signal (ACS) region). The ACS region is then used to estimate how each coil distorts k-space. This is then used to fill k-space with the missing lines.

2.3 MRI System

The MR scanner is dependent upon generating accurate magnetic fields across a range of frequencies. These magnetic fields are generated by some mechanisms within the scanner depending on their desired characteristics.

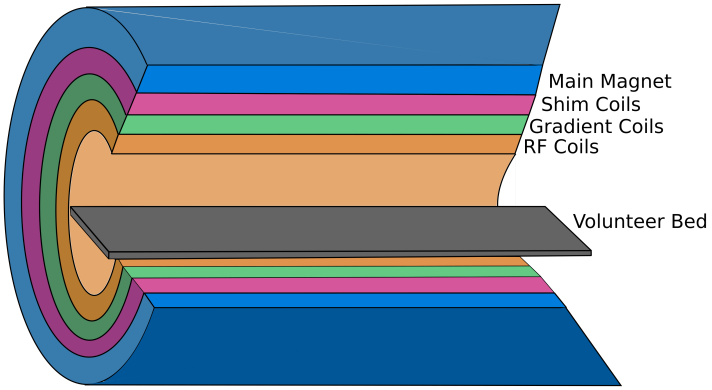


Figure 2.14: The typical coil arrangement used to generate the magnetic fields necessary for MRI

2.3.1 Static Field

The static magnetic field (B_0) of an MRI scanner is the largest outer-most component of the MRI scanner. Most modern MRI scanners use superconducting wires to generate the large magnetic fields required for MRI. A solenoid is made by wrapping super conducting material (commonly Niobium-Titanium or Magnesium Di-Boride) around a cylinder in a helix shape. This solenoid is then submerged in a bath of cryogenic fluid, usually liquid helium ($\approx$ 4 K) although there has been a recent drive to find high temperature super conductors that use liquid nitrogen ($\approx$ 77 K) [21]. When the solenoid's temperature drops below a critical temperature the electrical resistance of the material is essentially zero. A large electrical current can then be passed through the solenoid. This current induces a magnetic field in the direction of the solenoid. The strength of this field in the centre of the solenoid is given as,

$$B = \mu_0 \frac{IN}{L}, \tag{2.40}$$

where B is the strength of the magnetic field, μ_0 is the permeability of the vacuum, I is the current, N is the number of turns and L is the length of the solenoid.

Because the resistance of the material is low, extremely high currents can be used which generate very strong magnetic fields. Once the magnet is "at field" (at the field strength desired) the solenoid circuit can be closed creating a closed loop. With negligible resistance in the coil, current is allowed to flow indefinitely. This means that the static magnetic field can have high strength without requiring a continuous electrical current.

Fringe Field

The fringe field (the field that extends outside of the magnet) needs to also be considered especially when working at ultra high field strengths of 7 T and beyond. In general safety regulations restrict access to magnetic fields greater than 5 Gauss. To reduce the size of this area, active or passive shielding is employed [22].

In active shielding, the static field is countered by an opposing magnetic field.

This opposing field causes the field strength to decrease much more rapidly further away from the field. The field strength for an unshielded magnet declines with R^{-2} where R is distance from the magnet. An actively shielded magnet declines with R^{-4}. However, active shielding potentially decreases the homogeneous volume within the main magnet.

Passive shielding uses large quantities of iron and steel in the construction of the room housing the MRI scanner. These materials act to concentrate the magnetic flux lines within the walled volume. Care must be taken when using passive shielding to ensure that the scanner is placed about the symmetrical centre of the room to balance the magnetic pull of the magnet to the room. For the 7 T scanner at Nottingham the room housing the scanner contains 200 tonnes of iron and steel.

2.3.2 Shim Coils

Because homogeneity of the magnetic field is essential for MRI, additional magnetic coils are added to refine the static field. While great care is taken to generate a homogeneous field, due to physical limitations and factors perturbing the field create areas of inhomogeneity. These are corrected for during construction using passive shimming. In passive shimming the magnetic field strength within the bore is measured and areas of inhomogeneity identified. These areas are then corrected for by placing small pieces of iron into trays in the bore wall to smooth the field. This can reduce field inhomogeneity from 100 ppm to 1 ppm.

However,as previously discussed (Section 2.1.7) susceptibility causes perturbations in the magnetic field. These perturbations vary depending on the size, orientation and material of the object within the field. To overcome these subject dependent variables active shimming is employed to mitigate the effects as much as possible. Additional coils (shim coils) are placed within the bore that can be turned off and on at will. These active shim coils produced constant field gradients and higher order fields which can be adapted to a region of interest. Typically image based shimming is used, whereby a B_0 map is acquired and an

area of a scan is selected to maximise the homogeneity over a region by varying the amplitude of the current supplied to the shim coils.

2.3.3 Gradient Coils

As discussed in Section 2.2.2, the MR signal is localised by encoding the signal using spatially varying magnetic gradient fields. These gradients generate magnetic fields in the z direction that vary in strength spatially along each axis of the scanner (x, y, z). The gradients are assumed to be linear within the imaging field of view, the linear region of the gradients is typically of the order of $\pm$ 20 cm from isocentre.

At their simplest, the z gradient coil usually consists of a series of circular coils (Maxwell coils) along the cylindrical bore, while the transverse gradient (x and y) coils consist of a series of saddle coils (Golay coils) on each side of the bore. However often more complex "finger-print" designs are generally used by the manufacturers.

The 7 T scanner at Nottingham has a peak gradient strength of 40 mT/m and a maximum slew rate of 200 T/m/s. The slew rate is defined as the time it takes the gradient to achieve is maximum strength and is determined by the inductance of a coil and performance of amplifiers. Care must be taken when performing in-vivo scanning that the gradients do not produce excessive peripheral nerve stimulation (PNS) in the subject as a result of the rapidly changing magnetic field which induce electric fields and current within the body.

Concomitant Fields

Maxwell's equations for divergence and curl describe additional fields that are generated as a result of the desired linear gradient field. This accompanying field is known as the concomitant field (or Maxwell field) and is spatially non-linear [23–25]. For gradient fields orientated along z that are linear in strength along x (G_x), y (G_y) and z (G_z) the lowest order of the concomitant field (B_c) is

described as,

$$B_c(x,y,z,t) = \frac{1}{2B_0}(G_x^2 z^2 + G_y^2 z^2 + G_z^2 \frac{x^2+y^2}{4} - G_x G_z xz - G_y G_z yz). \quad (2.41)$$

However, given that $G << B_0$ the size of the concomitant field is generally small and is usually ignored during most imaging sequences. However for systems with strong gradients or at low field strengths, these fields must be considered.

2.3.4 RF Coils

The RF coil is responsible for the transmission and detection of the transverse RF field. Because of this the RF coils may be split into separate systems; the transmit and receive coils. Figure 2.14 shows the RF coils embedded within the magnet bore, however it is also common to use coils placed close to or directly on the subject. Coils placed closer to the subject offer a greater signal-to-noise ratio at the expense of patient comfort. In general, the transmit coil is large to provide RF to the entire imaging volume while the receive coils are much smaller so they can be positioned closer to the subject to increase the signal to noise ratio.

Often more than one receive coil is used for a given region of anatomy to provide a larger field of view or to enable parallel imaging. There has also been a recent push to use multiple transmit coils enabling a more homogeneous B_1 field at high field [26–28].

Chapter 3

Review of Motion Correction in MRI

3.1 Motion and MRI

The duration of an MRI scan varies greatly with scan type, but is typically of the order of minutes for a volume acquisition. This makes MR images extremely susceptible to subject motion. This is a major problem both in research where high resolution scans are limited by even the smallest of subject motions [29,30], and also clinically where it is estimated that motion artefacts in neuroimaging alone cost $140,000 (hospital revenue) per scanner per year [31].

Depending on the type of scan and movement, motion artefacts present in different ways. Generally motion during anatomical imaging (where the goal is to obtain one high resolution, high contrast image) results in ghosting and blurring [32] as seen in Figure 3.1.

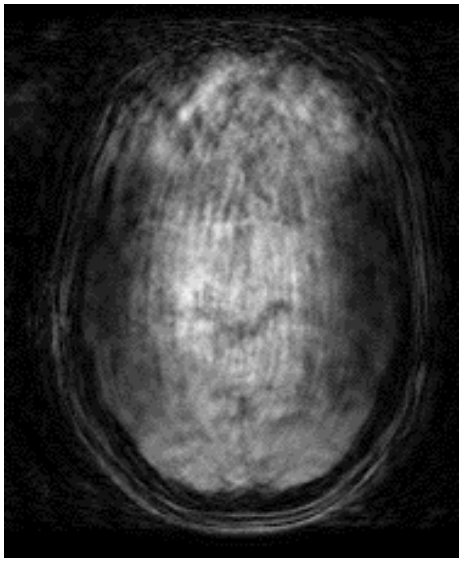


Figure 3.1: An MPRAGE image that has been distorted by subject movement, in this case the subject shook their head throughout the scan.

Because MR images are acquired in k-space, whole images can be affected by subject motion at any time during the acquisition. The severity of the image degradation is not only dependent upon the size of the movement, but also when the movement occurs: small movement while the centre of k-space is being sampled can be more detrimental to image quality than a large movement at the extremities of k-space. This is because the majority of the image information is

contained in the centre of k-space.

Rigid body motion can be split into two components; rotations and translations. Subject rotation results in a rotation of the underlying k-space data. MRI scanners depend on the correct assignment of data into k-space when sampling the NMR signal. When the subject rotates the data is misplaced in the k-space matrix. The transform of k-space due to a rotation is described as

$$\begin{bmatrix} k_{x'} \\ k_{y'} \\ k_{z'} \end{bmatrix} = \mathbf{R} \begin{bmatrix} k_x \\ k_y \\ k_z \end{bmatrix} . \tag{3.1}$$

where $\mathbf{R}$ represents a rotation matrix that describes the subject motion, $k_{x,y,z}$ are the positions of a k-space data point prior to subject rotation and $k_{x',y',z'}$ are the new co-ordinates for the same k-space data point after rotation.

Subject translations present slightly differently. Rather than altering the k-space co-ordinate of the data, they change the phase of the k-space data. The phase change can be described as,

$$\psi'(k_x, k_y, k_z) = \psi(k_x, k_y, k_z) e^{i\pi(T_x k_x + T_y k_y + T_z k_z)} \tag{3.2}$$

where $T_{x,y,z}$ are the translations in x, y, and z respectively and ψ is a particular k-space data point. The magnitude of the data point is unaltered.

Image artefacts occur due to errors in k-space (phase changes or rotations) rather than image space, motion during any part of the k-space acquisition causes distortions across the whole image. As such scans where images are gathered quickly are generally more resistant to subject motion, simply because less movement can occur during the acquisition of the image.

If many images are acquired during one temporal data set, such as in the case of fMRI, then inter-volume motion (or inter-slice for 2D multislice scans) can occur. For inter-volume motion, artefacts within an individual image are minimal, rather, subsequent images are misalligned [33]. This is a common problem in fMRI studies where scans can last long periods of time meaning subject motion is high.

These artefacts are often corrected to some extent with additional retrospective processing.

Additional artefacts can arise due to subject movement during gradients. As previously mentioned the gradients are used to encode the NMR signal based on the location of the spins, spins moving during spatial encoding experience a varying magnetic field which causes them to lose or gain phase, this change in phase causes spins to lose phase coherence with each other more rapidly. The phase shift is proportional to velocity, gradient strength and the square of the time it moves within the gradient.

Additional changes in contrast can occur due to spin history effects in the presence of subject motion. For multi-slice imaging, slices through the subject are excited and imaged successively. When a subject moves, regions of slices that were previously excited may be excited again. If the time between excitations is short this causes the longitudinal magnetisation to increase compared to stationary regions.

3.2 Motion by Anatomy

Motion is an issue in all *in-vivo* MRI. However different regions of the body can be characterised by different types of movement. In this thesis we consider rigid body motion of the head. However, it is worth first addressing methods that are commonly used for body imaging.

3.2.1 Movement in Body Imaging

Motion during body imaging is largely unpreventable because it arises from physiological body functions such as the heartbeat or breathing. These types of motion present as deformations in the image where regions of the body expand or contract, because of this rigid body motion characterisation is not sufficient.

Gating and Triggering

A common technique is cardiac/respiratory gating/triggering. Using this technique the state of the heart/lungs is monitored by some means and used to determine what data is used for reconstruction of the image. For the heartbeat, an electrocardiogram (ECG) or a peripheral pulse transducer is used to track motion [34, 35]. Care must be taken however, there is a signifcant delay between the contraction of the heart and the pulse wave to the finger, which makes cardiac MRI difficult using a peripheral pulse transducer. For lung imaging bellows can be placed on the subject's abdomen. When the chest expands due to the subject breathing, the pressure in the bellows increases. This pressure can then be monitored to determine the subjects respiratory cycle [36, 37].

During a gated sequence data can be continuously acquired regardless of subject motion. After acquisition the data can be retrospectively discarded so that only data where the subject was stationary is included, alternatively it can be binned to reconstruct images at different parts of the cycle. Triggering is a prospective version of gating whereby the physiological state of the subject is monitored and data is only acquired in a given state (lungs exhaled or heart expanded for example).

Cine MRI

Cardiac imaging contains extremely high amounts of motion due to the rhythm of the heart. Cine MRI uses a rapid imaging technique such as balanced-state free precession (BSSFP) gradient echo or EPI sequence to image the heart as it moves [38, 39]. BSSFP is commonly used because it has a high signal to noise ratio when imaging the heart due to the high T_2/T_1 ratio of blood. However any imaging sequence that allows very short echo and repetition times can be used (such as EPI) provided that lines of k-space can be acquired very quickly.

The cardiac cycle is split into several "frames" depending on their time through the cycle. K-space lines are continuously acquired over the cardiac cycle and assigned to one of the frames. Over the course of several heartbeats ($\approx$ 20

cycles) the k-space of each frame is filled and the image acquired. This method makes the assumption that the motion of the heart is identical through multiple cycles which makes it well suited for cardiac imaging, but unsuitable for lung imaging.

Navigator Echoes

An alternative to physiological tracking devices are navigator echoes. The simplest navigator echo is a 1D pencil beam that can be placed over the liver or diaphragm for example [40]. For a 1D echo, motion is presumed to be along 1 axis only (making it ideal for respiratory gating where organs move largely along the foot-head (FH) direction (scanner z axis)) with the echo acquired along the readout direction. This echo can then be reconstructed to give a 1D slice profile of the area. By gathering many of these echoes through time one can then determine motion along this axis. More complex navigator echoes are also used that can determine motion in 6 degrees of freedom by using a 3D navigator echo such as spherical [41] or cloverleaf [42] navigators. For navigator information a minimum of a 1D profile is acquired however if sufficient time is allowed a whole 2D or 3D image can be acquired to form a navigator image as described in Section 3.6.1. While these sequences don't require additional equipment, they do require additional time in the sequence and may modify spin history.

Reducing Signal From Moving Regions

Because data is acquired in k-space, each acquired k-space point contains information from the whole image. This means that a region of interest of an MR image can still be corrupted by motion even if it remains stationary if other regions of the image move. This effect is common for respiration and body imaging where the front of the body moves as the subject breathes while the spine remains stationary. In this case it may be beneficial to reduce the signal contribution from these moving areas using saturation bands. A spatial saturation pulse is a 90° RF pulse that is designed to suppress regions of the image through dephasing [43]. In this

case a spatially selective saturation pulse is applied to the front of the body prior to the main scan. A sequence diagram for a spatially selective saturation pulse is seen in Figure 3.2.

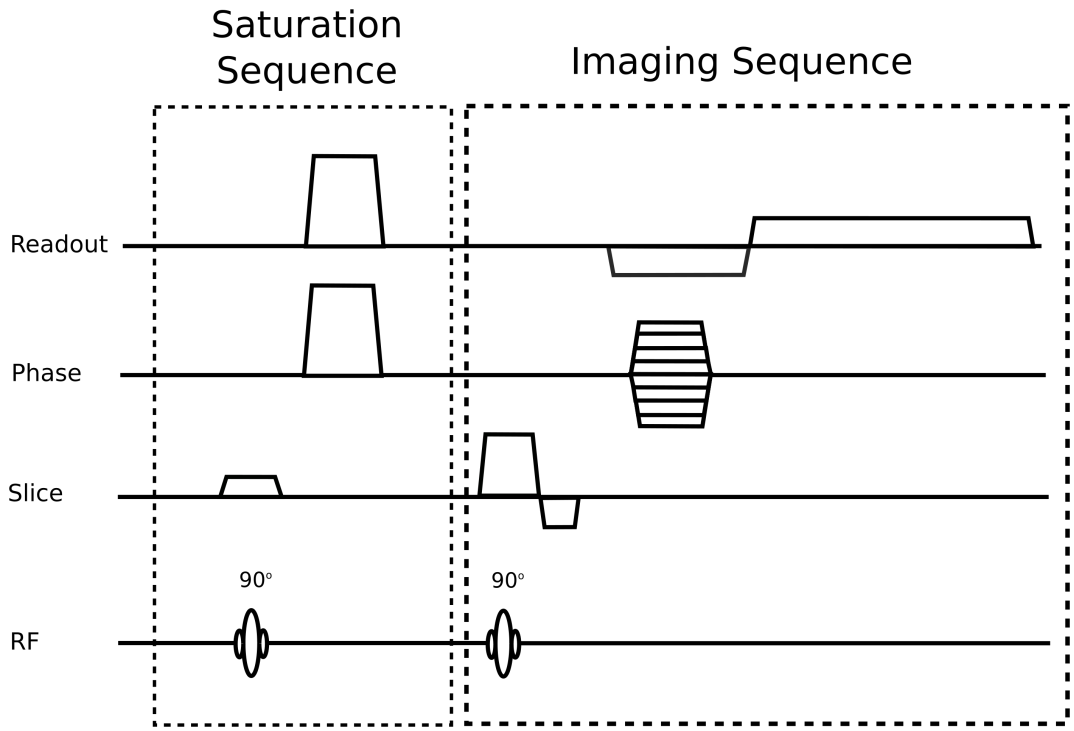


Figure 3.2: The gradient pulse diagram for a pre-scan saturation pulse.

The saturation pulse is set at a slightly different frequency to the excitation pulse and hence results in larger phase dispersion without the presence of a rephasing gradient. This effect is compounded by the use of spoiler gradients which act to remove the transverse magnetisation.

Alternatively the sequence contributions from moving anatomy can be reduced by turning off the receive elements located close to moving regions.

3.2.2 Movement in Neuroimaging

For neuroimaging almost all motion is caused by bulk head movement. With the exceptions of sudden impacts and accelerations, brain motion is generally described as rigid body movement [44,45] which can be described as a combination of translations and rotations about each axis applied uniformly to the whole

image. These transforms are usually described by a 4 $\times$ 4 matrix. A separate matrix is used to describe the translation and rotation around each axis :

$$\begin{aligned}
\mathbf{T} &= \begin{bmatrix} 1 & 0 & 0 & T_x \\ 0 & 1 & 0 & T_y \\ 0 & 0 & 1 & T_z \\ 0 & 0 & 0 & 1 \end{bmatrix}, \\
\mathbf{R_x} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & -\sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \\
\mathbf{R_y} &= \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \\
\mathbf{R_z} &= \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.
\end{aligned} \tag{3.3}$$

For the remainder of this chapter we will review different methodologies of determining subject movement assuming rigid body motion only.

3.3 Image Registration

Image registration is a method to reduce the misplacement between two images by applying a geometric image transform across to one of the images. Generally a transform is determined over many iterations of repositioning by minimising some cost function. The cost function varies depending on the various registration methods used, commonly used are mean square difference and mutual information. Mean square difference is a pure pixel intensity based method in

which the difference between each pixel in the proposed image (the image after the transform has been applied) and the fixed (the image one is registering a second image to) is taken, squared (to remove any sign differences) and then averaged across the whole image [46]. This is one of the quickest and simplest methods however it is only suitable for images from the same modality because it assumes the contrast weighting in each image is the same. For multi-modal registration mutual information is more commonly used. Mutual information is a measure of the dependency of one variable to another [47], that is to say if one variable is known, how much about a second variable is also known. If two variables are independent, then knowing one variable grants no information about the second and so the mutual information is zero. If one variable is a deterministic function of the other, then the mutual information is exactly the same as the entropy of the first function. Because mutual information is not directly related to image contrast, it can be used to register across multiple modalities or across images where scanning parameters have resulted in difference in image contrast.

Unlike most other methods which only permit rigid body transforms, image registration algorithms can be extended to correct higher order image geometric transforms such as shearing, stretching and scale. However they do involve interpolation.

3.4 Minimising Motion Effects

In general, the effects of subject motion can be reduced by decreasing scan times using methods such as parallel imaging (Section 2.2.4). However where scan times cannot be shortened additional measures must be taken.

3.4.1 Restraint and Sedation

The simplest method of reducing the effects of subject motion is to prevent the subject from moving during the scan. For subjects who are unable to remain still for the length of the scan, restraints or sedation may be necessary. In most head

MRI scans the subject is padded within the head coil. This padding helps dampen the noise of the MRI scanner and is normally sufficent to prevent a compliant subject from moving. If the subject is non-compliant then more extreme measures such as head or body straps are used to restrain the subject [48] .

Sedation is an alternative method to prevent the subject moving. Sedation is common in clincal practice to overcome voluntary subject movements as is the case in many paediatric scans, where children are either too young or nervous to remain still for the entire scan. In this case subjects undergo general anesthesia with chloral hydrate [49]. Sedation has been shown to reduce the need of scan repeats in infants. The side effects of chloral hydrate include dizziness, hyperactivity and vomiting for up to 4 hours after the scan however [50], and very little is known about the long term toxicity of the sedation agents in infants and neonates [51].

Despite these methods of reducing the subject motion it is often impossible to prevent artefacts entirely and as such subject motion needs to be directly corrected for.

3.4.2 2D vs 3D Imaging

If scans cannot be further accelerated then it may be possible to reduce motion artefacts by changing the scan sequence itself. In general 3D imaging has better SNR compared to multislice imaging however, it is also more susceptible to motion. As previously mentioned any motion during the aquisition of k-space distorts the entire image, in 3D imaging the entirety of k-space is captured as one scan and hence any motion can distort the entire volume. In multi-slice imaging each slice is captured seperately hence motion during acquition of one slice does not affect any other slice. However multi-slice imaging can result in inter-slice misalignment. Misaligned slices can be corrected retrospectively by acquiring a quick low resolution 3D volume and performing slice to volume registration between each slice in the multislice image, to the volume image [52].

3.4.3 Gradient Motion Nulling

If motion persists over the time of signal acquisition, gradient motion nulling can be used to reduce motion artefacts [53]. In this case additional readout gradients are applied before the main readout gradients (see Figure 3.3). These gradients act to cancel the phase gained by moving spins while preserving the encoding of static spins. For a spin which has constant speed, the phase gain during each gradient is described as,

$$\phi = \frac{1}{2}\gamma G v t^2, \tag{3.4}$$

where ϕ is the phase, γ is the gyromagnetic ratio, G is the gradient strength v is the velocity of the spin and t is the time of the gradient lobe. The net phase change is then the integral of Equation 3.4 with respect to time,

$$\phi = \gamma \int G(t)v(t)t^2 dt. \tag{3.5}$$

Each gradient lobe can be treated discretely so the integral becomes the sum of the contribution of each gradient lobe to the phase. That is to say the first lobe contributes $1kT^2$, the second contributes $3kT^2$ and the 3rd $5kT^2$. For a regular gradient read out sequence a moving spin experiences $-1kT^2 + 3kT^2 = 2kT^2$ phase change at the time of the echo, during a gradient nulled sequence the same spin experiences $+1kT^2 - 2(3kT^2) + 5kT^2 = 0kT^2$ resulting in no phase change. Gradient moment nulling however, is limited to movement that is constant throughout the encoding period, which is often not the case for typical subjects. In addition the added gradients mean longer echo times are needed to allow for additional gradients to be played out, these long TE's make sequences longer and hence potentially *more* susceptible to motion.

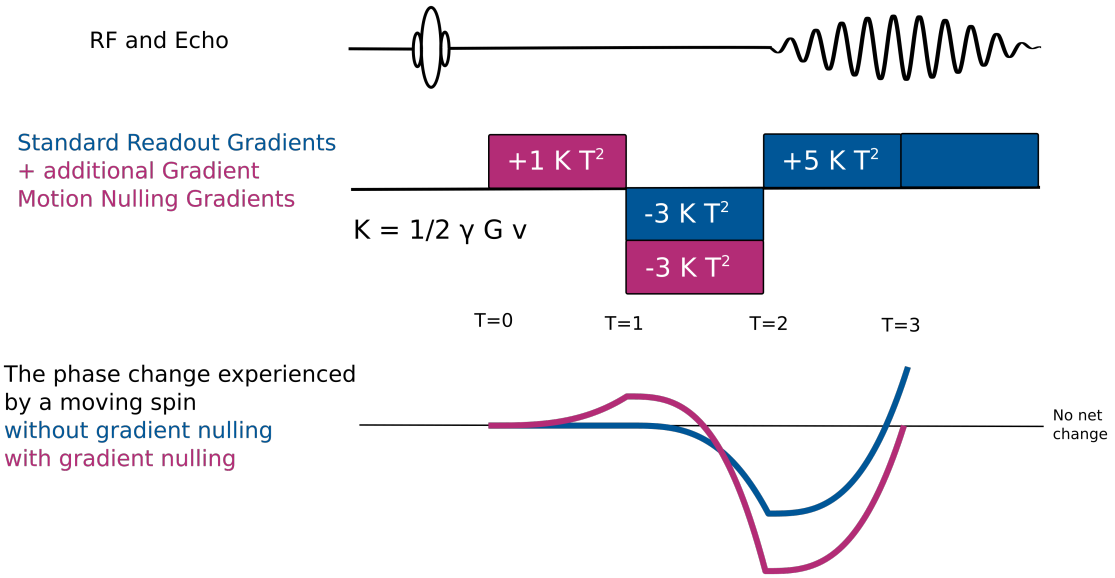


Figure 3.3: A gradient diagram for a typical gradient echo sequence. The standard readout gradients are seen in blue. For gradient motion nulled sequences additional gradients (pink) are included. These act to remove any phase change in spins that are moving.

3.5 Correcting Motion

In general, motion correction methodologies exist within two categories, prospective and retrospective. In prospective correction data acquisition is adjusted for subject motion, in retrospective correction data is altered to correct for motion.

3.5.1 Retrospective Correction

Retrospective correction is performed after image acquisition. Inter-acquistion motion correction can be corrected by applying image transforms directly to the image. Often these image transforms can be obtained from the images themselves via image registration [54] and do not normally require prior information about how the subject has moved. This type of correction is most common in rapid-volume scanning such as fMRI where the artefacts due to subject motion present largely as image misalignment rather than artefacts within the image.

For intra-acquisition motion correction, corrections must be done on the raw k-space data. Intra-volume correction requires knowledge of the how the subject

moved during the scan and also the scan trajectory (how k-space was sampled). This k-space trajectory can then be altered retrospectively by rotating each line based on the subject movement data and the phase of the k-space phase data changed by the subject translation. The equations describing how to alter k-space for rigid body motion correction are:

$$\begin{bmatrix} k_{x'} \\ k_{y'} \\ k_{z'} \end{bmatrix} = \mathbf{R}' \begin{bmatrix} k_x \\ k_y \\ k_z \end{bmatrix}, \tag{3.6}$$

$$\psi'(k_x, k_y, k_z) = \psi(k_x, k_y, k_z) e^{i\pi(-T_x k_x + -T_y k_y + -T_z k_z)}, \tag{3.7}$$

where $\mathbf{R}'$ is a matrix that describes the reverse rotation of a subject.

3.5.2 Prospective Correction

Prospective motion correction refers to any method of correction that is performed during acquisition. Many of these methods involve tracking subject movement during the scan, this information can be used in different ways depending on the correction method and the information available. One option is to discard data if large amounts of movement are detected during the acquisition. This data can then be re-acquired once the subject has stopped moving. This method is useful if a quantitative measure of subject movement is not available, however, it can also lead to very long scan times, especially if a subject is repeatedly moving during acquisition, which could cause the scan to be repeated many times.

Alternatively one can update scanner sequences periodically using the tracking information to counter subject movement. In this case data are not discarded but instead, the acquisition of new data is altered such that the scan tracks the subject. The gradient fields are rotated for subject rotations and the readout frequencies are adjusted for translations.

3.5.3 Prospective vs Retrospective Correction

Retrospective correction offers many advantages over prospective correction. The ability to perform the corrections to data once it has been taken offline means that scan sequences do not to need to be altered, and scan times remain constant. Some inter-volume motion correction can be performed without any additional information. Retrospective correction is not exclusive to motion, is it common to correct for geometric distortions in echo planar imaging (EPI) data retrospectively [55]. Unlike prospective motion correction, retrospective correction is not time limited so longer iterative corrections can be applied to data [56].

However retrospective correction can have limitations which prospective motion correction does not. When performing inter-volume retrospective motion correction, the original image may be rotated to a new position. This rotation causes the original pixels (on a Cartiesian grid) to be rotated to a non-Cartesian grid and then interpolated back onto a Cartesian grid as seen in Figure 3.4. This process results in image blurring.

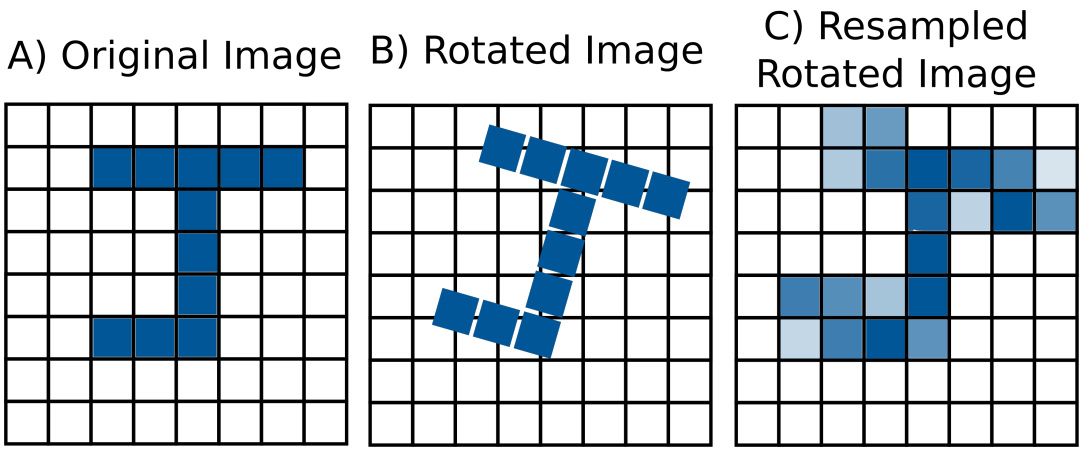


Figure 3.4: An example of how retrospective motion correction can result in blurring. The original image (A) is rotated to its new position (B). However the image is then displayed by re-sampling the image onto the underlying cartesian grid of pixels (C).

Retrospective intra-volume correction is applied after acquisition, therefore any data not obtained during acquisition is missing from the final reconstruction

(see Figure 3.5). This missing k-space data can result in image artefacts. Retrospective correction can also be difficult when performing multi-slice scans if the subject moves through the slices. In this case the slice of interest may not even have been excited and hence data will be lost entirely.

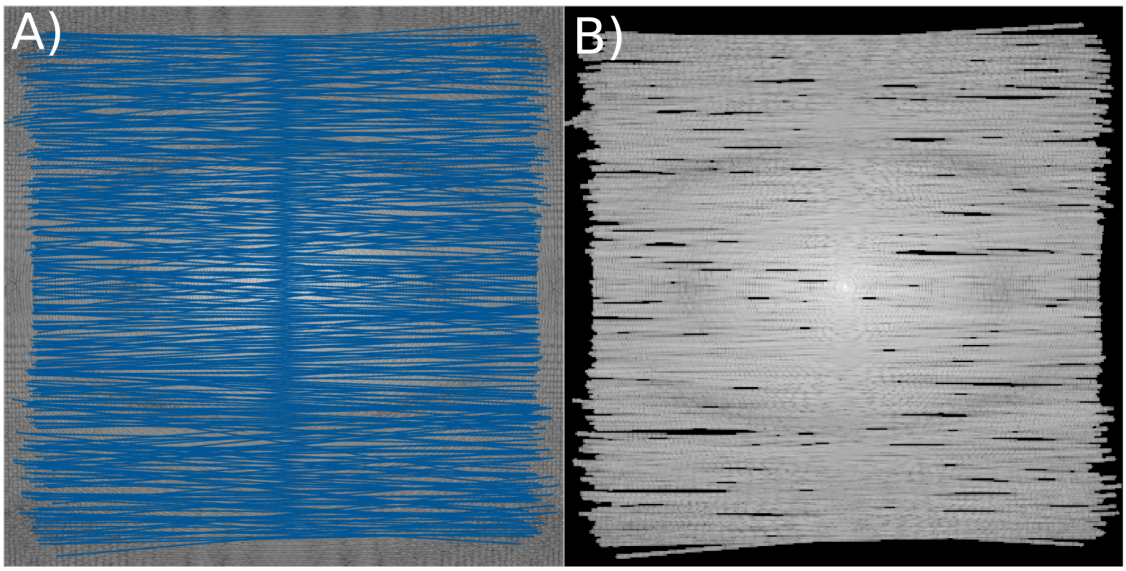


Figure 3.5: An example of of k-space after retrospective motion correction. A) Shows the new k-space trajectory after it has been adjusted for subject motion, B) shows the k-space data that is sampled and the missing areas of k-space where data was not captured.

Because prospective correction is performed as data is acquired, there are no missing k-space data. Prospective motion correction also allows scan parameters to be maintained relative to the head rather than just the imaging volume.

An additional benefit of prospective motion correction is it is less susceptible to spin history effects. If an object moves during acquisition areas of a volume can be inadvertently excited multiple times if the volume moves through areas of excitation. These present as variations in image contrast and can be difficult to remove retrospectively. Prospective correction ensures that only the desired regions of the volume are excited.

While prospective motion correction may result in slightly increased scan times, this is mitigated because the data that comes off the scanner is already corrected and does not require any additional resampling to allow for gaps in k-space.

3.6 MR Tracking Methods

For most correction methods (retrospective and prospective) knowledge of subject motion must be known to correct data. For retrospective inter-volume motion correction this can often be determined from registration of individual images, however for intra-volume motion correction or prospective motion correction alternative methods of determining subject motion must be used.

3.6.1 Navigators

As previously mentioned navigator echoes are sequences that capture one or more 1D profiles of an area to determine motion from physiological sources. These principles can also be applied during neuroimaging to determine head motion, provided a 3D navigator echo is used to capture motion in all 3 axes. Fully determining subject motion can be difficult from profiles because of the limited spatial information.

An alternative to navigator echos are navigator images. The principle of navigator images are similar to that of echoes however in this case an entire image is acquired periodically throughout the main scan. Image registration is then performed on the images to measure how the subject has moved, then used to correct the geometry of the main scan. These additional navigator images can increase scan time and change image contrast. It is important to make these scans as rapid as possible, while still retaining enough information to ascertain motion information. *Tisdall et al.* [57] proposed a volumetric navigator using EPI, Alternatively 3 2D orthogonal images can be acquired and registered to a volume image to measure motion in 3 dimensions [58].

In addition to increasing scan time navigators can have detrimental effects on image quality depending on how they are inserted within with the motion corrected scan. Low flip angles must be used in navigator images to prevent loss of contrast in the motion corrected scan, spin history from the navigators can result in additional artefacts as well. In order to reduce the effect of spin history, fat selective navigators have been proposed for head imaging [59]. In this case a fat

selective RF pulse is used during the navigator image causing only the fat around the skull to be excited. However this method still results in longer scan time.

3.6.2 PROPELLER

Self navigated sequences can be used to correct for motion, these sequences sample k-space in such a way that allows them to perform motion correction inherently; PROPELLER (Periodically Rotated Overlapping Parallel Lines with Enhanced Reconstruction) is such a method [60].

In PROPELLER sequences, k-space is sampled in a series of rectangular blades rotated around the centre of k-space as shown in Figure 3.6. Each blade contains a series of parallel linear trajectories that can be acquired either with a single or multiple excitations. These blades are periodically rotated around the centre of k-space until a full circle of k-space has been acquired. Because the centre of k-space contains the most information, and it is repeatedly sampled with each subsequent blade, each blade can be compared to the last.

This additional information is used during the reconstruction of the image. First, the phase of each blade must be corrected to ensure that each blade is rotated around the centre of k-space, only then is subject motion corrected for. Subject rotations can be determined by measuring the correlation between the overlapping k-space between each blade and the initial blade. The subject rotation can be corrected for iteratively, by rotating each blade to maximise this correlation. Translations are corrected for similarly by maximising the correlation between the phase of the k-space data. To reduce through plane motion artefacts each blade is weighted based on the correlation of the overlapping k-space centres (where low correlation indicates through plane motion).

While oversampling k-space offers additional information, the additional sampling time and the advanced reconstruction both add to total scan time. The area of k-space sampled by a PROPELLER sequence is described as,

$$\frac{2}{\pi}LN = M \tag{3.8}$$

where L is the number of lines, N is the number of blades and M is the full matrix diameter.

From this equation you can see to fully sample a matrix of size M with no gaps the overlap between blades must be 57% resulting in scans that are proportionally longer, if one opts to reduce the slice overlap so that the overall scan time is not increased, gaps between the blades result in image artefacts and so a balance is required between scan times and full k-space coverage.

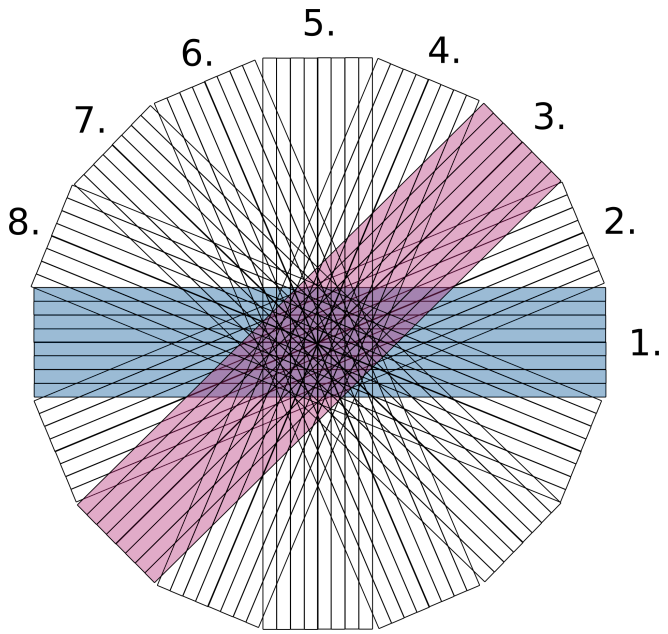


Figure 3.6: The PROPELLER k-space trajectory. The propeller scan constists of a series of blades (in this case 8). Upon completion of each blade, the trajectory is rotated to form the next blade. This continues until a circular k-space has been acquired

3.7 External Tracking Devices

Where high accuracy or minimal impact to the scan sequence are required, external tracking devices are often the preferred method of tracking subject motion. This is any device that is distinct from the scanner itself and provides geometry information, such as additional RF coils, NMR field probes and optical tracking.

3.7.1 RF Coils

RF coils as mentioned in Chapter 2 are used to detect the MR signal from a subject. However they can also be used to detect the gradient fields. These gradients generate spatially varying magnetic fields along each axis of the scanner. The gradient field is linear over typical imaging volumes with a null point (where the static field is unchanged by the applied gradient) at the gradient isocentre. When an RF coil is subject to the time varying gradient field a voltage is induced in the coil. The amplitude of this voltage is dependent upon the coils position within the gradient field. By applying gradients in each orthogonal direction a position of an RF coil can be determined.

In order to determine orientation of an object within the magnet multiple coils must be used with known orientations relative to each other [61]. The voltage induced in each coil from the applied gradient is dependent on the orientation and position of the coil, by using multiple coils where the orientation relative to the other coils is known, one can solve for the overall orientation.

The effective range of this method is limited to areas of the bore where the gradients are linear which varies between scanners. Because the RF coil determines its own location in space they are sometimes referred to as active markers.

3.7.2 Field Probes

NMR field probes can be used in a similar manner to RF coils to determine position in the magnet but use the NMR signal rather than induced field. In this case a small sample of ^{19}F is excited and then sampled using the same gradient encoding that the scanner would use to encode a 3D image. ^{19}F is used because of its relatively high gyromagnetic ratio (Table 2.1) and its stability. The free induction decay (FID) of the sample is measured while a gradient is applied and the magnetic field strength determined by measuring the frequency of the FID. By measuring the field induced by the gradients, the position of the probe within the magnet can be determined. The NMR FID is independent of the orientation of the probe however and so probe orientation can not be determined from a single

probe. Many (at least 3) rigidly connected probes must be used. Orientation can then be determined by fitting the individual probe positions.

3.7.3 Optical Tracking

Optical tracking methods vary greatly between methodologies but are generally more accurate than other methods of motion capture, only the field probe methods offer better accuracy [62]. Most methods involve the placement of one or more cameras within the bore of the magnet to track the subject, or some representation of the subject (a marker). Unlike other motion capture techniques which record motions in the magnet or image reference frame, optical tracking methods capture motion in their own external reference frame (normally that of one of the cameras) and must be converted to the magnet reference frame for use in motion correction. This transformation can produce errors in the measured motion as we investigate in Chapter 3.

Most optical tracking methods only allow the determination of rigid body movement with a notable exception of *Tracinnovations Tracline* system which generates a point cloud of the subject's head. This point cloud is then fitted throughout subsequent images to generate motion parameters. This system is able to achieve accuracies of approximately 0.3 mm, although accuracy is difficult to quantify because no ground truth is available. While this method represents the best option for patient comfort, the camera requires direct line of sight to the subject's face which in some scanner set ups is not possible.

Instead of tracking the subject directly, most optical tracking methods opt to track a marker representative of subject motion. This marker allows motion tracking even when the subject's head is obscured by a coil because the marker can be placed such that it protrudes out of the coil. Because a single optical camera only acquires 2D images, additional cameras or more complex markers are required to fully determine motion in 3D space. It is possible to capture 3D motion using multiple cameras by acquiring multiple images from different angles much in the same way eyes work [63]. By knowing the relationship between the cameras it is

possible to reconstruct multiple 2D images into 3D space and hence characterise motion with 6 degrees of freedom an optimised 2D representation of this can be seen in Figure 3.7.

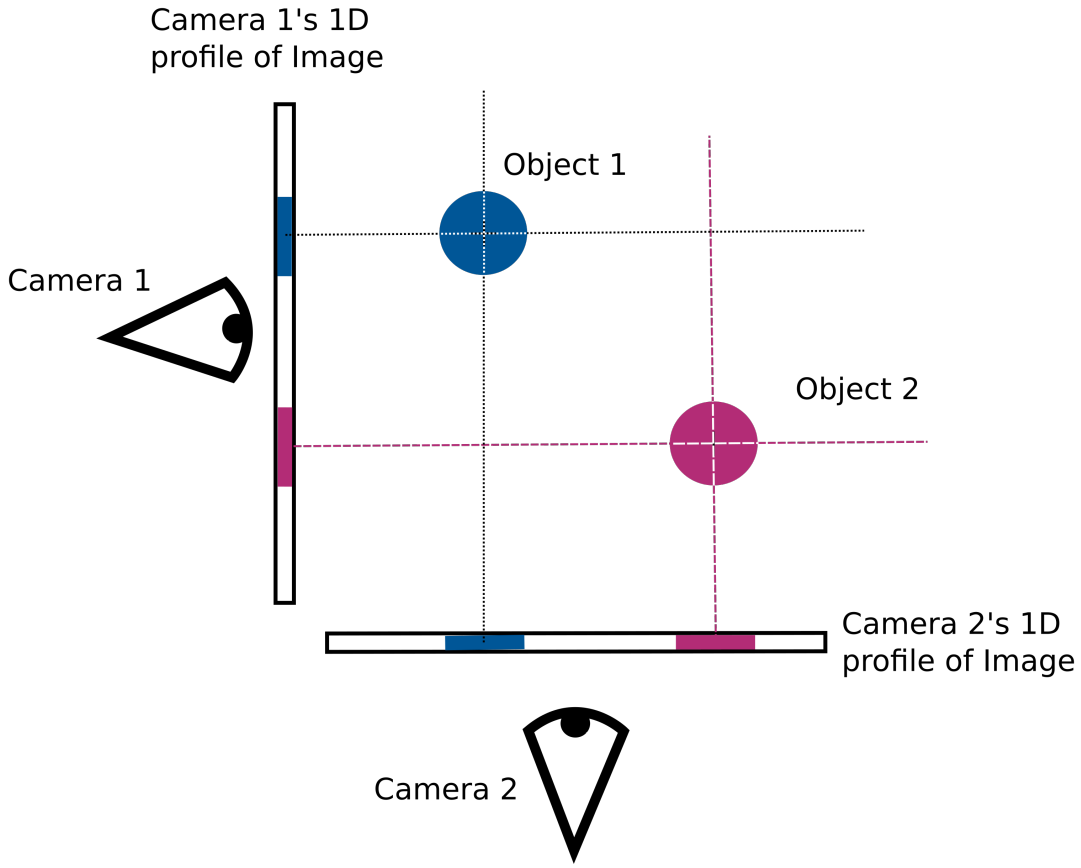


Figure 3.7: How multiple cameras can be used to determine an object's position. In this case each camera captures a 1D profile of a 2D space (much like real cameras capture a 2D image of a 3D space). Each object can then be projected back from each camera (assuming you know the transform between each camera) and the intersect is of each back projection then is the object location in 2D space

Alternatively multiple reference points can be used instead, assuming the reference points are rigidly connected [64]. Moiré Phase Tracking (MPT) is more closely related to the latter which uses a singular reflective marker consisting of multiple components to encode position and orientation.

3.7.4 Moiré Phase Tracking

The Moiré phase tracking (MPT) marker [65] is designed to capture both translations and rotations with a single marker by encoding the marker with different components; a starburst in the centre of the marker and 12 Moiré phase fringes. In the camera tracking co-ordinate scheme (Figure 3.8) tracking of the position of the marker in the x and y direction is performed using image registration of the starburst from the camera images. Tracking of the marker orientation and z position are obtained using the Moiré fringes.

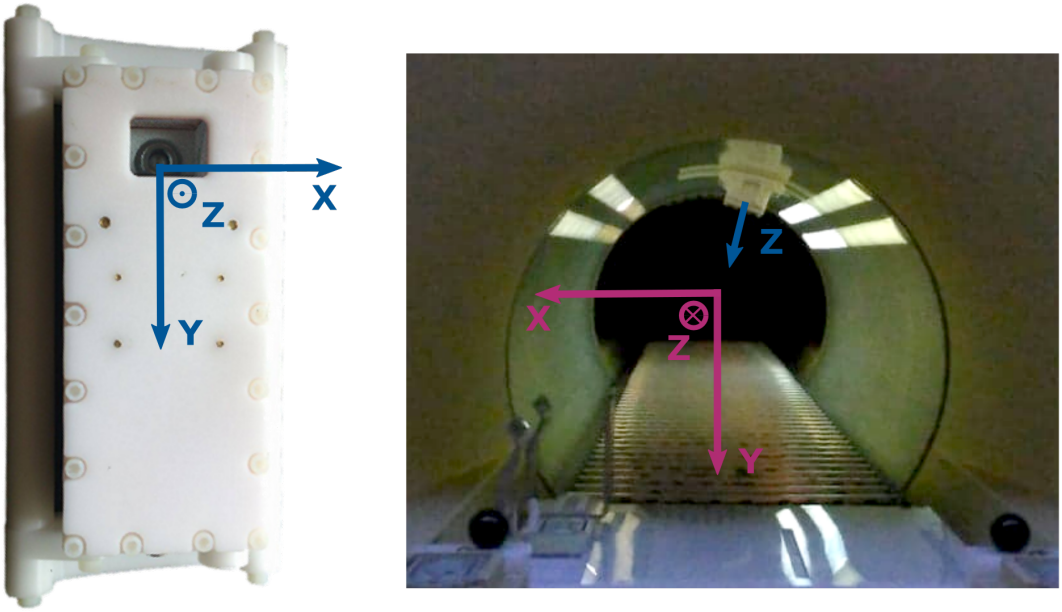


Figure 3.8: Photos illustrating the camera (blue) and scanner (pink) co-ordinate schemes.

These fringes are made by layering two equispaced transparent/opaque grids on top of each other in close proximity. Because a camera observes the grating from a single point, a larger amount of the rear grating is seen compared to the front grating hence from the camera's perspective more cycles are seen on the rear grating than the forward grating (Figure 3.9). This means that the rear grating appears to have a higher frequency than the front grating.

The overlapping frequency patterns interfere and form Moiré Fringe patterns. These fringe patterns are sinusoidal intensity patterns where the frequency is

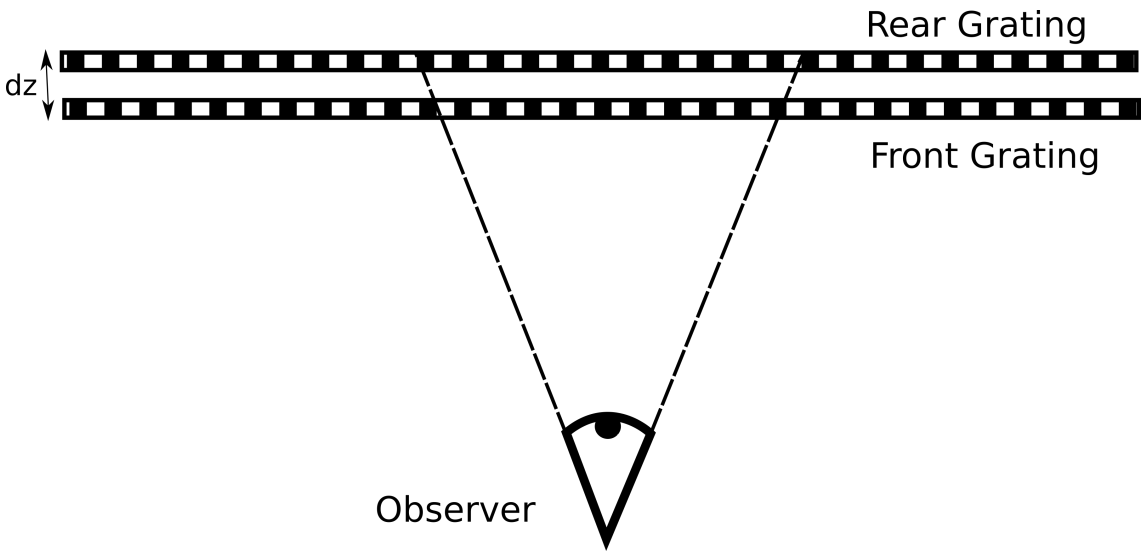


Figure 3.9: A diagram of the camera (observer) viewing two Moiré fringe gratings

dependent on the distance between the marker and the camera and the phase determines the orientation of the marker [66].

By fitting a sine wave to this intensity pattern the frequency and phase can be determined. Because the precision of one Moiré fringe pattern is not sufficient to determine absolute orientation, multiple Moiré fringe patterns are used, each with a slightly different spacing. One pattern describes the coarse marker orientation while the other describes the precise marker orientation. An example of the Moiré fringe patterns can be seen in Figure 3.10.

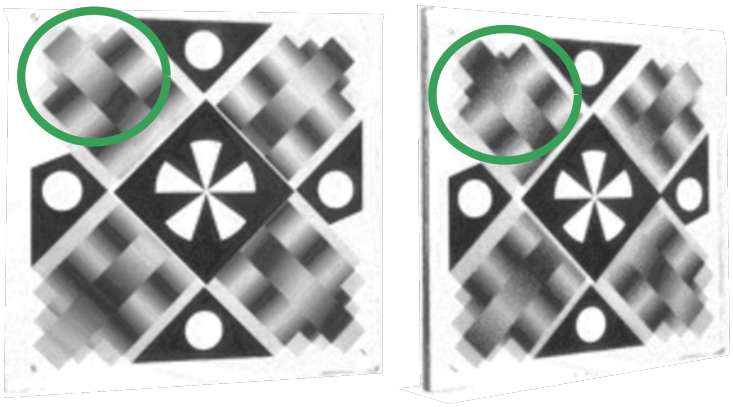


Figure 3.10: An example of the MPT marker in two different orientations, note the different phase of the fringe patterns.

The combination of the Moiré Phase Tracking and the conventional image tracking of the starburst, movement can be used to describe the position of the marker with 6 degrees of freedom and accuracies of up to 0.1 mm and 0.1 degree.

The accuracy of the Moiré phase tracking camera describes the accuracy with which the camera can track the marker, however this is not necessarily the accuracy with which the position of the imaging volume can be determined. Because the camera is an external tracking device, using it's own co-ordinate scheme, data must first be converted to the magnet reference frame before use in motion correction. This additional transform can amplify the measurement error and is investigated in Chapter 4.

For the remainder of this thesis the main motion tracking method used is the Moiré phase tracking camera described in Section 3.7.4. This method was chosen because of the minimal impact it has on scanner sequences while also offering highly accurate motion tracking.

Chapter 4

Strategies for Cross Calibration

4.1 Introduction

As described in Chapter 3, external tracking devices offer highly accurate motion measurements with minimal impact on the scanner sequence, however, because they are external devices, they capture motion in their own reference frame. This means that the co-ordinate scheme of the external device is often different to that of the scanner, which means motions detected by the external device must first be transformed to the magnet reference frame before they can be used in motion correction.

When tracking a marker using a camera there are 4 objects of interest, each with their own reference frame: the scanner (whose reference frame is defined relative to the gradient isocentre), the camera (whose reference frame is defined from the camera aperture), the marker (whose reference frame is defined by construction) and the imaging volume (whose reference frame is defined by the read-out, phase and slice directions).

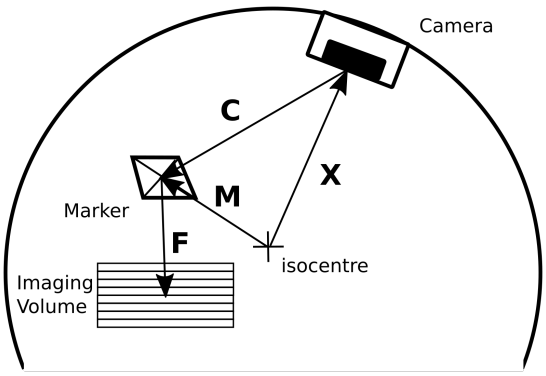


Figure 4.1: A vector diagram showing the transforms required for prospective motion correction. Isocentre-to-camera ($\mathbf{X}$), camera-to-Marker ($\mathbf{C}$) and marker-to-volume ($\mathbf{F}$).

To perform prospective motion correction, the transform between each must be known. In principle the volume-to-marker ($\mathbf{F}$) transform can be determined at the start of a scan. The imaging volume position in the magnet reference frame is set by the planning of the scan. The marker position is provided by the camera, assuming that the marker position in the magnet reference frame can be

determined using a calibration transform. Because the orientation of the volume can be inferred from the change in orientation of the marker, only the translational component of the marker-to-volume transform ($\mathbf{f}$) is required (see section 5.2). The camera-to-marker transform is provided by the camera tracking algorithm which tracks the Moiré fringe patterns on the marker to determine its position in the camera reference frame. The scanner-to-camera transform, however, must be determined experimentally prior to the scan. This transform is referred to as the "cross calibration" or just simply the "calibration".

In MRI, the cross calibration can be thought of as a quantitative description of the position of the camera within the magnet. The cross calibration varies every time the camera is placed in the bore and may even be susceptible to drift over time depending on the fixation method and thermal expansion of components.

There are several methods to determine the cross calibration. Most of these arise from the field of robotics where the aim is to convert a pose in the camera's reference frame to that of the room or "world" reference frame which is a reference frame describing the environment of robot. Although these methods have primarily been used to calibrate robot grippers many of them could be converted for use in calibrating an optical tracking camera in MRI. These cross calibration methods include iterative methods in which an initial estimate is made of the calibration, and then a series of movements is undertaken, these movements can either be specific movements (pure translation or pure rotation) [67] or natural movements (mixed translations and rotations about multiple axes) [68, 69]. For pure movements, the known movement performed is compared to the apparent movement measured using the calibration. The residual error between these values is then used to adjust the calibration until the error is below a specified threshold. This is repeated for additional movements with each iteration refining the calibration. More recently natural movements containing a mix of translations and rotations have been used as they are experimentally more simple. In this case, the principle is the same, but additional considerations are required since movement can now be about multiple axes.

The viewpoint method of calibration involves the placement of a known model

in the world reference frame, that is the reference frame you are converting camera data into. In *Motai et al. 2008* [70], they demonstrate this with a rigid grid of dots and additional landmarks. This known model can then be fitted to the camera observations. The orientation of the grid gives information about the camera's orientation and the spacing between the dots and the position of landmarks provides information about the camera's position.

Alternative methods involve recording multiple marker poses in both reference frames and solving for the calibration non-iteratively [71–73]. These approaches are similar to the iterative approaches in that they aim to reduce the difference between a pose in the known world co-ordinates and the apparent pose using the calibration, however, they rely on fitting across multiple poses to reduce the error. *Daniilidis et al. 1999* [73] describes a non-iterative calibration method that uses properties of quaternions to determine the cross calibration transform. Because quaternions can be used to describe screw motion, Daniilidis showed that only the line information is needed to determine the calibration, simplifying the problem. The *Daniilidis* method was chosen for use in this work because it is relatively easily implemented, particularly since the camera output is defined as a quaternion and so no conversion is needed. After poses have been acquired the calibration can be determined very quickly ($<$1 second). An additional advantage of this method is that is does not require prospective motion correction which most of the iterative methods require, and can be performed as a stand alone experiment.

However this non-iterative method does not provide any estimate of the accuracy of the calibration transform or the ultimate calculated position of the imaging volume. In addition this method requires gathering many phantom poses in the magnet and camera reference frames and the optimum number of poses to acquire is unknown. It seems that this method is not robust to errors in the pose determination, leading to calibration errors which can vary between metres to millimetres and hence can be very difficult to detect when a ground truth is not known.

Aim: The aim of this Chapter is to investigate the errors in calibration in three

stages. The first two stages involve the propagation of errors from calibration to motion tracking (Figure 4.2) and the final stage demonstrates a method of calibration transform validation.

In stage 1 we determine the impact that errors in the determination of poses have upon the subsequent calibration transform in simulation using a non-iterative cross calibration process [73] by varying parameters such as the total number of poses, the number of flawed poses and the type of pose error (translational or rotational).

In stage 2 we simulate the impact calibration errors have on the apparent imaging volume position and orientation when performing prospective motion correction. The aim was to determine acceptable levels of error in the calibration transform to define the position of the imaging volume within acceptable errors for a typical MRI experiment.

In stage 3 we conclude by demonstrating a novel method to measure the stability of the calibration, which does not require any additional data to be acquired. This new method can determine the reliability of a calibration before prospective motion correction to reduce wasted scan time due to errors in calibration while increasing confidence in the measured calibration transform.

Aspects of the work described in this chapter formed the basis of the following abstract: *James Smith, Olivier Mougin, Maxim Zaitsev, Benjamin Knowles, Richard Bowtell, Paul Glover, and Penny Gowland. From hand-eye calibration of optical tracking cameras to motion correction. In Proceedings of the International Society for Magnetic Resonance in Medicine, Number 1296, 2017.* [74]

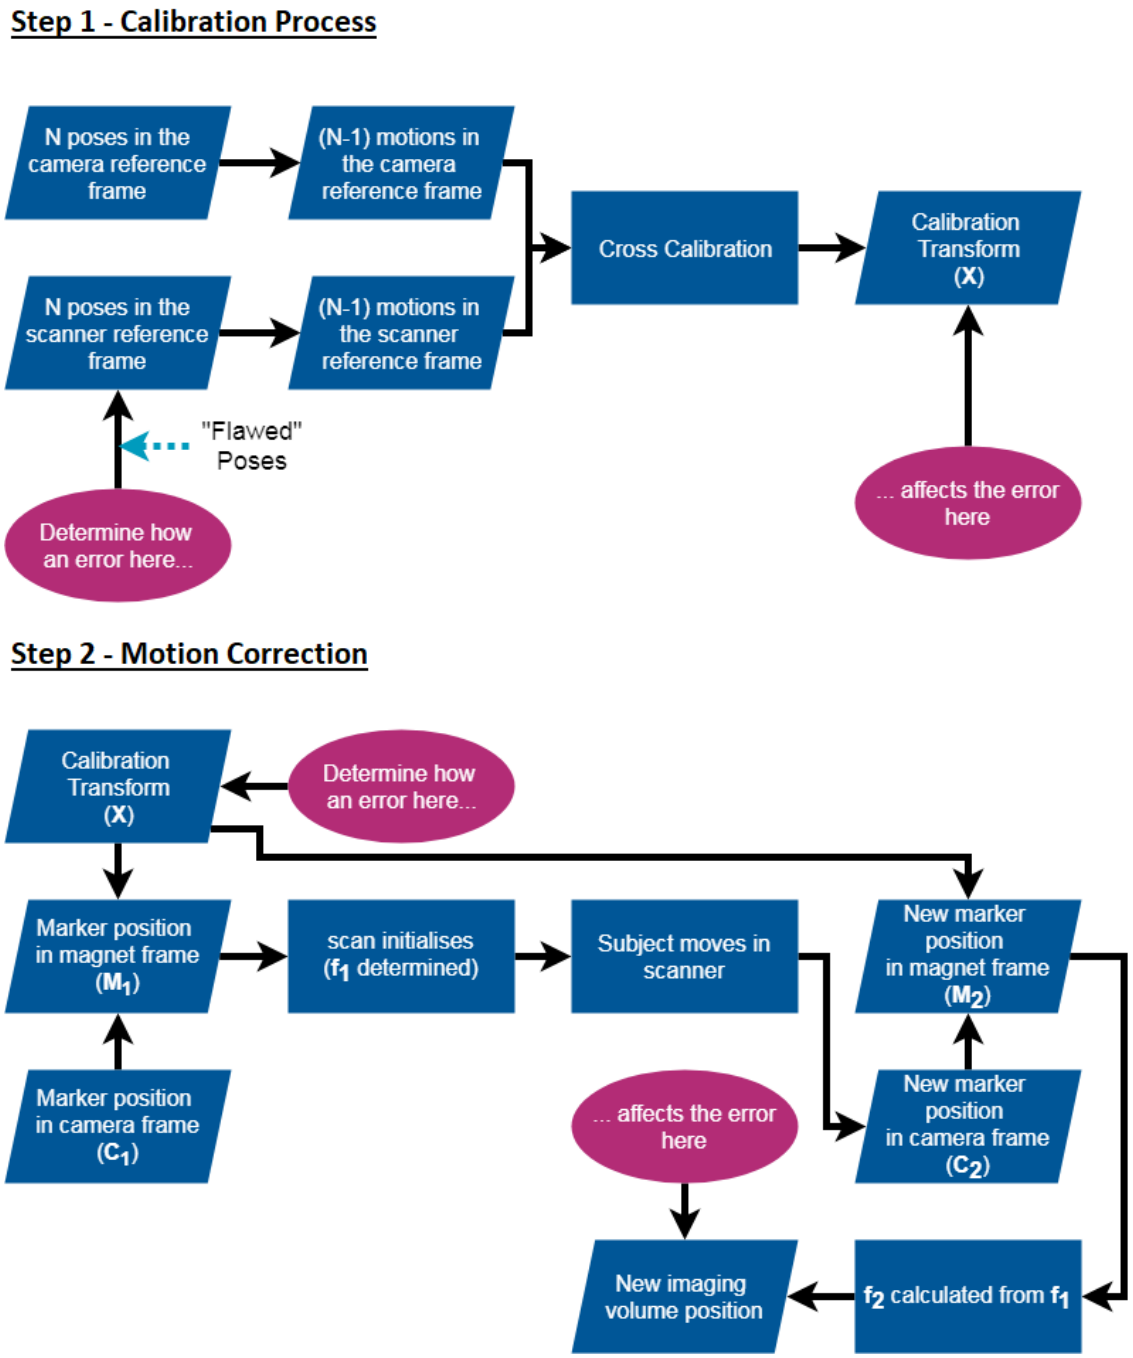


Figure 4.2: A flow chart of the propagation of errors from determining calibrations poses to the calibration transform and from calibration transform to imaging volume position. **X**, **C** **M** and are dual quaternions representing the transform between isocentre and camera, camera and marker (in the camera's reference frame) and isocentre and marker (in the magnet's reference frame) respectively. **f** is a vector describing the translation from marker to imaging volume in the magnet reference frame.

4.2 Theory

The hand-eye calibration transform is estimated experimentally using the non-iterative dual quaternion method described in *Daniilidis et al.* [73], which built on the work from Chou et al. [75]. This method draws on inherent properties of quaternions to quickly determine the transform.

Homogeneous transformation matrices are arguably the most common method to describe rigid body motion in physics, whereas quaternions are more often used in computer graphics design because of the speed at which quaternion maths can be performed. Daniilidis describes a modification to quaternions by combining them with dual numbers to form dual quaternions; which contain the same information as a 4x4 rigid body transformation matrix, but are a more compact form of transformation matrices (8 parameters rather than 16) and are therefore computationally more efficent [76] due to the fewer operations needed to perform dual quaternion multiplication compared to matrix multiplication.

Dual quaternions represent a combination of quaternions [77] and dual numbers [78] to encode both rotations and translations in an object. A quaternion is a 4 dimensional complex number and takes the form:

$$\mathbf{q} = w + xi + yj + zk, \tag{4.1}$$

where w,x,y,z are real numbers between -1 and 1, and i, j, k are the fundamental quaternion unit vectors. Generally w is referred to as the scalar component of a quaternion while x,y,z are referred to as the vector components. Quaternions obey the criteria :

$$\mathbf{q}^T\mathbf{q} = 1 \text{ and } \mathbf{q}^T\mathbf{q}' = 0, \tag{4.2}$$

where $\mathbf{q}^T$ represents the transpose of a quaternion and $\mathbf{q}'$ is the inverse.

Dual numbers have the form of:

$$z = r + d\epsilon \text{ with } \epsilon^2 = 0. \tag{4.3}$$

Dual numbers arise from linear algebra where ϵ describes an extension of a real

number into 2 dimensions (in a similar way that i extends a real number into the complex domain). Because $\epsilon^2 = 0$, ϵ is described as nilpotent. In Equation 4.3, $\mathbf{r}$ is the real part and $\mathbf{d}$ is the dual part.

In the case of dual quaternions r and d are quaternions $\mathbf{r}$ and $\mathbf{d}$. $\mathbf{r}$ is a quaternion representing the orientation component of the pose, while $\mathbf{d}$ represents the displacement of the pose which is defined as [79]:

$$\mathbf{d} = \frac{1}{2}\mathbf{t}\mathbf{r}, \tag{4.4}$$

where $\mathbf{t}$ is a quaternion containing the translational position of the object, where the scalar component is set to 0 and the vector components are the position of the object in 3D space.

A pose in the magnet reference frame is defined as the position of a marker relative to the gradient isocentre, while a pose in the camera's reference frame is defined relative to the aperture of the camera. Poses in the magnet frame can be obtained by transforming the pose in the camera reference frame by the calibration transform as seen in Figure 4.1 and defined as,

$$\mathbf{M}_{pose} = \mathbf{X}\mathbf{C}_{pose}, \tag{4.5}$$

where $\mathbf{M}$ and $\mathbf{C}$ are dual quaternions describing a pose in the magnet and camera reference frames, respectively, while $\mathbf{X}$ is a dual quaternion representing the cross calibration transform. The relation between a motion in the camera reference frame and a motion in the magnet reference frame is described as,

$$\mathbf{M}_{motion} = \mathbf{X}\mathbf{C}_{motion}\mathbf{X}^{-1}. \tag{4.6}$$

According to Chasles's Theorem [80] any rigid body transform can be described as a rotation about a defined axis in space and a translation along that axis; this type of motion is known as a "screw" motion. Screw motions are defined by 6 parameters, 4 describing the infintely long axis in space, 1 describing the rotation about this line and another describing the translation along this line (pitch). It can

be shown that a dual quaternion describing a screw motion can be written as [73],

$$\mathbf{q} = \frac{cos(\frac{\theta+\epsilon d}{2})}{sin(\frac{\theta+\epsilon d}{2})(\mathbf{l}+\epsilon\mathbf{m})}, \tag{4.7}$$

where $\mathbf{q}$ is a dual quaternion representing a screw motion. θ is the angle of rotation, d is the pitch (the translation along the line $\mathbf{l}$, $\mathbf{l}$ is a vector representing a line about which the pose is rotated and $\mathbf{m}$ is the moment describing the orientation of the line. Equation 4.7 demonstrates that for dual quaternions pitch and angle can be separated from the line information. *Daniilidis et al.* [73] showed that by separating out the components of the screw motion in the camera reference frame, and the corresponding motion in the magnet reference frame they can be written as a 6×8 matrix ($\mathbf{S}$). *Chen et al.* [80] showed that at least two non-parallel motions are needed to compute the hand eye calibration, where each motion has the same form ($\mathbf{S}$) . These matrices can be combined into a matrix $\mathbf{T}$ where:

$$\mathbf{T} = \begin{pmatrix} \mathbf{S}_1 \\ \mathbf{S}_2 \\ \vdots \\ \mathbf{S}_n \end{pmatrix}. \tag{4.8}$$

Singular value decomposition (SVD) of ($\mathbf{T}$) gives 8 singular values and their corresponding vectors. At least two of these singular values should be close to zero: the corresponding right singular vectors ($\mathbf{v_7}$ and $\mathbf{v_8}$) which span the null space of $\mathbf{T}$. In the absence of noise, $\mathbf{v_7}$ and $\mathbf{v_8}$ will equal exactly zero; if all the motions performed are exactly parallel then ($\mathbf{T}$), will have rank 5.

$\mathbf{v_7}$ and $\mathbf{v_8}$ can then be split into 4 4×1 vectors where $\vec{v}_7'$ = ($\vec{u}_1'$, $\vec{v}_1'$) and $\vec{v}_8'$ = ($\vec{u}_2'$, $\vec{v}_2'$).

Equation 4.9 can be derived using the constraints of unit quaternions (Equation 4.2),

$$\lambda_2^2(s^2(u_1^T v_1) + s(u_1^T v_1 + u_1^T v_1) + u_2^T v_2 = 0, \tag{4.9}$$

where $s = \lambda_1/\lambda_2$. In the absence of noise there is only one solution for s but in

practise there are two. Daniilidis proposes using only the larger value of s .

Using the value of s, λ_1 and λ_2 can be calculated. The calibration transform is defined as:

$$\mathbf{X} = \lambda_1 v_7 + \lambda_2 v_8. \tag{4.10}$$

4.3 Method

4.3.1 Determining Camera Calibration

The experimental process to determine the calibration involves scanning a phantom in multiple poses while simultaneously recording its position with the camera [81]. The relative movements between each pose can then be used to determine the calibration between the camera's and the scanner's frame of reference as outlined in Section 4.2.

The phantom contains 4 water filled spheres positioned unique distances away from each other creating a tetrahedral shape to remove all symmetry. The pose of the phantom in the MR frame is determined from the MR images. First, a binary mask is created from the MR images by thresholding at 0.15 of the maximum pixel intensity. The binary mask is refined by removing any regions that are smaller than 100 pixels, leaving only 4 regions corresponding to each ball. The position of the phantom is then determined by calculating the total centre of mass using,

$$\hat{g} = \frac{1}{N} \sum_{i=1}^{N} r_i, \tag{4.11}$$

where N is the number of pixels in the binary image with the value 1, i represents each individual pixel, r represents that pixel's position and $\hat{g}$ is the total centre of mass. The orientation of the object can then be determined by calculating the centre of mass of each distinct region of the binary mask g_i, where each region represents one ball. Because the separation between the balls are unique, the distance from $\hat{g}$ to each g_i is also unique. All movements by the phantom are rigid

body movements and so the relative distance between balls does not change and can be used to determine the movement between poses. The position of a ball in the second pose is described as,

$$\mathbf{m_i} = \mathbf{R}\mathbf{n_i} + \mathbf{t} + \xi, \tag{4.12}$$

where $\mathbf{m_i}$ represents the new position of each ball, $\mathbf{n_i}$ the original position of each ball, $\mathbf{R}$ is a 3x3 rotation matrix, $\mathbf{t}$ is a vector describing the object translation and ξ is some error. The translational component of the movement is the difference between the total centre of mass before and after movement. Equation 4.12 can therefore be solved for $\mathbf{R}$ by removing the translations (equivalent to setting the object to the origin, where rotations result in no additional translations). $\mathbf{R}$ can be found by minimising the difference between m_i and $\mathbf{R}n_i$,

$$\Sigma^2 = \sum_{i=1}^{N} |m_i - \mathbf{R}n_i + \mathbf{t} + \xi|, \tag{4.13}$$

or,

$$\Sigma^2 = \sum_{i=1}^{N} m_i^T m_i + n_i^T n_i - 2m_i \mathbf{R} n_i. \tag{4.14}$$

The corresponding poses in the camera frame can be determined from the marker which allows for tracking of position by the camera. In practice, it is difficult to measure the phantom position using the scanner at the same time as the marker position with the camera, because the coil blocks line of sight to the camera. Often a pose must be measured in one modality (camera or scanner) and then later acquired in the other under the assumption the the phantom has not moved (i.e. the phantom is in the same pose). However, the time difference between acquiring the pose of the phantom in the camera frame and acquiring the same pose in the magnet reference frame means that movement is possible, especially when the coils are moved to allow the camera to view the phantom. This creates an error between corresponding poses in the MR frame compared to the camera frame. These pose errors can then result in errors in the cross calibration transform. The errors in the calibration transform can then generate

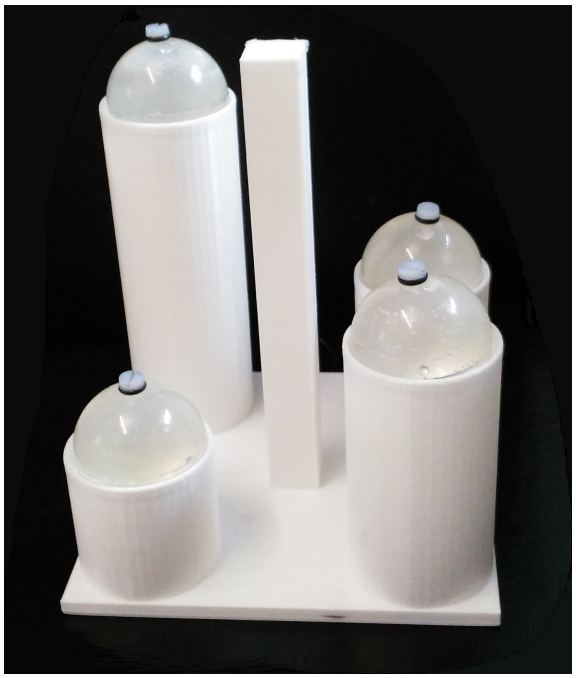


Figure 4.3: The phantom used for determining the camera cross calibration. The phantom consists of 4 water filled spheres arranged in a tetrahedron shape with no axis of symmetry

erroneous motions in any subsequently recorded movements when transforming to the magnet reference frame.

4.3.2 Simulation of Camera Calibration Error

A Monte Carlo simulation was performed to evaluate the impact that errors in calibration poses determined from the MRI scans have upon the resulting calibration transform. N poses in the camera reference frame were defined for $N = 4, 6, 8, 10, 12$. The simulated poses were centred about a point near the magnet isocentre. The movement from one pose to the next was limited to be less than $\pm$ 40 mm in any one direction and the rotations were limited to be no greater than $\pm$ 30° around x and y, and unrestrained ($\pm$ 180°) in z (camera frame). These were selected to reflect typical phantom poses possible in our experimental set up where it is difficult to rotate the phantom in the camera's x and y axis without obscuring the marker, and the overlapping field of view between the scanner and camera limits the size of the translations (see Figure 4.4).

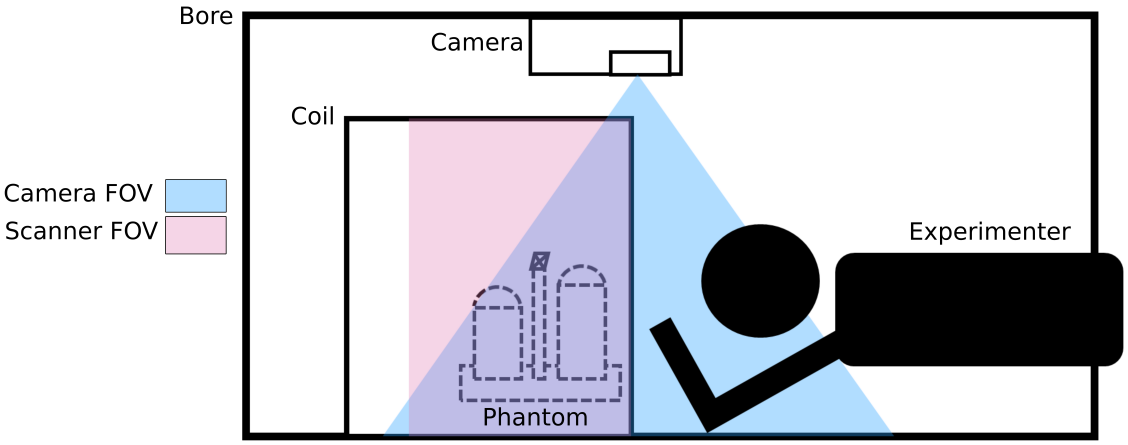


Figure 4.4: A diagram of the experimental process of determining the cross calibration transform. To record the pose of the phantom the phantom must be visible in both the scanner's field of view (pink) and the camera's field of view (blue).

A typical 7 T calibration was defined and was allowed to vary in each iteration by no more than 50 mm and 20°; this represents all reasonable positions one could place the camera within the bore and still be in a position to track the marker. In practise it was found that the calibrations movements vary far less than this; by only a few degrees or mm between placements. Each pose in the camera frame was then converted to a corresponding pose in the magnet frame using Equation 4.5.

Poses were converted from dual quaternions to translations and Euler angles (zyx) defining a set of N true calibration poses. Random errors were then added to either the translational (ϵ_T) or rotational components (ϵ_R). Errors were added to either one pose in a set of N poses (excluding the first pose), or all poses in N creating a set of N "flawed" calibration poses (see Figure 4.2). Random errors had standard deviations of up to 10 mm in each translational direction and up to 5 degrees in each rotational axis.

The motions between the first calibration pose and each subsequent pose of the flawed pose set were then determined, generating $(N-1)$ "flawed" motions. In addition the corresponding motions of the true magnet poses (without errors) were calculated. Motions for the poses in the camera frame were then calculated. This generated a set of flawed and true movements in the camera reference frame

and a corresponding set of movements in the magnet reference frame.

These motions were used to determine a “flawed” calibration transform. The root mean square error (RMSE) between the flawed calibration and true calibration was measured. The mean translational and rotational error was then recorded over a total of 1000 iterations for each pose data set size, error type and error size.

Calibration Window

In practice the calibration values must be within certain limits set by the MR environment. For example, the calibration cannot have an x or y translational component greater than half the width of the bore and still be within the magnet. Assuming an approximate calibration can be determined, an upper limit was set on calibration error of 145 mm and 45°, where if the error exceeded this, the calibration was said to have failed and the iteration would be repeated. This represents the user’s ability to identify failed calibrations and repeat experiments as necessary. The upper limits of the calibration error were chosen to represent the highest level of error that a user may mistake a “flawed” calibration for an accurate calibration. The total number of successful iterations was maintained at 1000 and the number of failed iterations was counted. The percentage of failed calibrations was then be determined where the total number of iterations was 1000 plus the number failed calibrations.

4.3.3 Calibration Error Propagation

To correct imaging volume geometry there are 3 successive transforms. Isocentre-to-camera, camera-to-marker and marker-to-volume. The isocentre-to-camera is the calibration transform ($\mathbf{X}$) and is assumed to contain the largest source of error and potentially has the largest impact to imaging volume location because it is the first transform. The camera-to-marker transform ($\mathbf{C}$) is obtained directly from the camera data, the errors associated with this are dependent on the tracking algorithm but for a Moiré phase tracking camera are quoted to be 0.1 mm and

0.1°. For the marker-to-volume transform ($\mathbf{F}$) only the translational component is needed ($\mathbf{f}$). If the marker's position is described in the magnet reference frame then a change in orientation of the marker is identical to the change in orientation of the imaging volume. The vector ($\mathbf{f}$) can be rotated as the marker rotates to describe the position of the imaging volume from the marker.

An error in the calibration transform propagates through each of these transforms, potentially increasing error size. A series of head rotations and translations were simulated and the subsequent impact that the calibration error had upon the new apparent position of the imaging volume after motion correction was investigated. A typical calibration for a camera placed inside the bore of a 7 T scanner was defined Tx = 50 mm, Ty = -250 mm, Tz = -150 mm, Rx = 0°, Ry = 0 °, Rz = 160 °. A series of head motions were also defined, a pure translation in the scanner x direction up to 30 mm and a pure rotation about x (head nod) up to 10 degrees. The imaging volume was simulated to initially be at isocentre with zero orientation, and the initial marker position just below that (Tz = -150 mm) in the FH direction to simulate its approximate position if the subject were to wear a bite bar.

Head motions were converted into marker movements assuming a rotation about isocentre. The marker position was then converted to the camera reference frame by applying the inverse calibration,

$$\mathbf{C} = \mathbf{X}'\mathbf{M}. \tag{4.15}$$

Two marker positions in the camera reference frame were generated, one before movement and one after movement (C_1 C_2). These parameters represent the output of the camera. Errors were added to the calibration. The camera poses were converted to the magnet reference frame using both the true calibration (X_T) transform and the flawed calibration (X_F) transform to generate 4 marker positions $\mathbf{M_{1F}}$ $\mathbf{M_{2F}}$ $\mathbf{M_{1T}}$ $\mathbf{M_{2T}}$ where 1 and 2 indicate the before and after poses and T and F indicate the true and "flawed" calibration used to determine that pose. The apparent transform from marker to imaging volume ($\mathbf{f}$) was calculated using the marker positions $\mathbf{M_{1F}}$ and $\mathbf{M_{1T}}$. This transform is determined using

the initial pose which assumes the location of the marker and the imaging volume are known. This creates one $\mathbf{f}$ for each calibration because, although the initial position of the imaging volume is independent of the calibration transform, the apparent position of the marker varies depending on the calibration error.

A rotation of the marker as described in the magnet reference frame is identical to the rotation of the imaging volume in the magnet reference frame so rotational transform from marker to volume is not needed, just the translational component. The rotation between $\mathbf{M_1}$ and $\mathbf{M_2}$ was calculated for both the true and the flawed poses. This rotation was taken as the rotation of the imaging volume. Each $\mathbf{f}$ was rotated by the apparent rotation of the marker and then added to the translational component of the marker to give the apparent position of the imaging volume, $\mathbf{v_{2T}}$ and $\mathbf{v_{2F}}$. The Euclidean separation from the true volume position to the flawed volume location was then measured, as well as the RMS error in volume orientation.

This process was then repeated using several calibration error sizes in both the rotational components and translational components to determine the relationship between error in imaging volume and calibration transform for a several of head movements.

Using this relationship the maximum permitted error in calibration to yield imaging volume positions accurate to 0.5 mm and 0.5° was determined.

4.3.4 Assessing Calibration Reliability

There is often no ground truth to the calibration transform and thus it can be difficult to determine if a calibration you have obtained is reasonable. Calibration errors can be of the order of tens of metres which can be identified by eye. However calibration errors of tens of mm can be difficult to spot. We theorise that in instances where the calibration has failed, the calibration will be unstable, i.e. altering the poses within the calibration pose data set would result in radically different calibrations.

To test this a set of calibration poses in the camera and the scanner reference

frame were generated. In one instance, these calibration poses had relatively small noise-like errors added to the rotational and translational components (0.1° and 0.2 mm). This data set represents typical results from a 'good' calibration. In another instance, the same poses had large errors added to them (1.5° and 6 mm); these represent a bad calibration data set where the error is large enough to cause large calibration errors of the order of 15 mm. These errors are large enough to cause errors in the prospective motion correction, but small enough to not be identified by the experimenter.

Each of these calibration data sets underwent a bootstrap simulation. In this simulation, a random subset of poses was selected from the larger calibration pose set. The subsets were selected to always contain the first pose of the full set, and a random assortment of all other poses. The number of poses used ranged from 4 to 16 poses; in each instance, 32 permutations of the random subset were used to determine the apparent calibration. Each of the resulting calibrations was then split into its translational and rotational components. The variation in each of the calibration components was then used to measure the stability of the calibration.

4.4 Results

4.4.1 Simulation of Camera Calibration Error

Using the Monte Carlo simulation, the calibration error was simulated as a function of pose error, for data sets containing 4 to 12 poses. Figure 4.5 shows the calibration error when **one** of the poses in the data set contains an error. There is a strong linear relationship between the size of the pose error and the resulting error in calibration. A first order polynomial line of best fit was plotted to show this relationship. The error bars represent the standard deviation over 1000 successful iterations.

This linear relationship was also observed in Figure 4.6 where **all** the poses contained errors. As expected, the size of the calibration errors in this case were

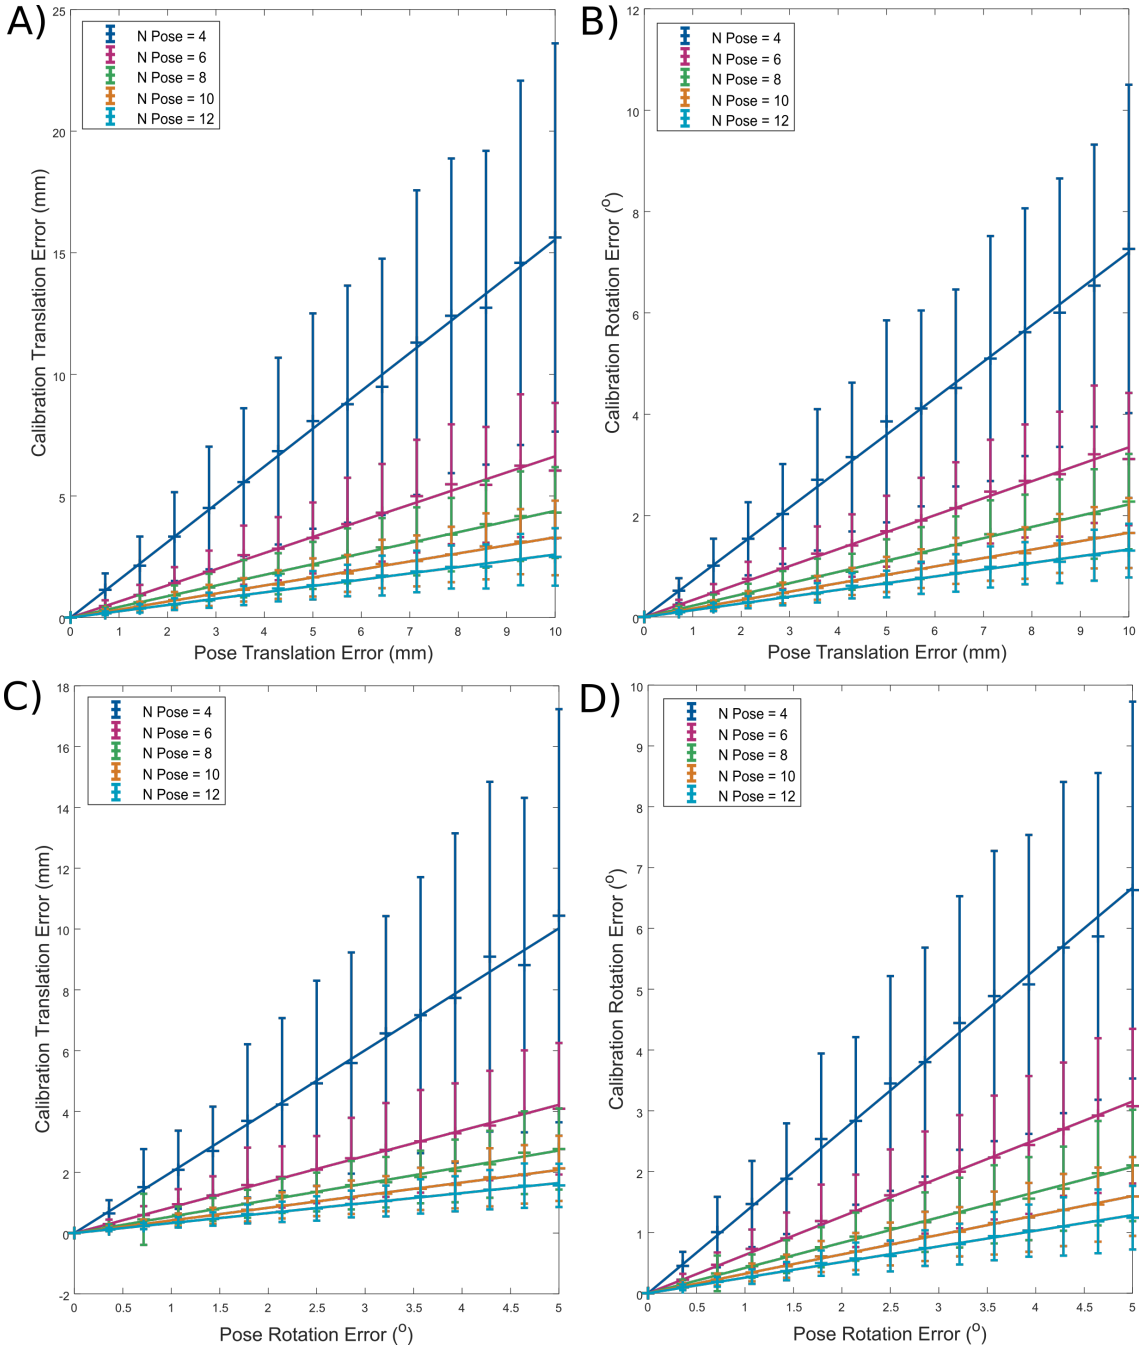


Figure 4.5: The calibration error (both translational (A, C) and rotational (B, D)) as a function of pose error for a set of of N poses where **one** pose contains either translational (A, B) or rotational error (C, D).

much larger than those observed for data sets containing one error.

In all cases, the number of times a calibration had to be repeated because the calibration lay outside of the expected window, grew with the size of the pose error (Figure 4.7). However, the percentage of failed calibrations was much greater for

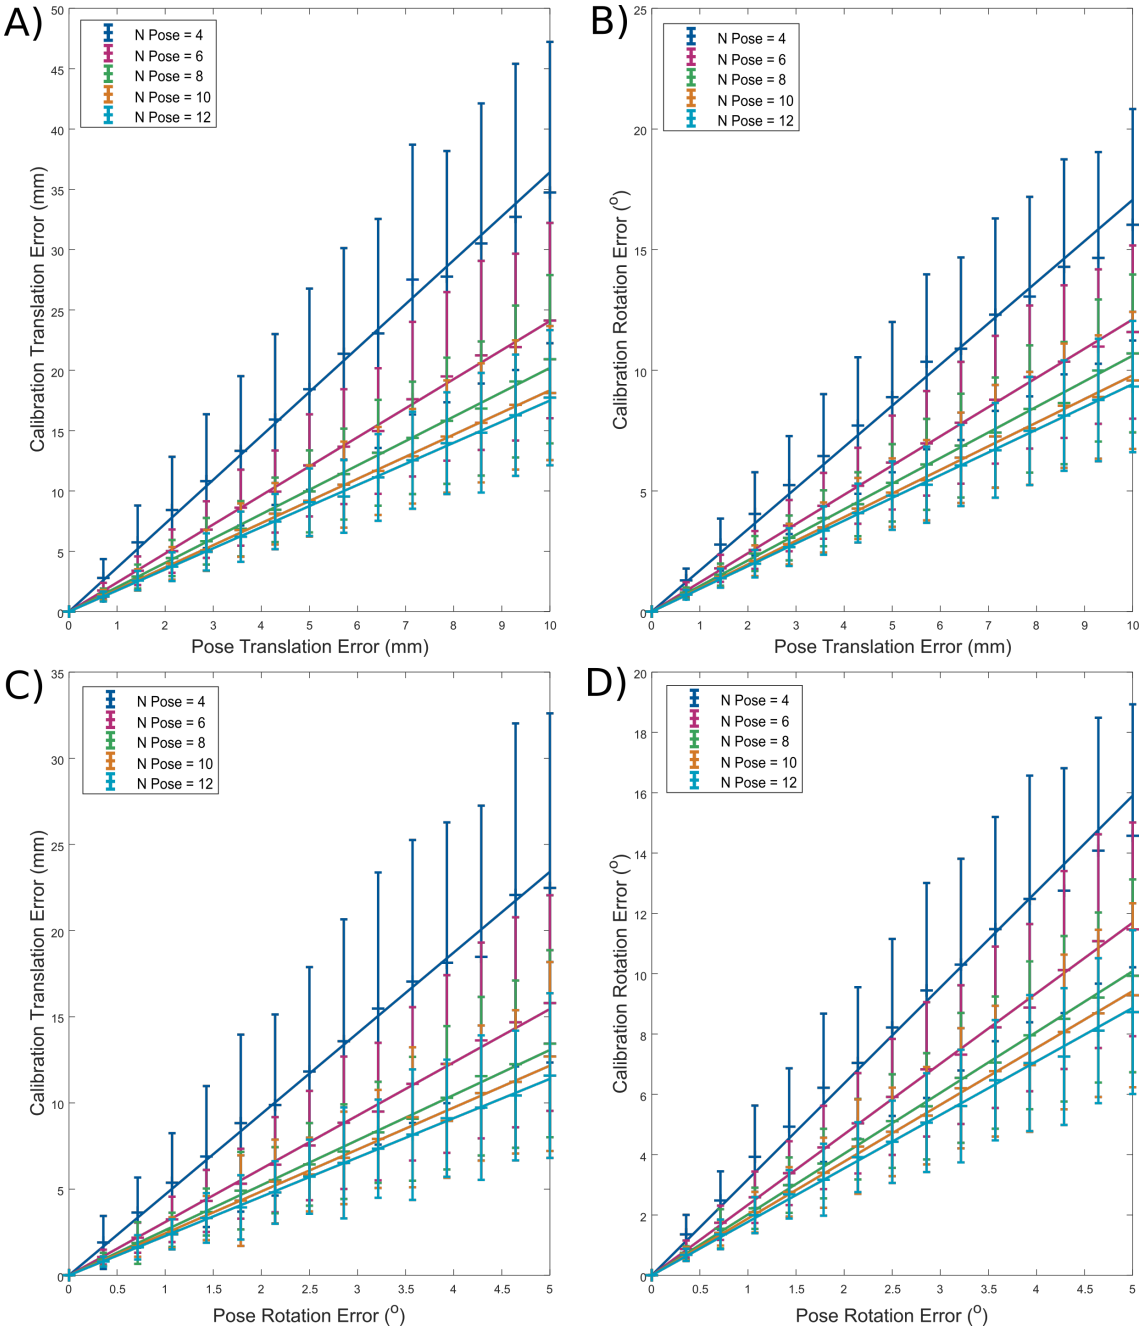


Figure 4.6: The calibration error (both translational (A, C) and rotational (B, D)) as a function of pose error for a set of of N poses where **all** poses contain either translational (A, B) or rotational error (C, D).

rotational errors than it was for translational errors. It was observed that in the case of high error calibrations there would also be a rotational error of 180° along one axis suggesting that the calibration was flipped along that axis. These results were excluded out from Figures 4.5 and 4.6 by the calibration window.

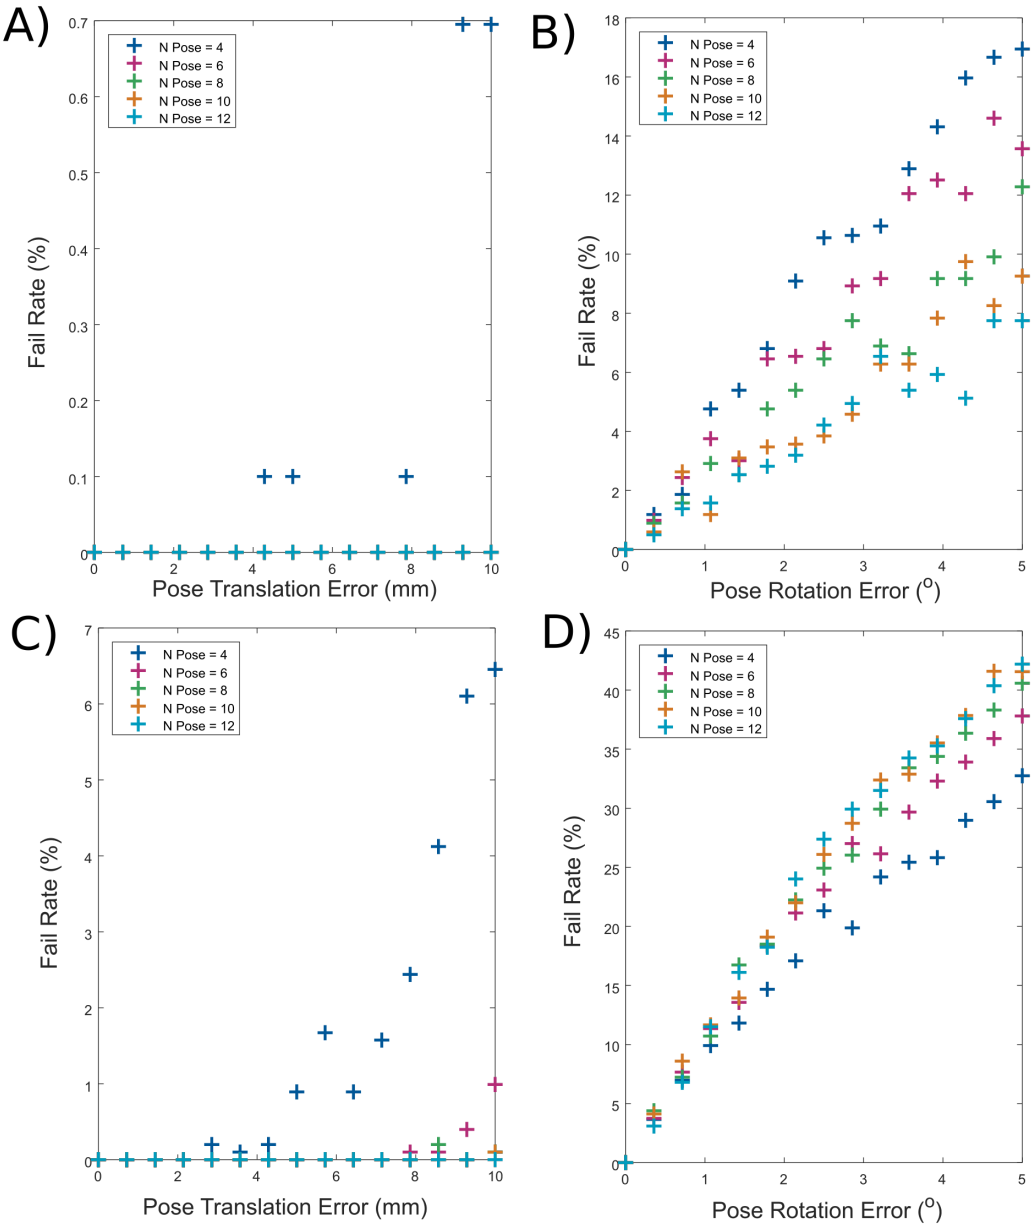


Figure 4.7: The percentage of failed calibrations as a function of pose translational error (A,C) and pose rotational error (B,D) for a single flawed pose (A, B) and a full set of flawed poses (C, D).

4.4.2 Calibration Error Propagation

Errors in the calibration propagate through to errors in the apparent imaging volume. The errors in the imaging volume are dependent upon the type and size of error in the calibration and the size of the movement. Figure 4.8 shows the errors

in imaging volume position and orientation as a result of calibration error in the translational and rotational x direction, for a series of head translations of different amplitude. Marked on each graph is the pose calibration error to achieve volume position and orientation accuracy to 0.5 mm and 0.5°. They show that only errors in the calibration rotational components affect the position of the imaging volume. However the orientation of the imaging volume remains unaffected. Translational errors in calibration had no effect on the imaging volume's apparent position or orientation.

Figure 4.9 shows similar results to Figure 4.8 but a series of head rotations of different amplitudes. These results show that head rotations are susceptible to errors that the translations were immune to. Only translational errors in calibration do not propagate to the imaging volume orientation. In all other cases error in the calibration leads to errors in the imaging volume position and orientation.

Using Figures 4.8 and 4.9, the maximum calibration errors permitted to achieve accuracies of $\pm$0.5 mm and $\pm$0.5° in the imaging volume are shown in Table 4.1. The required pose accuracy to determine the calibration of this accuracy is also shown.

Calibration Accuracy	2.0 mm	0.4°
Minimum Calibration Pose Accuracy	1.0 mm	0.2°
Maximum Calibration Pose Error	2.5 mm	1.3°

Table 4.1: A table of the maximum errors in the calibration to give imaging volume position accurate to to $\pm$0.5 mm and $\pm$0.1 °. Also shown are the pose accuracy required to achieve this calibration accuracy using an 8 pose data set.

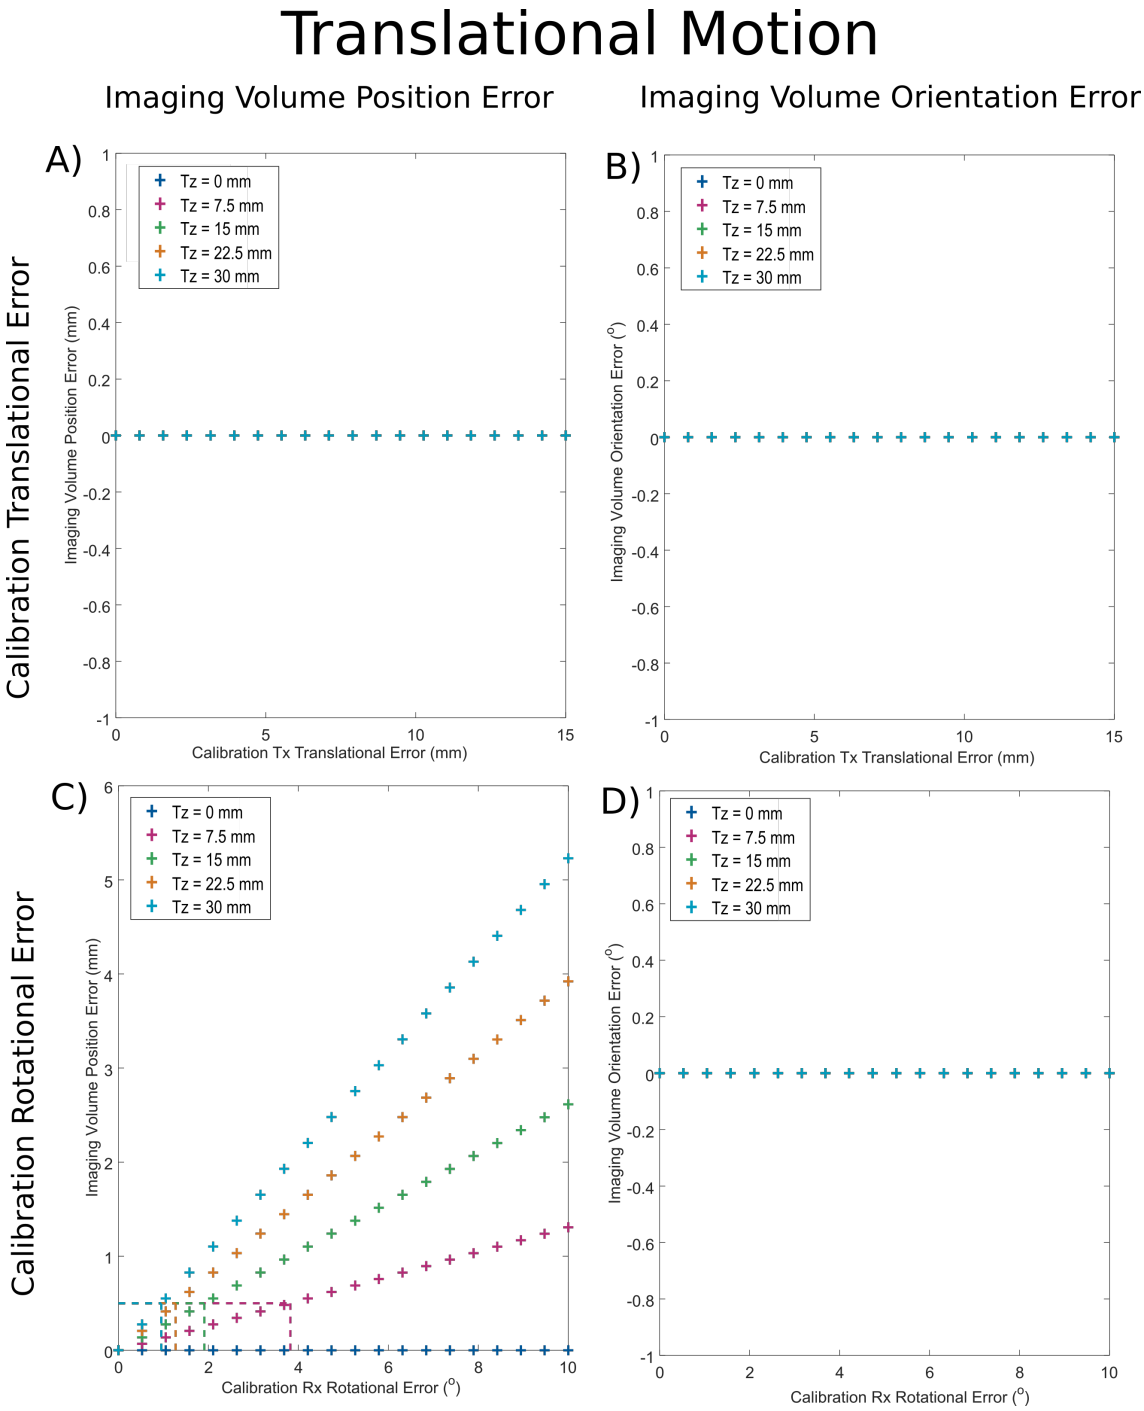


Figure 4.8: The measured error in a simulated imaging volume's position (A,C) and orientation (B,D) using a calibration that contains translational (A, B) and rotational (C,D) errors in calibration. Several head translations along z were simulated from 0 mm to 30 mm .

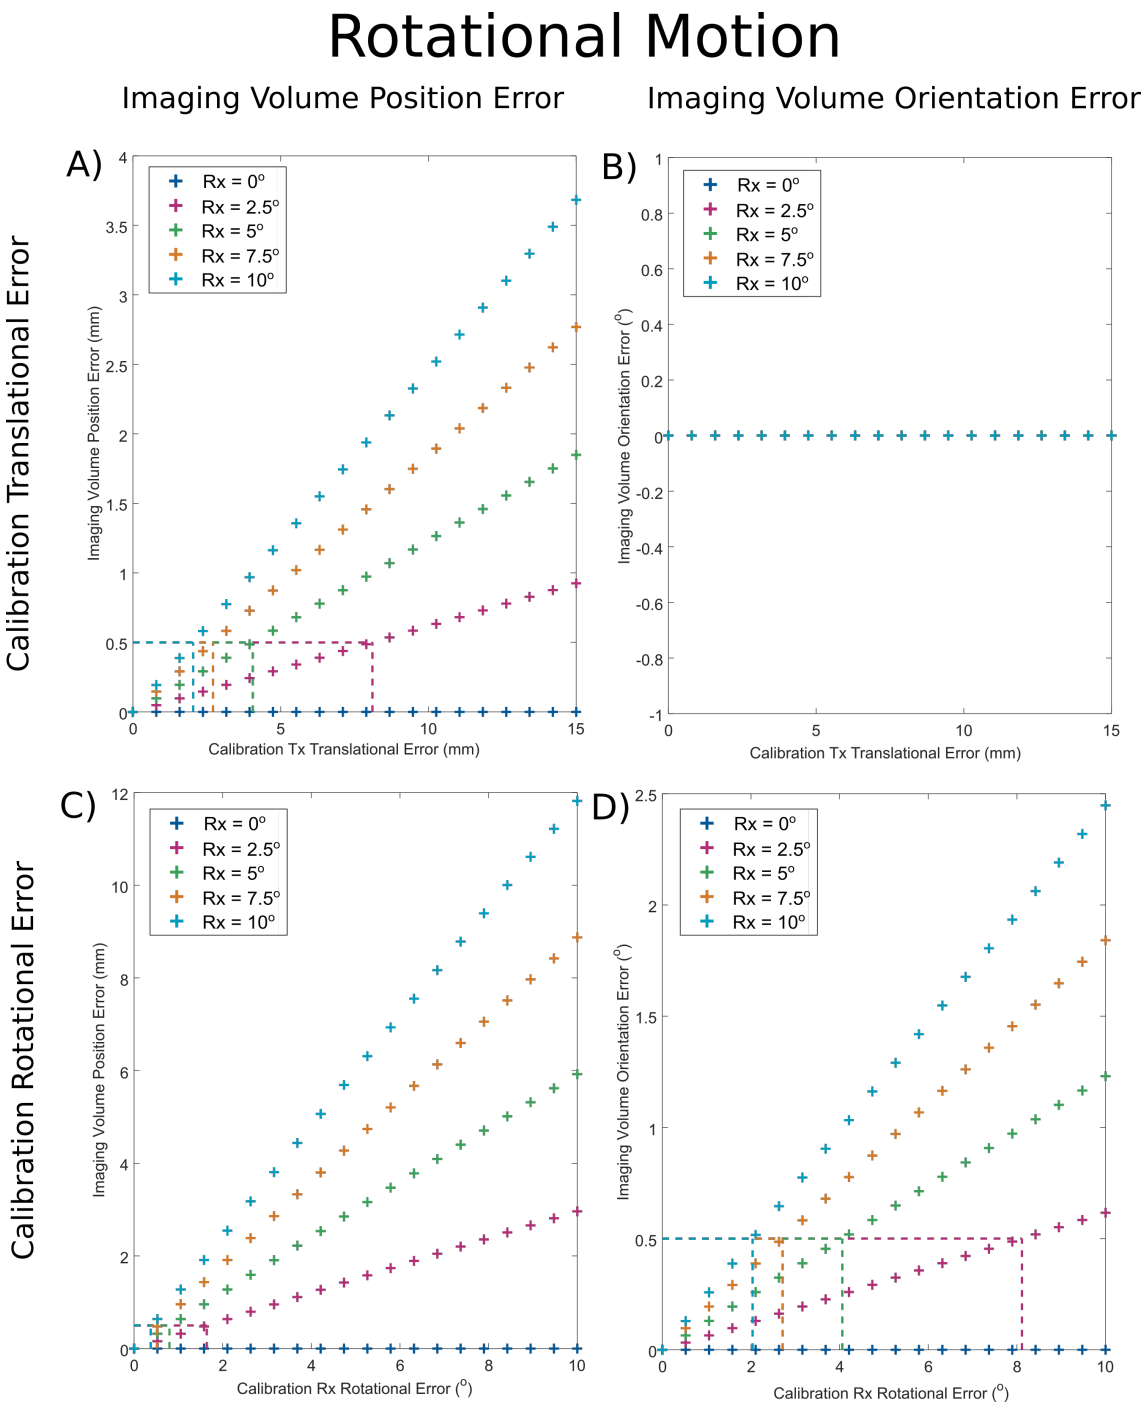


Figure 4.9: The measured error in a simulated imaging volume's position (A,C) and orientation (B,D) using a calibration that contains translational (A, B) and rotational (C,D) errors in calibration. Several head rotations about x were simulated from 0°to 10°.

4.4.3 Assessing Calibration Reliability

For each iteration of the bootstrap simulation a calibration was determined with 6 components. The variation of each of those 6 components is shown in Figures 4.10 and 4.11.

Figure 4.10 shows the variation in the calibration as a result of small errors in the calibration poses. These errors were small enough to give a relatively stable calibration where the maximum variation for an 8 pose calibration is approximately 1 mm and 0.4°. The calibration determined using all 16 poses was also accurate.

Figure 4.11 shows the variation in the calibration as a result of large errors in the calibration poses. In this case variations were of the order of 30 mm and 8°. In the event of large errors removing a pose from the data sets result in a large difference in calibration.

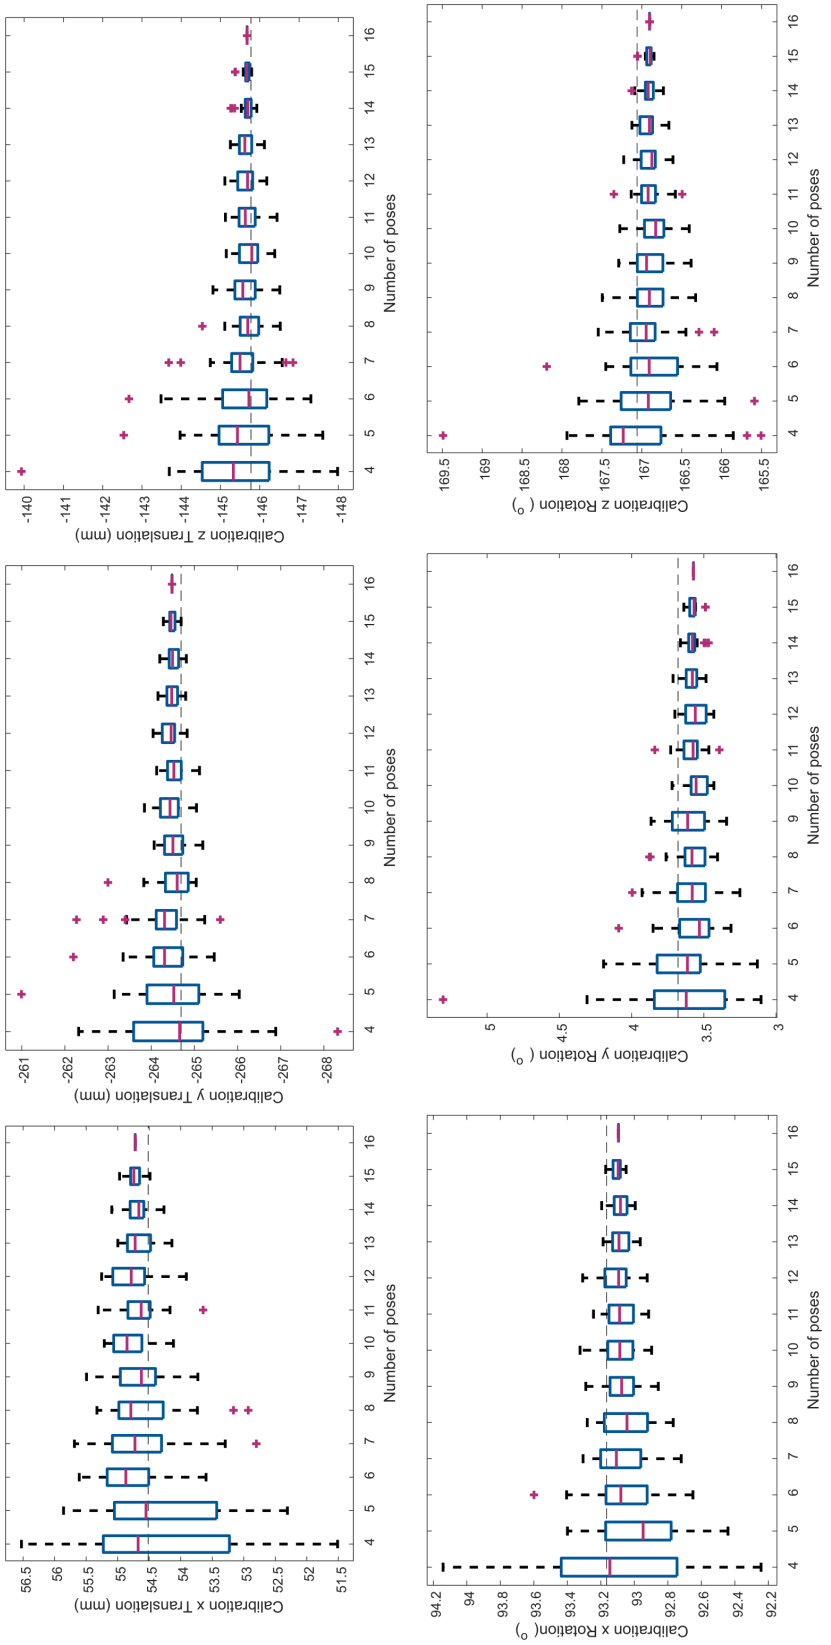


Figure 4.10: The calibration parameters returned from a bootstrap simulation using poses that contain **small** errors. The dashed line shows the true value for that parameter.

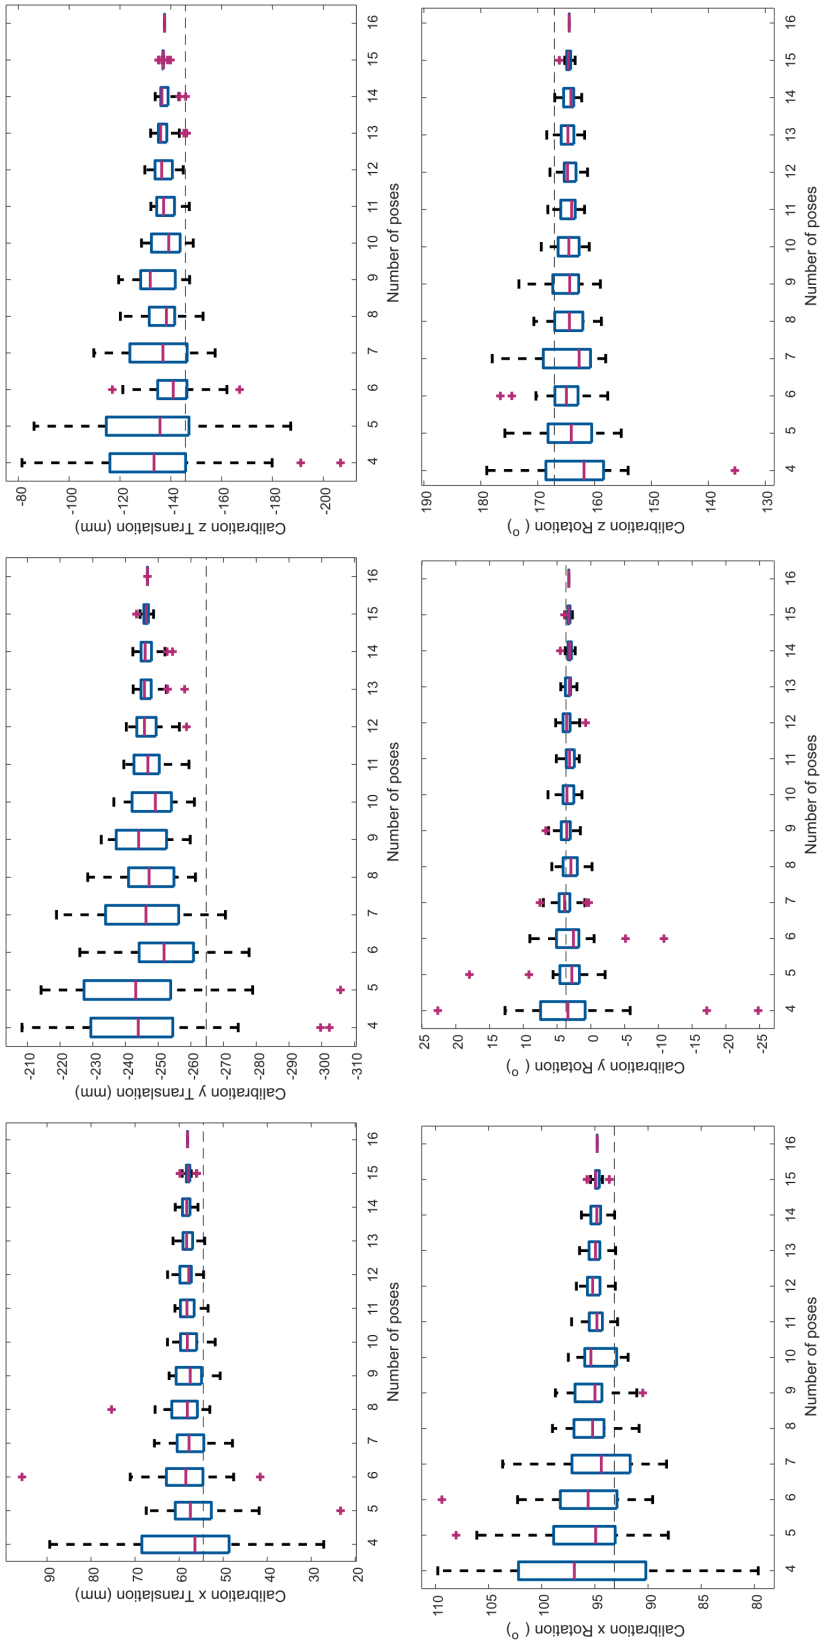


Figure 4.11: The calibration parameters returned from a bootstrap simulation using poses that contain **large** errors. The dashed line shows the true value for that parameter.

4.5 Discussion

4.5.1 Simulation of Camera Calibration Error

The strong linear relationship between pose error (rotational and translational) and calibration error (rotational and translational) is shown in Figure 4.5. This relationship held for data sets containing a single flawed pose and data sets containing entirely flawed poses. It is worth noting that this relationship assumes one can identify and repeat calibrations where the calibration algorithm failed, because of this calibrations with very high errors were excluded from the average.

For both the fully flawed and single flawed pose data sets the calibration error was reduced by including more poses but the increase in accuracy seemed to vary non-linearly, with there being a large increase in accuracy by moving from 4 poses to 8 poses, but a minimal increase moving from 8 poses to 12. However, additional poses reduced the standard deviation across the calibrations, suggesting that a higher number of poses may lead to a more stable calibration. This was most noticeable in the single flawed pose simulation, but was also true for the fully flawed pose calibration.

As expected, adding errors to all the poses resulted in a higher error in calibration than just adding error to a single pose. Typically the errors seen in a fully flawed calibration data set were twice as large as those seen with a single flawed pose.

Results in Figures 4.5 and 4.6 are of calibrations that were within the calibration window. (errors smaller than 145 mm and 45°). It is also useful to determine how often errors exceeded this value. The fail rate of the calibration is shown in Figure 4.7. This shows that the calibration method is actually extremely robust to translational errors in pose with the calibration failing only when all poses contained translational errors greater than 5 mm and only contained a total of 4 poses. Less than 1% of calibrations failed for a single flawed pose of less than 10 mm regardless of how many poses were used to determine the calibration. However rotational errors result in a large number of failed calibrations. This fail-

ure rate reached nearly 50% in the case of many large rotational errors in pose. When these calibrations were investigated the majority failed for not reaching the rotational requirement of being within 45° of the true calibration, rather than the criterea of being with in 145 mm of the true calibration (Most of these calibrations had errors of 180°). This suggests that there is another local minima where the calibration can be thought to contain a 180 degree flip but the translational components remain the same.

In the case of a single flawed pose increasing the number of poses decreased the fail rate, however when all poses contained error there was a noticeable increase in fail rate as a result of using additional poses. The exact mechanism behind this is unknown, however it is possible this is in response to the overall error of the pose set (the sum of all error added to the data set) being higher if more poses are used.

Given the experimental setup to determine the calibration within our 7 T environment, calibration errors are mostly likely to occur along the z axis. These errors are caused by the movement of the top of the coil back and forth between scanning which can potentially cause the entire coil to shift in z (Figure 4.12).

To overcome this, the average position before and after the MR scan is taken for the pose in the camera frame. In the event of large differences between the before and after poses the entire pose is discarded. Figure 4.7 suggests that translational errors are unlikely to result in failed calibrations while Figure 4.5 suggests that the translational error is still likely to increase the calibration error overall.

Figures 4.5, 4.6 and 4.7 suggest that overall increasing the number of poses in your calibration data yields higher accuracy calibrations. The high failure rate suggests that the calibration should not be run unsupervised and the output must always be checked for reliability. These results suggest a calibration consisting of 8 poses is a good compromise between calibration error and time to perform calibration. 8 poses also offers the user some degree of flexibility to discard poses if it is determined that one of them contains a large error.

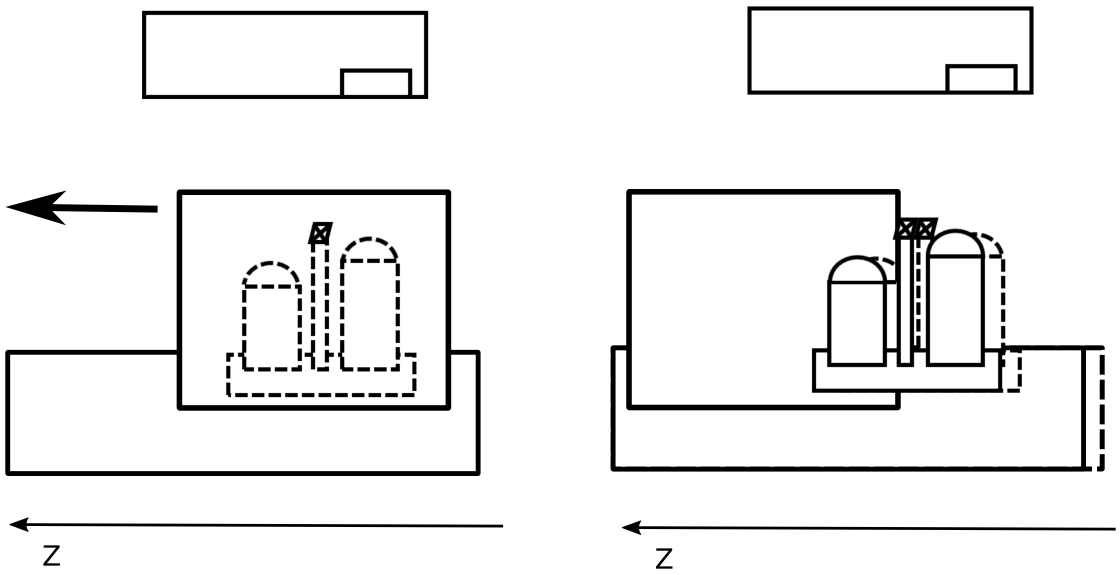


Figure 4.12: When moving the top of the coil so that the camera can view the phantom, the coil is prone to shifting in its entirety. This can cause errors in the position of the phantom along the direction the coil was moved in.

4.5.2 Calibration Error Propagation

The error in calibration was seen to propagate through to errors in the imaging volume for most but not all head movements and error types. No error was observed for head translation in the presence of translational calibration errors. This is likely because the imaging volume is determined based on the change in marker position, while the absolute position (r) of the marker may be different due to the presence of calibration errors. The change in marker position (Δr) does not vary because the marker contains the same error before and after movement. Translational error was not seen to impact the rotational accuracy of a head rotation. This is because the orientation of the co-ordinate scheme is still correct, hence a rotation around x is still a rotation around the same axis. However the position of the head *is* affected by translational error in calibration for head rotations. The translational error in calibration causes an apparent difference in marker position. Since the transform between imaging volume and marker ($\mathbf{f}$) is determined based on the position of the marker, $\mathbf{f}$ is wrongly defined as a result of calibration transform error. When the head is rotated, $\mathbf{f}$ is also rotated, creating a

lever effect positioning the head in the wrong place which in turn results in errors in the imaging volume observed in the simulation (as shown in Figure 4.13).

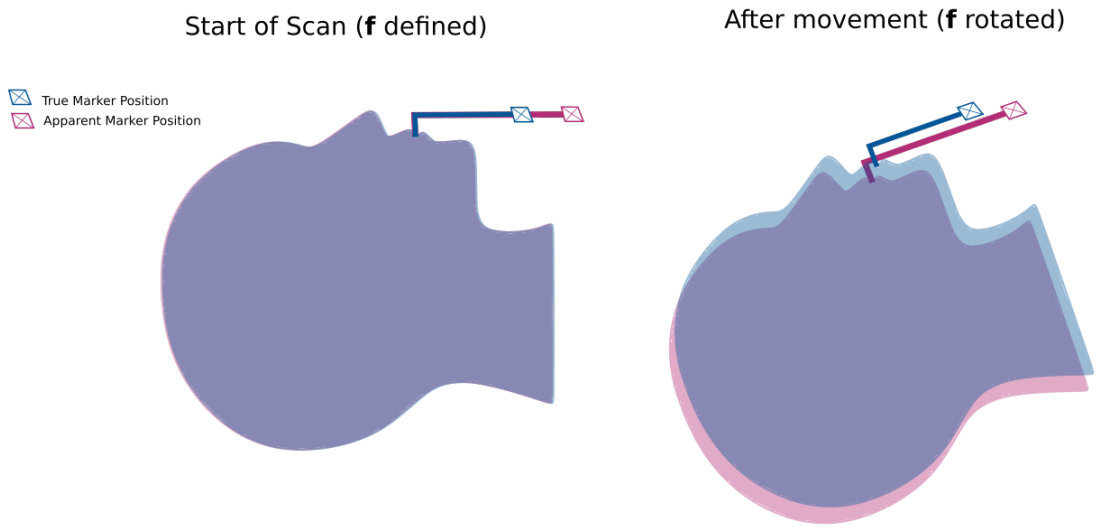


Figure 4.13: The difference in head position as a result of the incorrect starting position of the marker. The $\mathbf{f}$ vector is wrongly defined which propagates to head position when the head rotates. Here it can be seen that the correct locations are in blue, and the false locations are in pink after the marker undergoes the same movement (a rotation of 20°)

4.5.3 Assessing Calibration Reliability

The bootstrap simulation contained one full pose set containing 16 poses without any error. Using the data set with no errors, the calibration gave the exact error free calibration transform even if only a subset of 4 poses was used. The determined calibration varies slightly in the presence of small errors in translation and rotation of the calibration poses. The calibration varied most only when a subset of 4 poses were used. No variation was seen when all 16 poses were used because all motions were relative to the first pose, Pose 1 was always included so while the order of the motions changed, the motions themselves did not change.

The variation of the calibration was much higher for large errors in pose (4.11) reflect this. For example an 8 pose calibration had a variation of 40 mm in the y

direction, while for small errors this variation was only 1.4 mm.

Bootstraping represents an easy method for determining how accurate the calibration transform is. It can be performed at the time of calibration by removing a pose from the calibration set and monitoring how much this changes the calibration transform.

Interestingly, bootstrapping cannot be used to improve the accuracy of the calibration. Even in the case of 16 poses, the median calibration was still significantly different to the true calibration. Instead, the bootstrapping method should only be used to determine the stability of the calibration. In the event of an unstable calibration the experiment should be repeated.

4.5.4 Practical considerations

Calibrating the camera takes approximately 45 minutes and requires at least two people for a 7 T scanner; one to lie inside the magnet and move the phantom and coil, the other to operate the scanner. Because of this it may be tempting to acquire fewer poses. Figures 4.6 and 4.5 suggest while that can be done, it is not recommended as the errors in calibration are likely to increase. The reduced number of poses would also limit the user's ability to perform a bootstrap simulation, which means the stability of the calibration cannot be measured.

Calibrating at lower field strengths, where the bore is shorter and more accessible may be practically easier and reduce the calibration time considerably. At 3 T the transmit coil can be embedded within the bore casing. This means that the camera is not obscured by the coil and simultaneous scanning and camera monitoring can occur during the calibration. Additionally, the smaller bore means that an experimenter can lean into and move the phantom between poses eliminating the need for a second experimenter.

Additional complications occur as a direct result of the field strength at 7 T. Typically the coil and phantom are moved by an experimenter lying on the bed in a face down prone position during the scan, at 7 T this position is known to be particularly disorientating for the experimenter [82, 83] limiting the amount of time

a typical experimenter can comfortably remain in the scanner, and hence the total number of poses that can be gathered for a calibration.

Alternative calibration methods exist including iterative approaches [68,84,85]. These methods offer advantages over the *Daniilidis* method. Firstly an estimation of motion correction accuracy is known from the determination of calibration. Because the iterative calibration aims to reduce the error seen during volume-to-volume prospective motion correction. The error of a prospective motion correction scan then is the residual error during the calibration, which can be further reduced by performing more iterations. Secondly provided the initial estimate of the calibration is sufficient, convergence occurs very quickly, typically taking only 3 to 4 scans. This represents a much faster method of determining the calibration.

Despite its advantages, this method is not suitable for the Philips 7 T environment at the moment. Iterative calibration methods require prospective motion correction to be performed on the images which was previously not available. However with the Geocam patch (see Chapter 5) it may be possible to implement in the future. Amendments would have to be made to the patch to make the process quick. Volume-to-volume motion correction is currently available on the scanner as part of the Philips prospective motion correction, this could be used to measure the residual error after Geocam prospective motion correction. Provided the transform that was applied during the Geocam correction was recorded, the calibration can be updated to reflect the residual error. This method however would require simultaneous acquisition of a phantom in the camera and magnet reference frame so a different phantom would have to be used which could limit the range of movements that could be performed. However it could potentially reduce the number of experimenters required to calibrate. The phantom can be removed from the magnet between volumes, moved and and then placed back into the magnet.

4.6 Conclusion

In this Chapter it has been shown that for a non-iterative cross calibration method, 8 poses represents a reasonable compromise between time to acquire scans and calibration accuracy. The impact that pose determination has upon the cross calibration in the presence of translational errors and rotational errors for both one pose and all poses within a calibration data set has been detailed. For small errors, the error in calibration varied linearly with the error in pose for translation and rotation. However, as error increases the fail rate (the number of calibrations outside of the expected calibration window) of the calibration rapidly climbs. In many cases these calibration failures generate unrealistic calibration values, that could easily be spotted by the user.

The effect that calibration error has on determining the position and orientation of the imaging volume, using methods one might see during prospective motion correction have been described. These results suggest a linear relationship between error in imaging volume position resulting from error in in calibration for head nods and translations. These errors increase as the size of the head movement increases. From these results one can estimate the required calibration transform accuracy to achieve a given accuracy of imaging volume position. For large head movements it was determined that calibration error must be no greater than 2.0 mm and 0.4° for large head movements.

A novel method for determining the reliability of the calibration transform has been shown. This method, which requires no additional data to be acquired, measures the stability of the calibration transform using a bootstrap simulation. For good calibrations the stability is high, while poor calibrations the stability is low.

Chapter 5

Geocam Software Patch

5.1 Introduction

As discussed in Chapter 3, prospective motion correction offers many advantages over retrospective correction including improved k-space coverage and reduced image blurring as a result of interpolation. Many MRI scanner vendors already offer some prospective motion correction.

Prospective image regstration is a common method of motion correction available on scanners for multi-volume imaging. Prospective image registration works in a similar method to navigators; each volume in a multi volume sequence is registered to the previous volume under the assumption of rigid body motion. The transform between the images is used to describe the subject motion, which can then be used to update the scanner geometry to the subject's new position. Variations of this technique are available on Siemens scanners where it is referred to as Prospective acquisition correction (PACE) [86] (or on Philips scanners known simply as prospective motion correction). These methods can only be used in a limited set of imaging sequences, for instance dynamic EPI scans where multiple volumes are acquired such as fMRI, and then can only correct for inter-volume motion.

Navigator methods can be used to correct for intra-volume motion, however the acquisition of additional data throughout the main scan with sufficient spatial resolution, can drastically increase scan times. The use of external tracking devices such as optical tracking cameras allow for frequent geometry updates with minimal impact on scan times or image contrast. However, prospective motion correction using an external tracking camera has until now only been available on Siemens scanners through integration with the XSPACE motion library but this motion correction method has been shown to be very promising both clinically and in research.

Zaitsev et al. [87] demonstrated the power of optical tracking to image moving objects. In this case the optical tracking system consisted of two IR cameras which tracked 4 spherical markers mounted on a bite bar. They performed geometry updates after every excitation using geometry data from the cameras. They

demonstrated prospective motion correction by rotating a phantom by 15° during the scan. Results showed minimal motion artefacts despite the relatively large movement. They did however note some residual artefacts probably caused by gradient non-linearity.

Singh et al. [88] demonstrated the use of optical tracking clinically. In general bite bars are the preferred method to affix markers to a subject because they give strong coupling to the skull (a requirement for motion tracking of the brain) and markers can be positioned so they are outside of the head coils which would normally restrict the line of sight of the camera. To maximise patient comfort, Singh opted to affix multiple Moiré phase tracking markers to a subject's forehead. The advantages of this approach are two-fold. Firstly the additional marker provided redundant information about the subject location, which could be used if one of the markers was obscured by the coil. Secondly two markers can be used to measure non-rigid motion. This is especially important when affixing markers directly to the skin because facial expressions and other motions can cause the skin to move locally causing the markers to move. With only one marker this movement is inseparable from the bulk head motion, however by using two markers the separation between each marker can be used as a measure of non-rigid motion because under rigid bulk motion the Euclidean separation between the markers should not change.

Strucht et al. [89] used optical tracking to obtain the highest resolution whole brain images at 7 T. They used prospective motion correction to update the geometry of the scan using a Moiré phase tracking camera, and updated the geometry every k-space line (every TR). Using this they were able to obtain images with a resolution up to 0.12 mm, with a total scan time of 50 minutes. To obtain resolutions that high motion correction is essential as natural unconscious head movements are often an order of magnitude greater than this. This is even more important when scan times are as long as 50 minutes because subject drift over this time period can result in large differences in position between the start and end of the scan (simply because they have had more time in which to move).

Aim : To create a software patch for a Philips 7 T MR scanner that accepts ge-

ometry data from a Moiré phase tracking camera to track subject movement and prospectively update the scanner geometry. The software patch (Geocam) was tested against existing prospective motion correction available, and its additional functionality was demonstrated both in phantoms and *in-vivo*.

5.2 Implementation

5.2.1 Camera output

Moiré phase tracking uses the marker based tracking method described in Chapter 3. Tracking of the marker is performed on a closed external system which receives data from a camera placed in the bore. Data is transmitted from camera to computer via a fibre optic cable and from computer to scanner via high speed ethernet connection. Images from the camera are processed into position and orientation information using the methods described in Section 3.7.4. The structure of this data packet can be seen in Table 5.1, the position information is contained in the x, y, z parameters while the orientation information is contained in the quaternion parameters qr, qx, qy, qz. The additional information is the packet can be used for quality control. The status is used to indicate if there is a fault with tracking the marker (marker not found, or partially excluded), if the marker was well tracked then status is 0. The frame number can be used to retrospectively compare data received on the scanner to data written to log on the camera computer.

The position and orientation information can be combined into a dual quaternion ($\mathbf{C}$). A full description of a dual quaternion can be found in Chapter 4. $\mathbf{C}$ is the location of the marker in the camera reference frame.

Let ($\mathbf{C}$) denote a dual quaternion containing information about the marker's position and orientation. The individual quaternions are then defined as ($\mathbf{c}_q$) for the rotational quaternion and ($\mathbf{c}_t$) for the translational quaternion.

By default the MPT computer exports data via UDP (User Datagram Protocol) packet. However as standard the scanner does not contain the libraries required

Type	Name	Description
Integer	Status	Describes the status of motion tracking.
Integer	TargetIDNumber	ID of the marker that is being tracked.
Integer	FrameNumber	Camera frame number.
Float	x, y, z (mm)	Describes the position of the marker.
Float	qr, qx, qy, qz	Describes the orientation of the marker.
Unsigned Integer	frame Time (s)	Time the frame was acquired.
Unsigned Integer	frame Time (ns)	Time the frame was acquired.
Unsigned Integer	irq Sequence Number(s)	Counter for the number of TTL triggers.
Unsigned Integer	irq Time (s)	Time of last TTL trigger.
Unsigned Integer	irq Time (ns)	Time of last TTL trigger.
Float	tSScknot	Position of the marker in spherical co-ordinates.

Table 5.1: A table containing the data within the UDP packet sent from the camera

to accept UDP packets. Instead, the UDP packet was appended to a TCP packet on the tracking computer before forwarding the data onto the scanner.

5.2.2 Determining Image Volume Location

The geometry data within the packet is recorded in the camera reference frame and requires transformation into the magnet reference frame before use. This transform is referred to as the cross calibration transform (written as a dual quaternion $\mathbf{X}$) and describes the position of the camera relative to the isocentre of the magnet in the magnet reference frame. This transform is determined experimentally before the main experiment and is made available to the scanner as a readable text file. The method by which this transform is determined is described in Chapter 4. The marker position in the camera reference ($\mathbf{C}$) frame can be converted to the magnet reference frame ($\mathbf{M}$) using Equation 5.1.

$$\mathbf{M} = \mathbf{X}\mathbf{C} \tag{5.1}$$

In this case dual quaternions are used to express the transforms $\mathbf{X}$, $\mathbf{M}$ and $\mathbf{C}$ (The same equation would apply if using transformation matrices).

The goal of prospective motion correction is to maintain the position of the imaging volume relative to the brain, by tracking a marker which is rigidly connected to the skull assuming strong coupling between the the brain and skull. It is important to remember that the camera data describes the position of the marker, not the position of the head. By maintaining the position of the marker relative to the imaging volume, one also maintains the position of the marker relative to the brain provided that the marker and brain are rigidly connected.

The geometry data from the camera is provided as a vector of translations and a quaternion describing the orientation, where the description of a change in orientation and position are independent. This can be thought of as translating an object to the origin point, rotating it and then translating it back; so that position and orientation are separable variables.

Before camera data can be used to update the scanner geometry, the marker position in the magnet reference frame must be determined as well as the relationship between the marker position and the imaging volume. A vector diagram of the transforms required is shown in Figure 5.1.

To perform prospective motion correction the change in orientation and position of the imaging volume must be determined. This information can be determined by calculating the marker's position from isocentre.

A translation from isocentre to marker in the magnet reference frame can be described mathematically as

$$\mathbf{m}_t = \mathbf{x}_q \mathbf{c}_t \mathbf{x}'_q + \mathbf{x}_t \tag{5.2}$$

where $\mathbf{m}_t$, $\mathbf{c}_t$ and $\mathbf{x}_t$ are the translational components only of the dual quaternions $\mathbf{M}$, $\mathbf{C}$ and $\mathbf{X}$ respectively, as seen in Figure 5.1, $\mathbf{x}_q$ is the rotational component of the dual quaternion $\mathbf{X}$ only and $\mathbf{x}'_q$ denotes the inverse quaternion of $\mathbf{x}_q$. This equation can be split into two parts, the first transforming the translational component of $\mathbf{C}$ to the magnet frame and the second describing the addition of the isocentre-to-camera translational component.

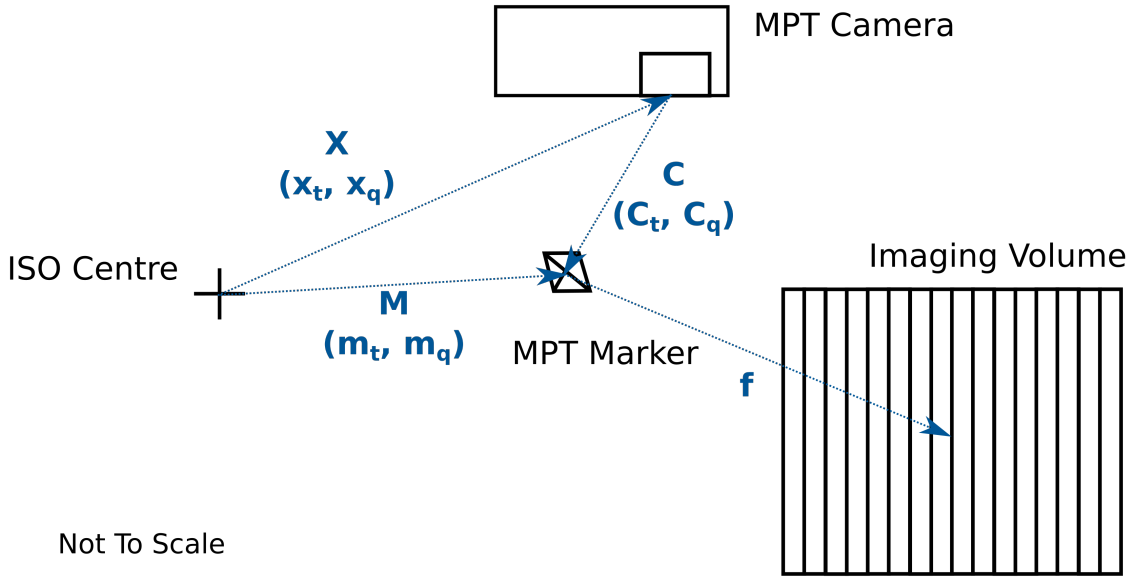


Figure 5.1: A vector diagram showing the transforms used for prospective motion correction; from isocentre-to-camera known as the cross calibration (represented as the dual quaternion $\mathbf{X}$, or a quaternion describing the translational component $\mathbf{x_t}$ and a quaternion describing the rotational component $\mathbf{x_q}$), camera-to-marker as measured by the camera ($\mathbf{C}$ or $\mathbf{c_t}$ and $\mathbf{c_q}$). These can then be used to describe the position of the marker-from-isocentre ($\mathbf{M}$ or $\mathbf{m_t}$ and $\mathbf{m_q}$). The transform from marker-to-imaging volume ($\mathbf{f}$) can be obtained at the start of the scan, however only the translational component of this is needed as the imaging volume orientation can be determined directly from $\mathbf{m_q}$.

The rotation of the marker relative to isocentre is described as,.

$$\mathbf{m}_q = \mathbf{x}_q \mathbf{c}_q \mathbf{x}'_q \tag{5.3}$$

$\mathbf{m}_q$ can then be converted to extrinsic Euler angles to describe LR, AP and FH rotations .

Once the marker position and orientation are known the position and orientation of the imaging can be determined. It is important to note that because the camera describes the position and orientation of the marker and not the imaging volume, the transform between marker and imaging volume must be determined. At the beginning of a scan the location of the head is known from planning and the marker position is known at all times from the camera data. Given these known

positions of the imaging volume and marker, the translational transform from one to the other can be determined. We describe this as the $\mathbf{f}$ vector (describing the location of the field of view from the marker), which contains only translational information because the rotational transform between the marker and imaging volume is not needed. In the case of 3D imaging there is a single $\mathbf{f}$ vector describing the single imaging volume, however for multislice imaging there are N (where N is the number of slices) $\mathbf{f}$ vectors each describing the position of each slice, from the marker. The $\mathbf{f}$ vector for a particular slice is defined in Equation 5.4.

$$\mathbf{f}_{slice} = \mathbf{m}_t(n=0) - \mathbf{s}_t(n=0) \tag{5.4}$$

where $\mathbf{s}_t$ is the position of a particular slice, and n is the number of shots into the scan. The determination of $\mathbf{f}$ occurs during the initialisation phase which is usually the first shot of the scan for shot-by-shot correction.

During subsequent shots the desired position and orientation of the imaging volume is unknown assuming the subject has moved. However the marker position in the magnet frame can always be determined using the camera data and the calibration.

Given the assumption that the marker and the head are rigidly connected, the change in orientation of the marker, is identical to the change in orientation of the head in the magnet co-ordinate frame, so the change in orientation since the last shot can then be described simply as:

$$\mathbf{s}_{\Delta r} = \mathbf{m}_r(n) - \mathbf{m}_r(n=0) \tag{5.5}$$

where $\mathbf{s}_{\Delta r}$ is the change in orientation of the slice about each axis(x, y and z). and $\mathbf{m}_r(n)$ is the rotational quaternion of $\mathbf{M}$ as Euler angles.

As mentioned previously the camera describes the orientation and position of the marker separately, and so the imaging volume's position is also determined separately. In the case of a pure translational this can be determined by subtracting the $\mathbf{f}$ vector from the marker's current position directly. However if the marker

has rotated the $\mathbf{f}$ vector also rotates, so the new $\mathbf{f}$ vector can be described by:

$$\mathbf{f}(n) = \mathbf{m}_{\Delta q}\mathbf{f}(n=0)\mathbf{m}'_{\Delta q} \tag{5.6}$$

where $\mathbf{m}_{\Delta q}$ is a quaternion describing the change in orientation of the marker.

The change in the imaging volume location can the be determined using:

$$\mathbf{s}_{\Delta t} = (\mathbf{m}_t(n) - \mathbf{f}(n)) - \mathbf{s}_t(n=0) \tag{5.7}$$

where $\mathbf{s}_{\Delta t}$ describes the change in position of the imaging volume (or of each individual slice).

$\mathbf{s}_{\Delta t}$ (Equation 5.7) and $\mathbf{s}_{\Delta r}$ (Equation 5.5) describe in full the change in imaging volume position and orientation required to track the subject movement.

5.2.3 Updating Scanner Geometry

The scan geometry was updated using the existing prospective motion correction update procedure present on a Philips, 7 T Achieva system. Unlike the existing method however, the geometry data was provided by the camera rather than image registration, which enabled more regular geometry updates. In the Geocam patch the scanner geometry was either updated during the volume acquisition loop (in the case of volume by volume correction), or over the shot acquisition loop (in the case of shot by shot correction), as seen in Figure 5.2

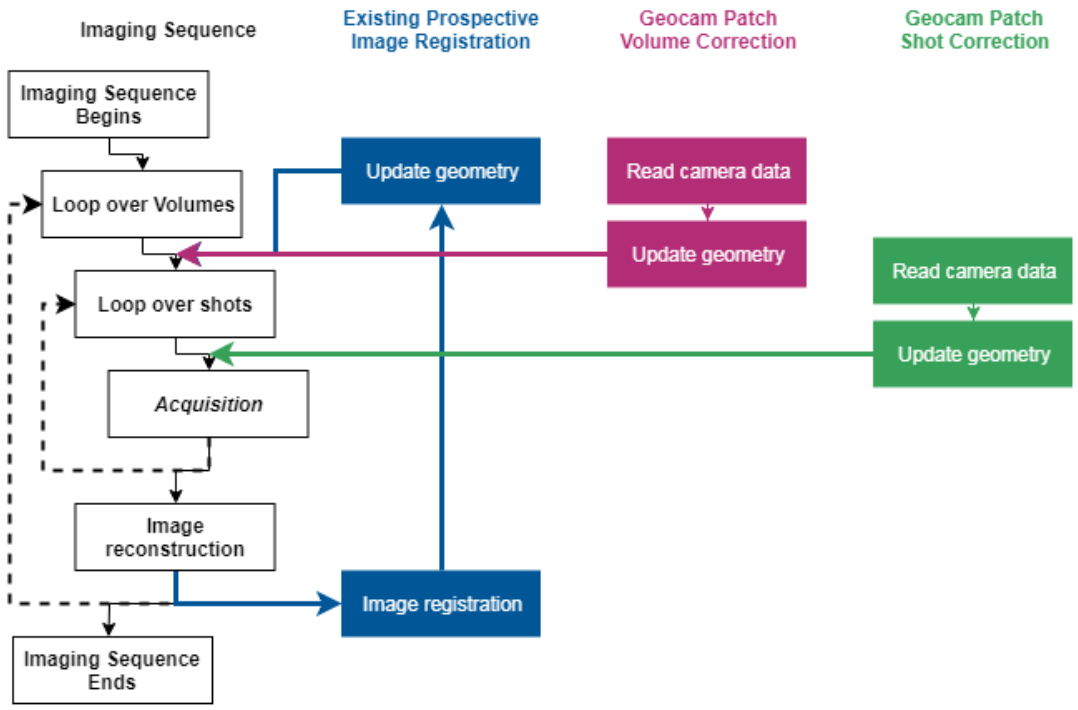


Figure 5.2: A diagram showing where the geometry was updated during the imaging sequence for the existing PMC (blue) the volume by volume Geocam PMC (pink) and the shot by shot Geocam PMC (green).

5.3 Method

5.3.1 Demonstration of Geocam and Prospective Image Registration

The functionality of the Geocam patch was compared to existing Philips prospective motion correction by applying volume-by-volume correction to a moving phantom. The Philips prospective motion correction method is a variation of PACE [86] in which each volume image in a multi-volume scan is registered to the previous image. This registration is used to determine how the subject has moved, and update the next scan's geometry. Scanning was performed using a Philips 7 T Achieva MRI scanner with a Nova transmit coil and a 32 channel receive head coil. The phantom was scanned using a single shot EPI scan with an isotropic FOV 190 mm and isotropic voxels of size 2 mm, a TE/TR of 35/2000 ms and a

SENSE factor of 2. A total of 10 volumes were acquired. EPI was chosen because it is most commonly used for multi-volume imaging such as fMRI, and as such is the mostly likely type of scan to be used with Philips prospective motion correction. The phantom was scanned for 5 dynamics in a stationary "resting" position, between the 5th and 6th dynamic the phantom was rotated by approximately 25° about the FH direction (Figure 5.3). A delay in geometry updates is present, where movement is detected by the camera, but the scanner has not updated the scanner geometry for the next volume. For the purpose of validation this means that the phantom can be scanned once in its new moved position using the original scan geometry. The scan then proceeds with the phantom in its now "rotated" position with the updated scanner geometry. This experiment was repeated using both the Geocam patch and the Philips prospective motion correction. The size of the rotation was monitored using the Moiré phase tracking camera to ensure the movement size was similar between motion correction methodologies. A timeline of this experiment can be seen in Figure 5.3

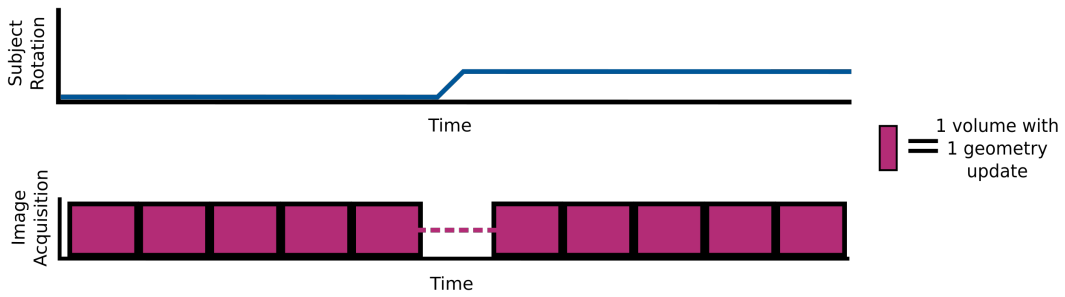


Figure 5.3: The timeline of the demonstration of the Geocam patch. The phantom was imaged a total of 10 times, and rotated between the 5th and 6th volume.

5.3.2 Determining Geocam Accuracy

To quantify the accuracy of the prospective motion correction a custom phantom was designed. The phantom was a 150 mm diameter sphere containing six 10 mm thick square plates equally spaced and parallel to one axis of the sphere

but with a random in-plane rotation. Each plate contained randomly positioned 15 mm holes. The phantom was filled with salt water at 0.5% concentration to mimic the salt concentration of the body. The random position of the holes and orientation of the plates removed any geometric symmetry and allowed accurate image registration to be performed. A bite bar with a MPT marker was attached to the phantom using pressure sensitive adhesive putty (Blu Tack™) so that the phantom could be tracked from within the coil. An image of the phantom can be seen in Figure 5.4.

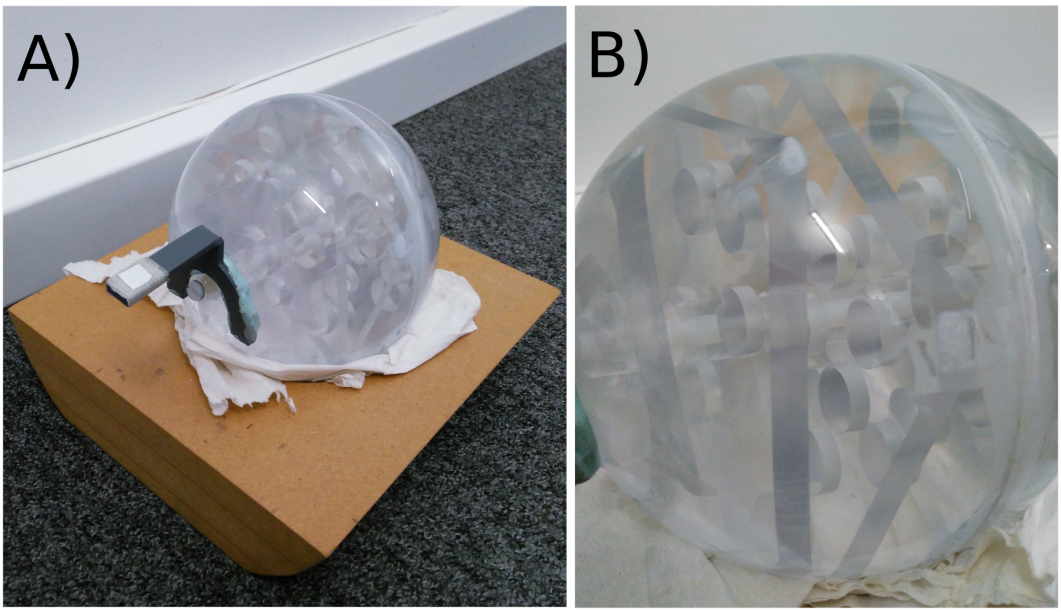


Figure 5.4: An image of the phantom created to determine the accuracy of the Geocam patch. (A) shows the phantom in full with the bite bar, (B) is a closer view of the rotated plates within the phantom.

The phantom was scanned using a Philips 7 T Achieva MRI scanner using a Nova transmit/receive coil, the 32 receive channel coil was removed to allow for additional equipment to be placed into the coil to control the movement more accurately. The phantom was then scanned using a TFE FFE scan with a FOV = 240, 240, 150 mm with a isotropic pixels of 3 mm, a TE/TR = 1.48/6 ms and 25 volumes using optically tracked prospective motion correction. A magnetization prepared gradient echo sequence used in this case rather than an EPI sequence due to phase artefacts observed in the EPI images (Figure 5.8 when rotating the

phantom underwent a large rotation. As before, the phantom was scanned in it's initial "resting" position for 5 volumes, before being rotated around the scanner's LR axis, it was then rotated again after a further 5 volumes, and then once more an additional 5 volumes. The phantom was placed back into the resting position after a total of 20 volumes and the remaining 5 volumes were then acquired. A total of 5 volumes were acquired per phantom position (25 volume images total). The prospective motion correction updated the scanner geometry to correct for the movement.

The accuracy of the PMC was then determined by performing image registration between the images in the moved pose and the initial pose. The error of the registration was measured by registering images of the phantom in the same pose (volumes 1-5, 6-10 etc) to each other and determining the standard deviation of the residual motion. Rigid body image registration was performed in MATLAB using a pixel intensity based method. Any transform performed as a result of the image registration was taken as the error of the prospective motion correction. The error in the registration was determined by registering the 5 images of the phantom in each pose to themselves, since in theory there is no movement of the phantom within a pose.

A timeline of this experiment can be seen in Figure 5.5

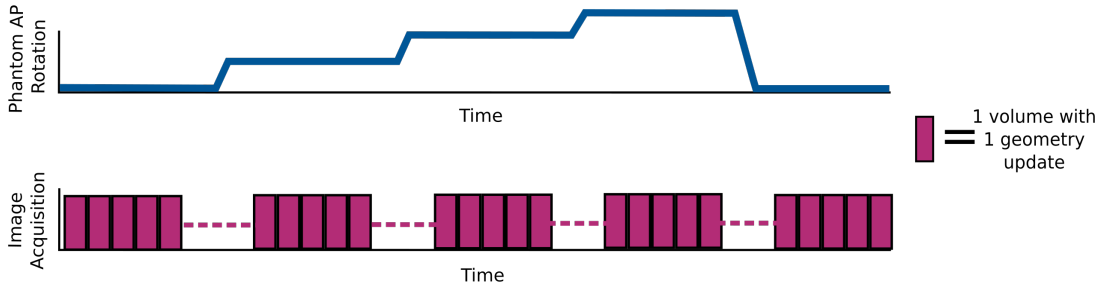


Figure 5.5: The timeline of experiment for determining the Geocam accuracy. The phantom was imaged 25 times with an increasing rotation after every 5th volume.

5.3.3 Intra-volume Motion Correction

The Geocam patch was then tested for intra-volume motion correction. A cylindrical water filled phantom was used, this phantom contained a series of distinct geometric shapes along the length of the cylinder, these shapes acted to remove the symmetry in the RL and AP directions. This phantom was scanned using a multi-shot 3D magnetization prepared gradient echo sequence repeated 3 times where each volume contained 128 shots. This scan was chosen to reduce phase artefacts seen with EPI and to maximise the number of shots per scan, thus increasing the temporal resolution of the motion correction. The total scan time was 100 seconds (33 seconds per volume). During the first volume acquisition the phantom remained in it's resting position. During the second volume the phantom was rotated about the FH direction (approximately 8 seconds into the dynamic, taking 1 second to complete the movement), the phantom was then translated in FH (approximately 13 seconds into the dynamic, taking 2 seconds to complete the movement). The phantom then remained it its rotated position for the remainder of the volume as well as the following volume. A timeline of this experiment can be seen in Figure 5.6:

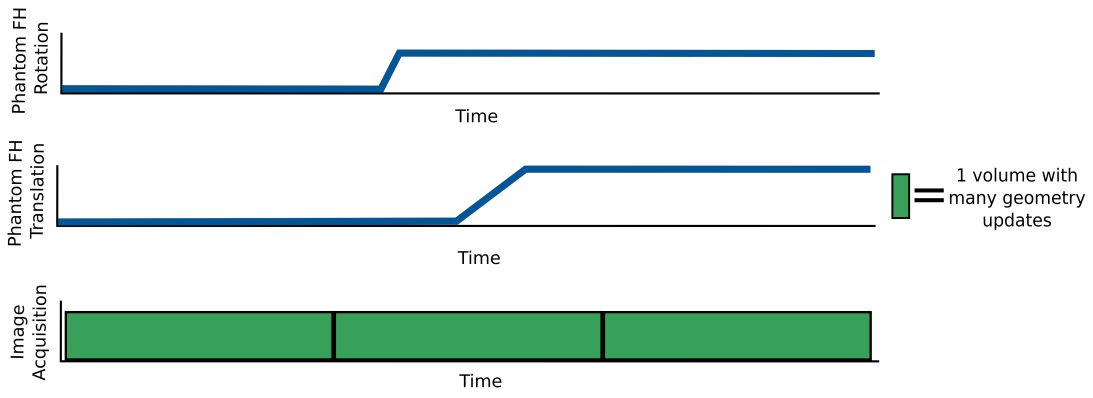


Figure 5.6: The timeline of experiment demonstrating intravolume motion correction. The phantom was imaged 3 times with FH translations and rotations during the second volume. 20 geometry updates were performed per volume.

The Geocam patch was also used to correct anatomical images *in-vivo*. Using

a MPRAGE sequence and an optimised fast field echo (FFE) sequence.

The MPRAGE scan used a FOV = 256 $\times$ 256 $\times$ 180 mm^3, a matrix size of 256 $\times$ 256 $\times$ 180 and a TE/TR = 3.68/7.71 ms. The shot interval was 2067 ms. The subject was instructed to rotate their head (head shake) halfway through the scan by the experimenter. The subject then remained in this rotated position for the remainder of the scan.

The FFE sequence used a FOV = 256 $\times$ 256 $\times$ 2, a matrix size of 256 $\times$ 256 $\times$ 1 and a TE/TR = 4.28/50 ms and 10 volumes. A total of 256 shots were acquired with a shot interval of 50 ms. The subject was instructed to shake their head half way through the 3rd volume, remain stationary for the 4th and 5th volume, and then nod their head half way through the 6th volume and remain stationary for the remaining volumes. The scanner and camera were synced using a TTL trigger at the beginning of the scan. This enabled the movement performed during the scans to be synced to the data acquisition.

Both scans were repeated with motion correction enabled and disabled.

The subject wore a custom-made bite-bar made of a rigid nylon plastic and thermosetting dental plastic. The bite-bars were made in-house by warming the thermosetting dental plastic and mounting it to the nylon plastic. A dental imprint was then taken by asking the subject to bite down on the dental plastic, before allowing the bite bars to cool and harden over 24 hours. Bite bars were sterilised using dilute sodium carbonate peroxide denture cleaner. It was necessary to mount the MPT marker onto the bite-bar rather than the subject for two reasons: firstly, the marker had to be visible to the camera during the acquisition (a marker placed directly onto the subject would be obscured by the coil). Secondly, a marker mounted directly to the subject is prone to non-rigid motion between the brain and marker due to the poor coupling between the objects.



Figure 5.7: An example bite bar worn by subjects to allow tracking of the MPT marker from within the MRI environment.

5.4 Results

5.4.1 Demonstration of Geocam and Prospective Image Registration

EPI images comparing the Philips and Geocam motion correction methods are shown in Figure 5.8. They show that the Philips prospective motion correction and the Geocam method give nearly identical results. Both methods were able to accurately correct for phantom movement. However both methodologies suffer heavily from EPI phase artefacts after motion correction.

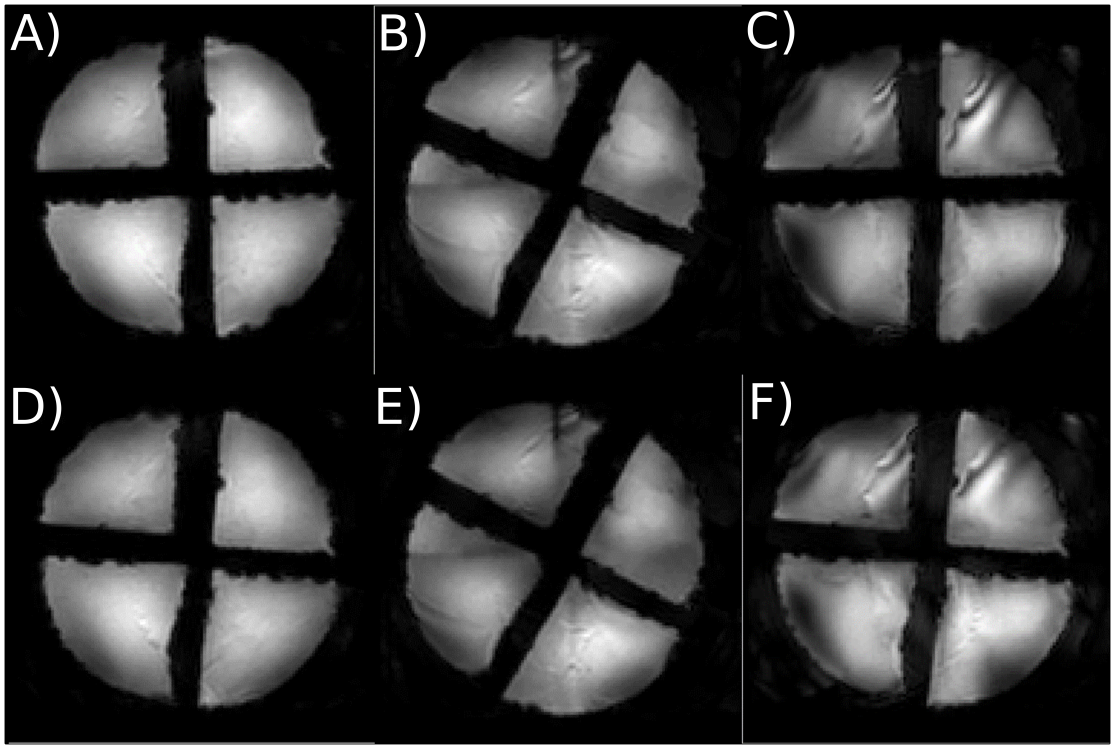


Figure 5.8: EPI images of a moving phantom corrected using prospective motion correction. A-C show the phantom when scanned using the Geocam patch, while D-F show the phantom scanned with the Philips prospective motion correction. A,D show the phantom in its resting position. B, E show the phantom in its rotated position *before* the geometry update and C, F show the phantom in its new rotated position *after* the geometry update.

5.4.2 Determining Geocam Accuracy

Figure 5.9 shows the images obtained for a series of different sized rotations. The residual errors measured using image registration and are shown in Figure 5.10. The error in the registration varied between poses, the largest error observed was 0.2 mm along the z direction (image slice direction). The camera is quoted to be accurate to 0.1 mm and 0.1 °. The error in registration was determined by registering images in the same pose and was typically of the order 0.1 mm and 0.1 °.

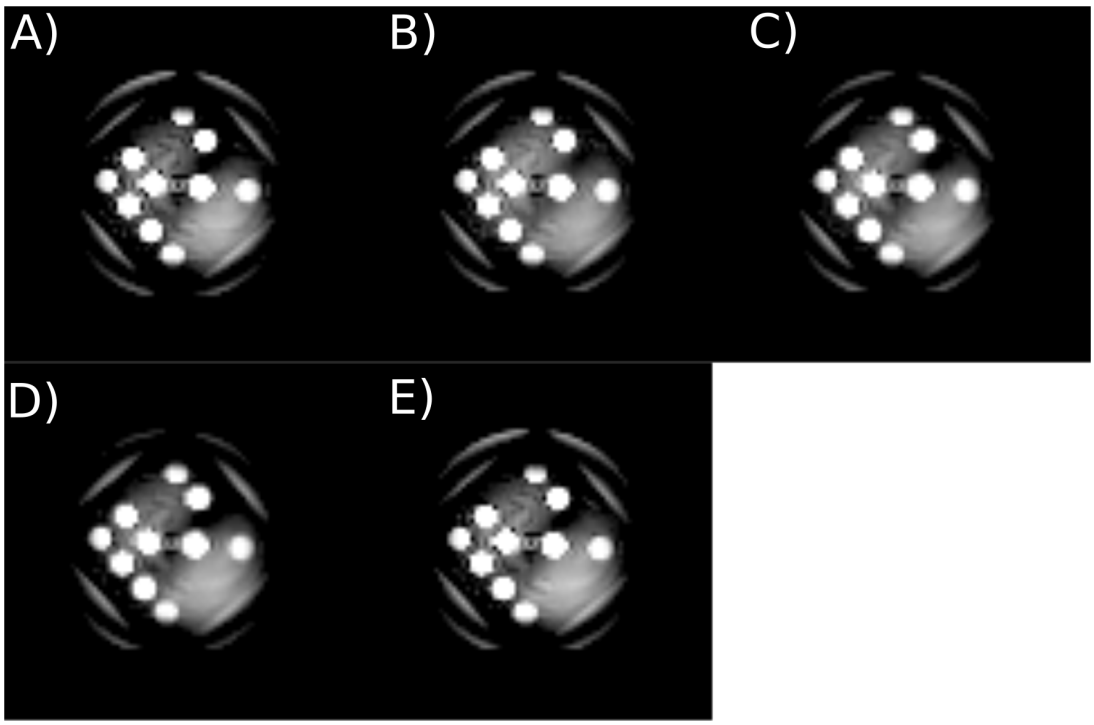


Figure 5.9: The images of the phantom used to determine the residual error in the GEOCAM patch. A) shows the phantom in its resting position, B), C), D) show the phantom with a $\approx 5°$, $\approx 10°$, $\approx 15°$ rotation about the AP direction (in plane), E) shows the phantom when it was returned to its approximate resting position. Note that the Geocam patch ensures that, despite the phantom rotation, the images are all aligned

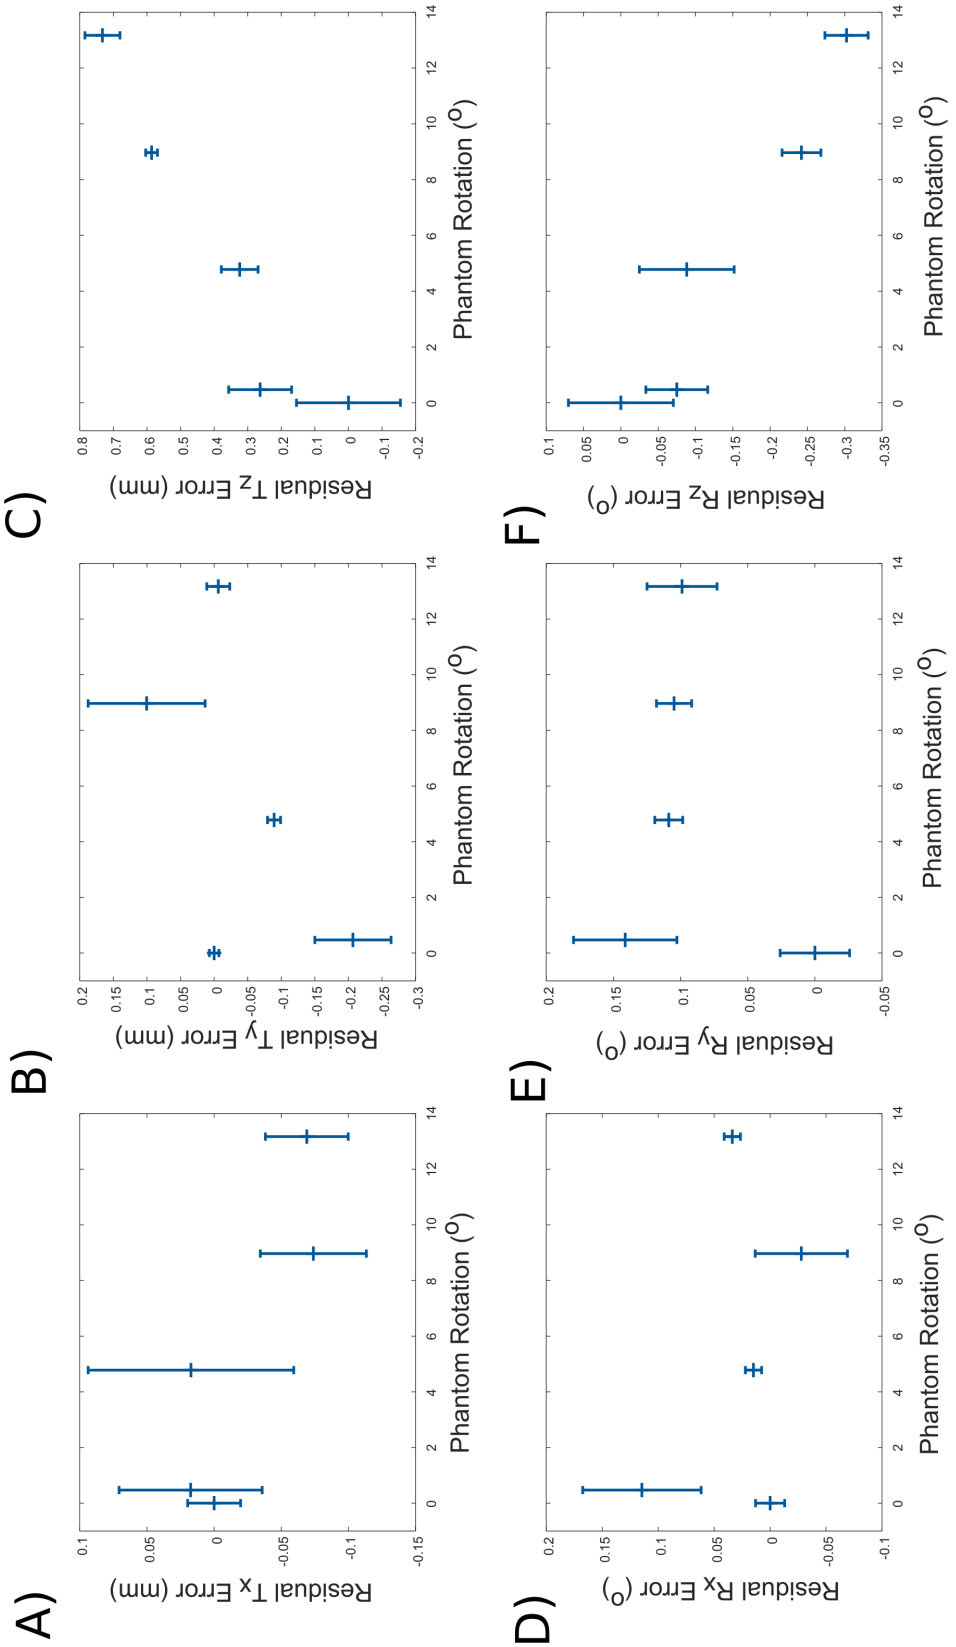


Figure 5.10: The residual error in the phantom position determined using image registration after prospective motion correction. The residual error in translation (A, B, C) and rotation (D, E, F) is shown as a function of the size of rotation.

5.4.3 Intra-volume Motion Correction

The Geocam patch was shown to prevent intra-volume motion artefacts as seen during an extended acquisition in Figure 5.11. No movement occurred during the acquisition of the images in Figure 5.11A and Figure 5.11C. Registration between Figure 5.11C and Figure 5.11A gave a mean error of 0.4 mm and 0.4°.

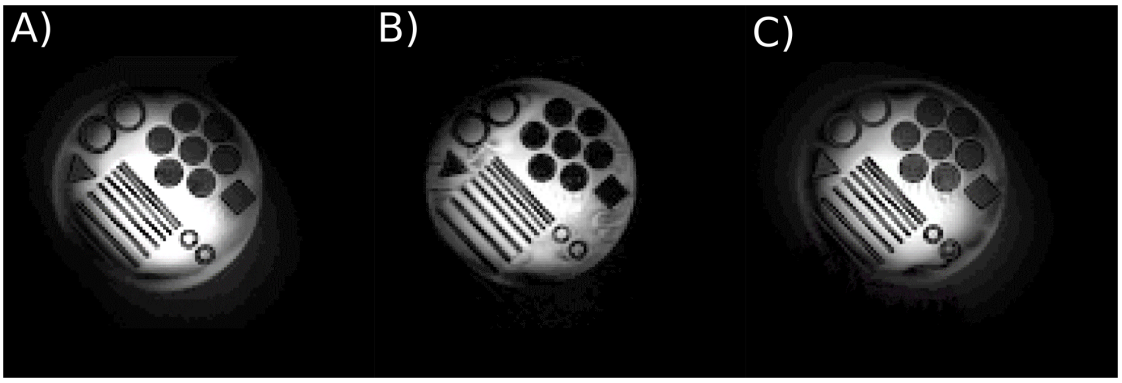


Figure 5.11: A phantom that has been prospective motion corrected for a movement that occurred during acquisition. A) shows the phantom in its resting position, B) shows shows the phantom as it moved during the acquisition and C) shows the phantom in its new position

The Geocam patch was used with an MPRAGE sequence with the geometry updated each shot in the presence of a single large subject movement. The scan was repeated twice; with (Figure 5.12A) and without (Figure 5.12B) motion correction. Figure 5.12 shows that the image is improved from the use of the Geocam patch compared to no correction at all, however significant artefacts still remain.

The MPRAGE sequence represents a typical 3D imaging sequence used for anatomical imaging. The FFE image was selected to contain many shots to maximise the frequency of geometry updates. Figure 5.13 shows the images acquired during an FFE imaging sequence where the subject was instructed to move twice during the acquisition. The sequence was repeated with and without the Geocam PMC, the subject movements are shown in Figure 5.14.

In addition to blurring, Figure 5.13 also showed changes in image contrast in images with high subject movement. A region of white matter (5 mm $\times$ 5 mm)

towards the front of the head was selected and the pixel intensity averaged for each image in Figure 5.13.

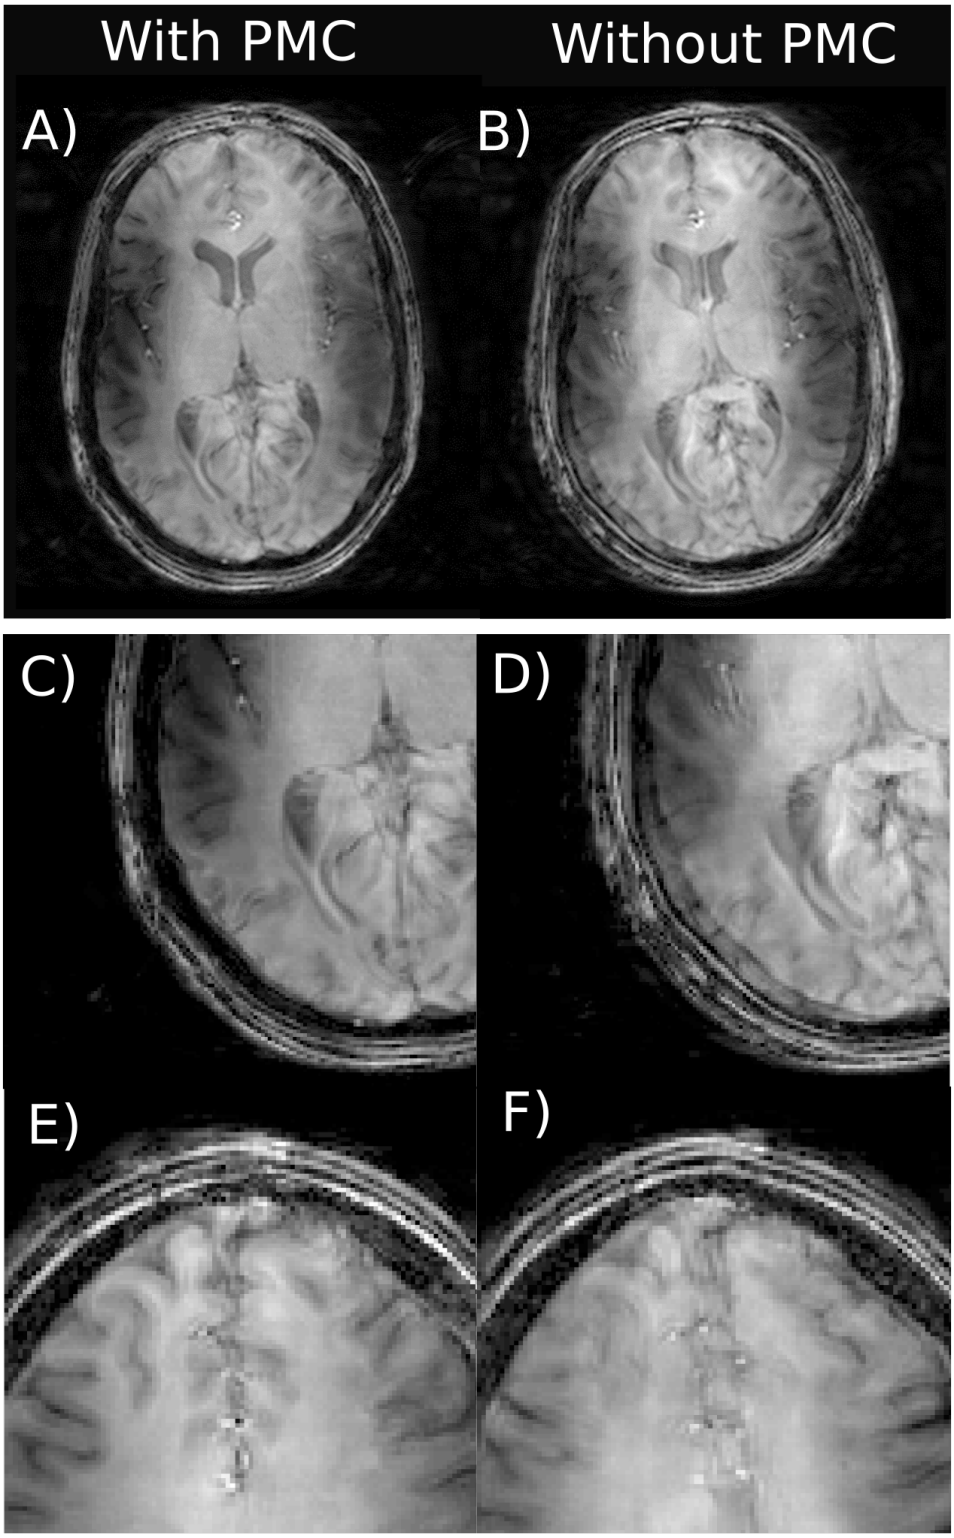


Figure 5.12: An In-vivo MPRAGE image with (A, C, E) and without (B, D, F) motion correction. The subject moved their head approximately half way through the scan in a “shaking” motion. (C), (D), (E) and (F) show regions of the brain where there were significant motion artefacts in the no motion correction case.

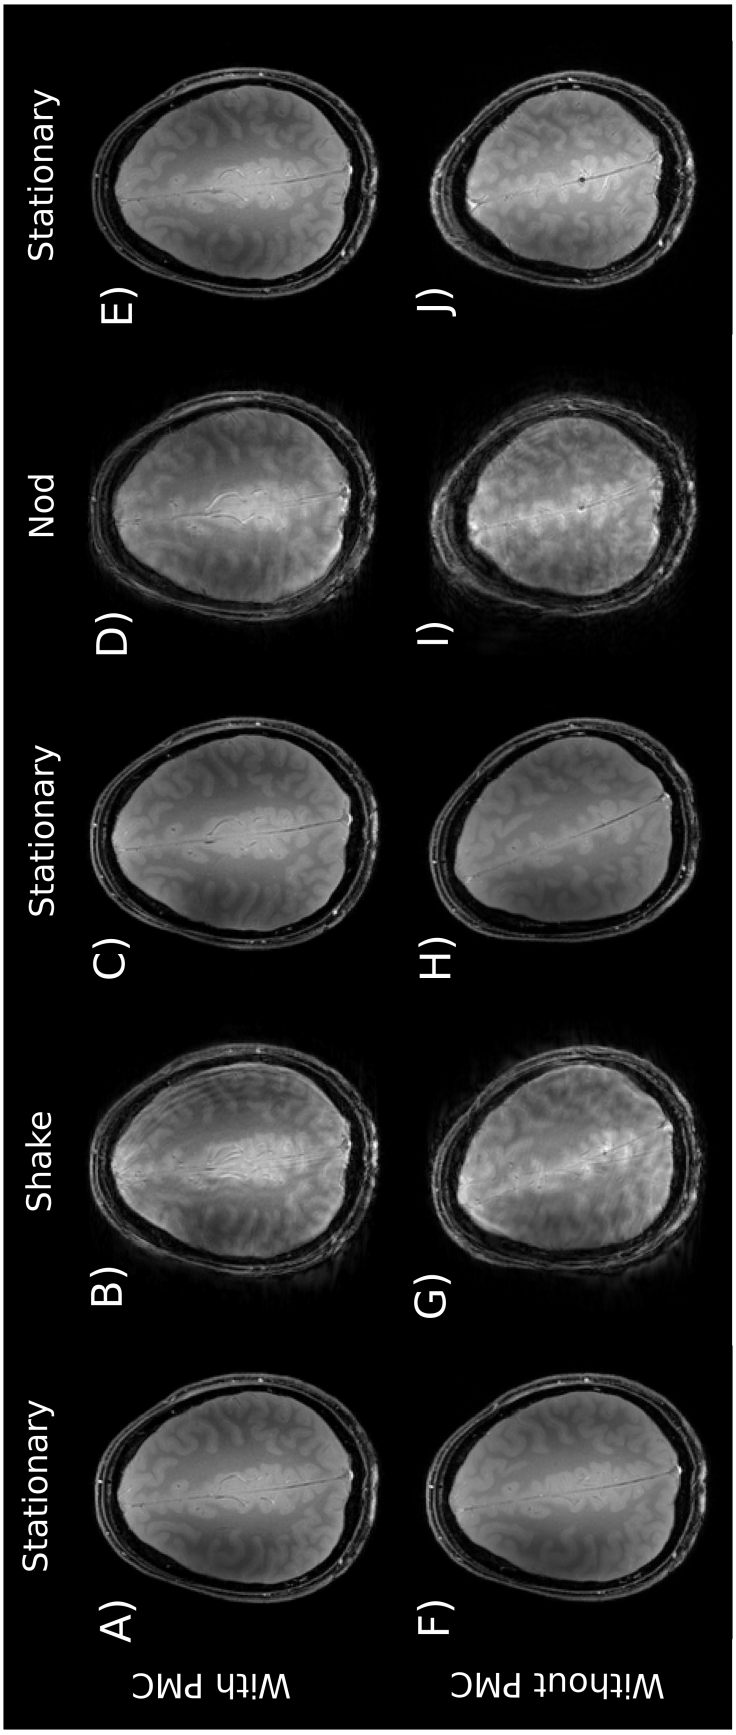


Figure 5.13: The FFE images showing the images acquired with (A,B,C,D,E) and without (F,G,H,I,J) motion correction. The subject was instructed to shake their head during the acquisition of image (B) and (G) and nod their head during (D) and (I).

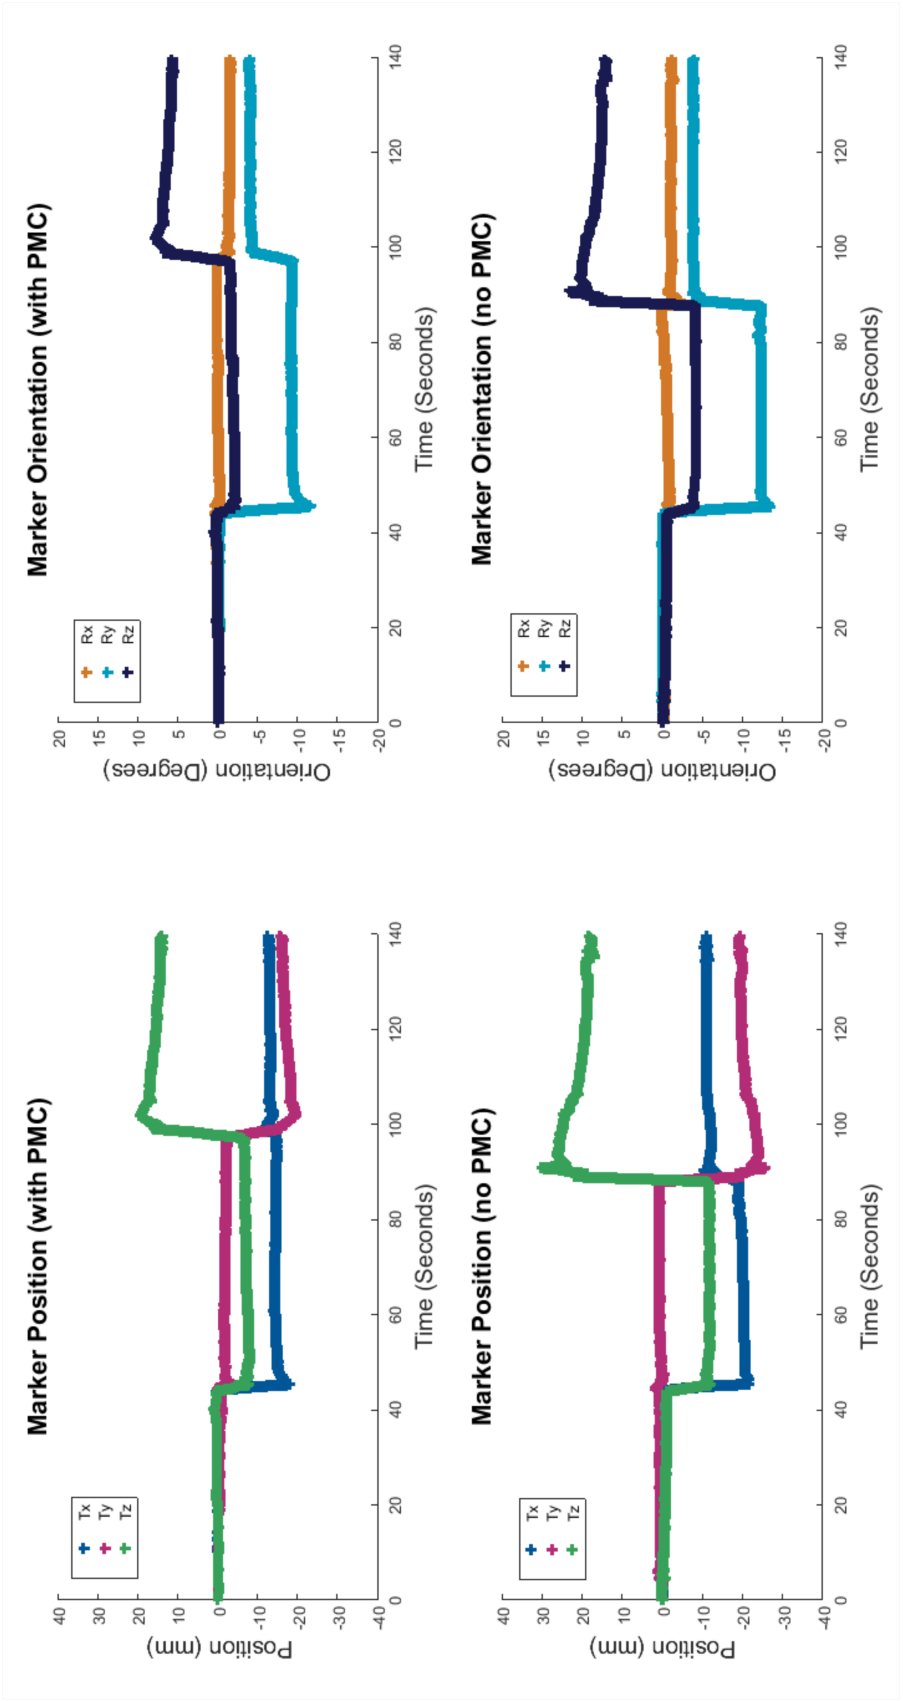


Figure 5.14: The size of the subject movement during the FFE imaging sequence. The subject was asked to nod their head during the 3rd dynamic (approximately 1/3 into the sequence), and shake their head during the 6th dynamic (approximately 2/3 into the sequence).

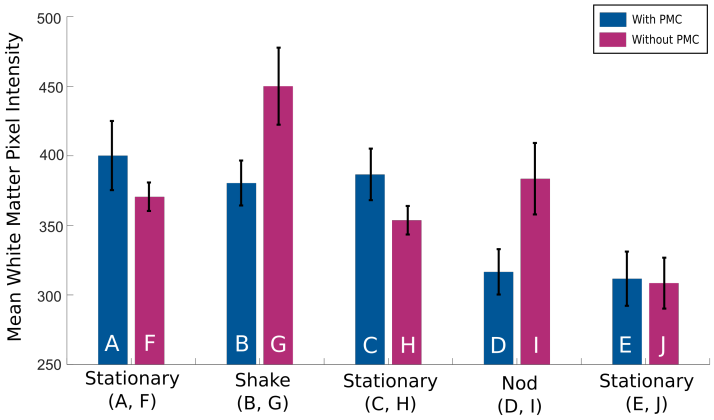


Figure 5.15: The mean pixel intensity of an area of white matter for each image in Figure 5.13. Due to spin history effects the pixel intensity increases when the subject moved their head

.

5.5 Discussion

5.5.1 Demonstration of Geocam and Prospective Image Registration

The Geocam patch resulted in nearly identical images compared to current Philips prospective motion correction as demonstrated in Figure 5.8. The phantom rotation was performed by hand and so the exact size of the rotation could be different between the Geocam patch scans and the Philips patch test, however because of the delay in the geometry updates we can see from Figure 5.8B and Figure 5.8E that the size of the rotation was very comparable. In both cases the phantom was successfully motion corrected back to its original position.

While the rigid body motion was successfully corrected for in both cases, both corrected images suffer from large EPI phase artefacts. These phase artefacts arise from a failure in the EPI phase correction algorithm. The EPI phase correction was determined before the start of the scan via a prescan designed to calculate the first and second order phase correction parameters. However the phase correction parameters appear to be geometry dependent, when the phantom, was rotated new phase correction parameters were not recalculated. This means that the phase correction that is applied in Figures 5.8C and 5.8F were not suitable for the phantoms rotated position (remember that the phantom is in a different location in Figure 5.8C compared to 5.8A, but the scanner geometry is changed to make it appear as if they are in the same position). This artefact presents a potential limit on the size of the motions that can be corrected for when performing EPI imaging, however it may be possible to correct for these artefacts retrospectively by regridding k-space, if errors can be captured in a prescan. In reality the motions used here were much larger that would be seen for a reasonably compliant subject, but when doing motion studies where the subject is asked to deliberately move their head the images should be inspected for phase artefacts and if necessary the size of the movement reduced. Subsequent *in-vivo* scanning (Figure 5.16) the phase artefacts were seen to be most prominent for

head shakes (in plane rotations) rather than other head movements.

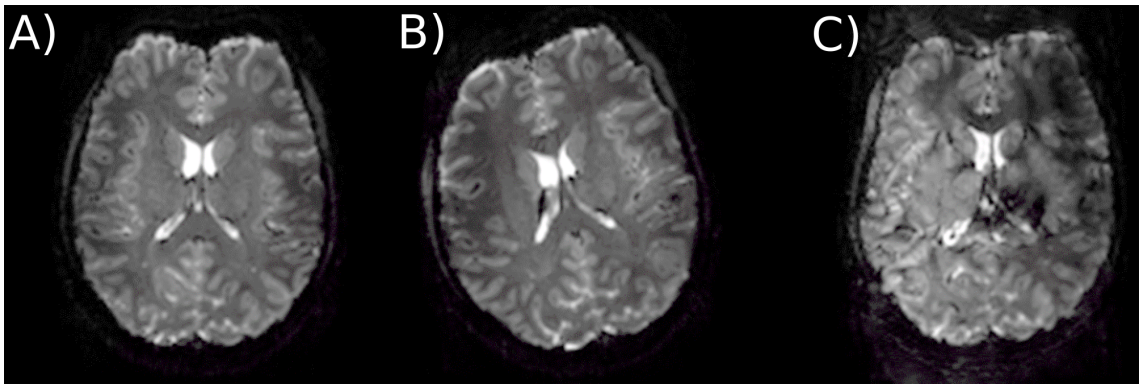


Figure 5.16: An example of the phase artefacts seen during an EPI scan with the Geocam patch. The subject was asked to move from their resting position (A) to a large shake position. (B) shows the position of the head in the new shake position before geometry updates, (C) shows the head position after the geometry update.

There is an additional artefact present in both the Geocam and Philips motion correction dataset in the rotated image before the geometry update (Figure 5.8B and Figure 5.8E). This cross shape can be see through the 4 quadrants of the phantom. We theorise that this is caused by the coil sensitivity profile used in the SENSE reconstruction. The coil sensitivity profile is determined before the scan begins, with the phantom in its initial resting position, because of the shape of the phantom used in this experiment there is no information about the coil sensitivity in the centre cross region. This is not an issue during stationary imaging because the there is no signal to reconstruct in that area anyway. However when the phantom had moved, there was then signal in this region but still no coil sensitivity profile corresponding to that area which resulted in signal drop out. This is supported by the fact that these artefacts were removed when the geometry update takes effect. Because we are using the existing Philips prospective motion correction methodology, in addition to changing the scan geometry it also applies the same geometry changes to the coil sensitivity profile. These artefacts are unlikely to be an issue during *in-vivo* scanning because the head does not contain large regions devoid of signal, and once the geometry updates take effect the artefacts are removed anyway. If this artefact does present itself as an issue

it is possible an offline SENSE reconstruction could be used. By doing this one could apply additional smoothing the coil sensitivity maps to extend the sensitivity maps into areas devoid of signal and reduce SENSE artefacts.

5.5.2 Determining Geocam Accuracy

The prospective motion correction using the Geocam patch is shown to be accurate to less than 1 mm over a 13° rotation as seen in Figure 5.10. Errors in registration were determined by co-registering images of the phantom in the same pose (no movement). The registration errors were similar in magnitude to the quoted camera accuracy (0.1 mm and 0.1°). The residual error in the PMC correlates linearly with the size of the motion performed at least for AP rotations and errors in FH directions. This suggests that the optical tracking method performs better for realistic smaller head movements than larger head movements, although in either case the prospective motion correction patch would yield better results than no correction at all, so all movement benefits from the Geocam patch. The likely cause of this error is the calibration: In Chapter 4 it was shown that for errors in the calibration, the error in the imaging volume position grew linearly with the size of the movement, which is how the errors in Figure 5.10 are presented.

The images in Figure 5.9 show an interesting artefact as a result of the PMC, there is an RF artefact that moves around the edge of the phantom as it rotates. This artefact is caused by the physical location of the RF coils relative to the phantom, they appear to move around the phantom when in reality it is the phantom that is moving and the RF coils that remain stationary. This highlights another potential downfall of the prospective motion correction; it is limited by immoveable physical constraints. One possible solution to this would be to use a dynamic RF pulse which can compensate for the new subject position, possibly using multiple transmit coils to ensure a smoother RF profile in the presence of subject motion.

5.5.3 Intra-volume Motion Correction

The Geocam patch can be used to correct for intra-volume motion as seen in Figure 5.11. In this case the phantom was rotated in plane (FH axis) by approximately 16° and then translated by 30 mm. There is no ground truth with which to compare this image, unlike for inter volume motion correction where a delay could be added to view the rotated phantom. The phantom position shown in Figure 5.11C is prospectively corrected to appear to be in the same pose as the phantom in Figure 5.11A, and by this co-registration we can obtain the error of the prospective motion correction. Registration gave a mean error of 0.4 mm and 0.4°. This agrees with the typical errors we expected from Figure 5.9 in terms of the accuracy of the prospective motion correction. Figure 5.11B does show artefacts which relate to the phantom motion. Typically direct motion artefacts cause ghosting as the same object is captured in multiple spaces during phase encoding, that does not appear to be the case here, more likely these artefacts are the result of flow because they are not seen when the phantom is stationary in Figure 5.11A and 5.11C.

The *in-vivo* motion corrected MPRAGE scan shows that the Geocam patch does reduce motion artefacts in anatomical imaging. Figure 5.12B which does not use prospective motion correction shows significantly more artefacts than Figure 5.12A which was acquired using the Geocam patch, this is most notable towards the left anterior part of the brain. In both cases the subject rotated their head around z (in plane) half way through the scan after being instructed to do so by the experimenter. However, despite the improvement, Figure 5.12A still shows motion artefacts. For this MPRAGE sequence the TR = 2000 ms, which means a geometry update was applied once every 2 seconds. This allowed a significant amount of motion to occur before the update took place.

Fluctuations in the pixel intensity were observed between the images in Figure 5.13, these are highlighted in Figure 5.15. These changes in intensity are likely to be caused by two effects, spin history and changes in B_1 due to the new head position. In the case of no PMC, the increase in white matter intensity during the shake movement (G) is likely caused by spin history effects. We can tell this

because when the subject remains stationary in the new head position (H), the white matter intensity returns to its previous value in the original head position (F), suggesting head pose makes very little difference to B_1 for shake movements. The Geocam patch was able to prevent spin history effects during the shake movement , which is why the variation from stationary, to shake, to stationary (A to B to C) is small. Variations in intensity between stationary to nod to stationary (C to D to E) suggest that there is a difference in B_1 in these head positions making it difficult to separate contributions from spin history and B_1 effects. However, with PMC the white matter pixel intensity closely matched the last stationary pose, which could suggest reduced sensitivity to spin history effects, if the majority of k-space was captured in the latter head position.

An alternative imaging sequence was used to generate Figure 5.13. In this case the scan was optimised using a very short TR which allowed many shots to be performed during the acquisition. The scans required 256 shots per volume and the scanner geometry was updated every 50 ms. Despite this the images still contained artefacts related to the subject movement as seen in Figure 5.13B where the subject shook their heads. However k-space was acquired using a linear k-space trajectory, because the subject shook their heads approximately half way through the scan, these images represent the worst case scenario. There is still a significant improvement for intravolume motion correction between the Geocam correction and no correction at all, as seen in Figures 5.13D and 5.13E. Comparison of Figures 5.13E and 5.13J highlights another advantage of the Geocam patch. In this case the subject performed a nodding motion, however because this is a single slice scan and the movement was through slice. This meant that the original slice that was intended to be imaged (the slice in image 5.13F) was not excited and hence not imaged at all. The Geocam patch ensured the desired region of interest was always acquired despite the through slice movement.

5.5.4 Geocam patch Limitations

A limitation with the current Geocam framework is that it processes geometry by passing the updated geometry through the existing prospective motion process. The advantage of this is that all geometry is safety checked before being implemented. Equation 5.8 shows the changes in gradient waveform in response to a subject rotation [87].

$$\mathbf{g}'(t) = \mathbf{R}\mathbf{g}(t) \tag{5.8}$$

where $\mathbf{g}$ is a gradient waveform and $\mathbf{R}$ is the a matrix describing the subject rotation. This equation shows the contribution of each physical gradient in generating the waveform. The gradient strengths are restricted to prevent damage to the coil or excessive peripheral nerve stimulation and so the safety checks ensure that the scan aborts if the requested gradient strength exceeds these limits.

An additional advantage mentioned earlier is the geometry updates are also considered during reconstruction as well, this means that SENSE artefacts can be compensated for. The disadvantage of using the existing motion correction pipeline is the delay between the camera data being available and when the scan is updated. Because the geometry checks are run in parallel with the current acquisition the geometry update is not applied during the current shot, but instead is applied during the next shot, meaning there is always a delay of one shot before the geometry update is applied. While this is a minor issue for some 3D scans where many shots are acquired rapidly, this can be a problem for MPRAGE sequences where the shot time is of the order of seconds, because large subject motion can occur between updates. A key point is that the Geocam patch performs the correction per shot *not per excitation*, this means that for a single shot multi-slice scan, there is only one update for the entire volume, despite each slice being excited separately (see Figure 5.17).

Additional update delays exist because of the use of TCP packets rather than the intended UDP packets. The TCP packets require a handshake protocol to ensure no data is lost. The scanner has a buffer that stores data before it is

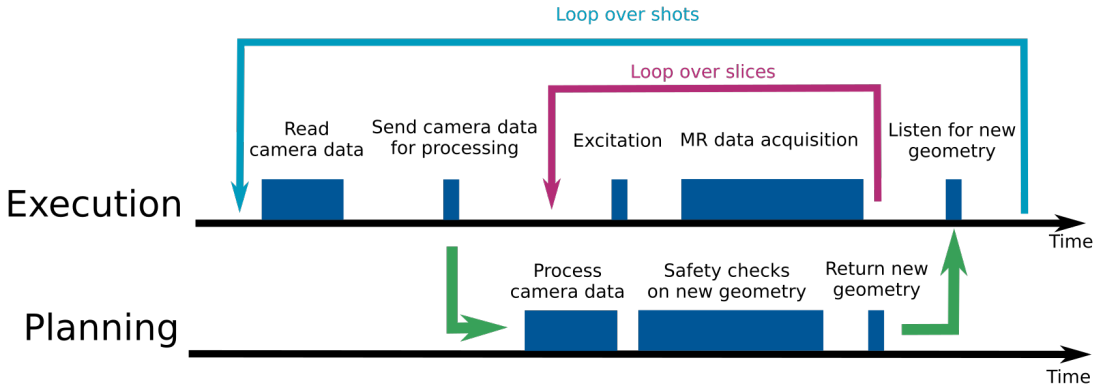


Figure 5.17: A diagram showing the actions performed by the planning (plans the scan, including geometry) and execution (carries out the MR sequence) systems. This sequence of events results in only one geometry update in a single shot, multislice scan.

processed. When full it does not accept any additional data and the handshakes are denied. The data is then stored on a buffer on the camera computer waiting to be sent. This means that when the data is read on the scanner and the scanner buffer cleared, the camera sends the old data it has backlogged in its buffer. This can create delays of several volumes before the real time geometry is sent to the scanner. This is a problem with scans that have a long time between shots. For example multi-slice multi-dynamic scans, where the entire volume must be scanned before the buffer is next read.

To fix this delay, the Geocam patch could be altered in two ways. Firstly the call to read geometry data from the camera could be made more frequent, for example every line of k-space. At the moment this update is performed per shot beacuse it is infrequent enough to allow the geometry checks to be performed for each update as discussed above. Performing updates every k-space line would mean that the buffer would be cleared more regularly preventing a backlog of old data. The second potential change would be to change the scanners gradient and readout frequencies directly. To do this gradient waveforms could be adjusted using Equation 5.8 and the phase of the transmit and receive RF signal adjusted to compensate for the offsets introduced by subject motion [90]. This would drastically speed up the geometry update process. Simplified safety checks could

still be performed by monitoring the requested gradient waveform strength and aborting if necessary.

5.6 Conclusion

Here we have created new functionality for a Philips 7 T MRI scanner by using an optical tracking camera to perform prospective motion correction. We have demonstrated prospective motion correction for both inter-volume and intra-volume movements. The motion correction is accurate to less than 1 mm for large head movements but is may vary depending on the accuracy of the camera calibration (See Chapter 4). This motion correction methodology results in increased image quality by removing most motion artefacts, however some artefacts may still remain for large movements (EPI phase, flow, RF inhomogeneity and shimming etc.).

Chapter 6

EPI Geometric Distortion Correction

6.1 Introduction

Echo planar imaging (EPI) is the corner stone of functional magnetic resonance imaging (fMRI) and diffusion weighted imaging (DWI) because of the speed with which volume data sets can be acquired. During fMRI scans it is common to measure brain activity as a function of some stimulus. Because of this, total scan times are of the order of several minutes, with many volumes being acquired in this time to maximise temporal resolution. This means that the images are subject to large inter-volume movements. It is common to retrospectively correct these volumes by registering these data sets together to ensure that pixels within a time series are consistent in space. Retrospective correction can be hampered by poor image registration however [91]. This can be caused by localised warping of images due to changes in local magnetic environments which are not typically taken account of during retrospective motion correction. These artefacts occur in the phase encoding direction. Prospective motion correction maintains the phase encoding direction relative to the head, ensuring that warping always occurs in the same direction (relative to the head) and hence images are consistent (Figure 6.1).

EPI has a strong sensitivity to the frequency of the MR signal in each voxel. During EPI imaging k-space data is acquired in a 'raster-like' fashion. The readout direction (k_x) is acquired quickly as the gradients switch back and forth rapidly. Because a line of k-space data is gathered quickly, readout bandwidth is said to be relatively high. The time to advance in the phase encoding direction (k_y) is long because each small step in k_y must wait for a full line of k_x to be acquired, hence the phase encoding direction is said to have low bandwidth. Because of this EPI images are more susceptible to B_0 field inhomogeneity which causes significant dephasing along k_y throughout the acquisition, discussed further in section 6.2.

Generally the static B_0 field in the scanner varies very little ($\approx$1 ppm over a 22 cm spherical volume [92]), however, objects within the bore perturb the field depending on their magnetic susceptibility. Shimming can improve the homogeneity by applying a low order polynomial field to remove some of these variations [93].

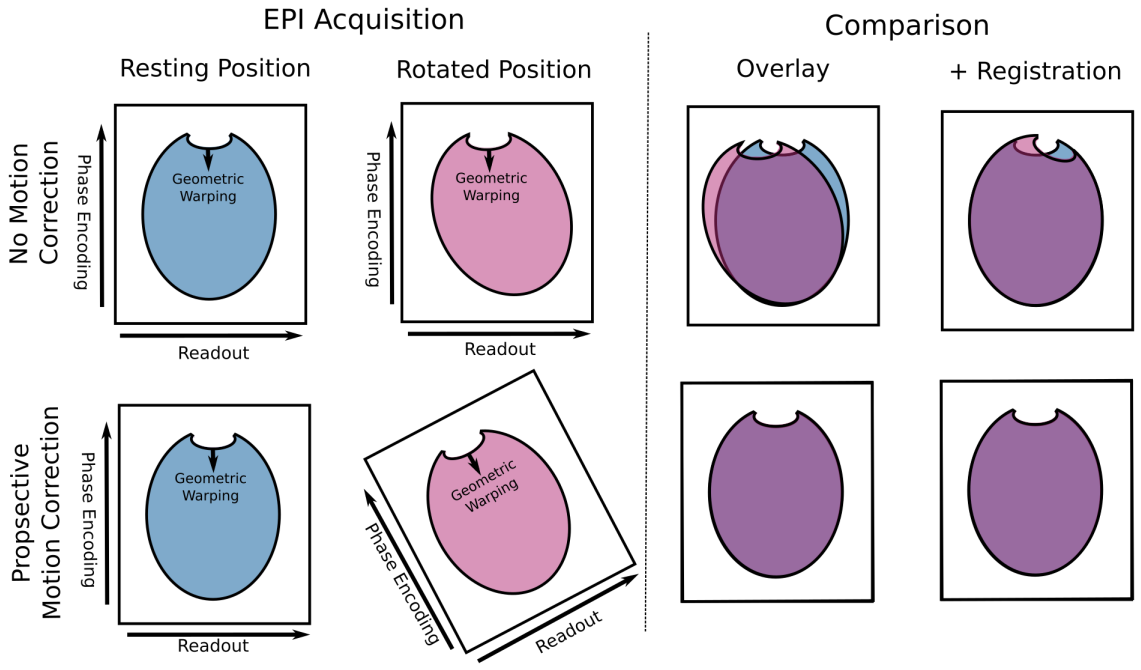


Figure 6.1: How prospective motion correction maintains the phase encoding direction relative to the head. This results in more consistent images in the presence of subject motion compared to images acquired without motion correction.

Shimming effectiveness can be limited over large areas and cannot compensate for areas of high inhomogeneity near air-tissue interfaces such as the sinuses, which can often result in perturbations in the static field of the order of 100s of Hz at 7 T.

Many different EPI distortion correction methods exist, such as those available as part of the FMRIB Software Library (FSL) [94]. “TOPUP” uses two EPI images acquired with opposing phase encoding directions to estimate the perturbations in field due to susceptbility [95]. It does this by estimating a B_0 field which when used to unwarp each image in opposite directions, results in the greatest similarity between the images as measured using a least mean squared difference. Because the distortion is estimated directly from the images, no prior information regarding B_0 is required.

TOPUP considers subject movement between acquisitions, and factors rigid body motion into its model. However the assumption is made that B_0 and head position are independent. This not the case since B_0 changes depending on

head orientation [96] which causes additional warping of the images. This can be significant for large movements.

If the internal B_0 field is known, by aquiring a B_0 map at the start of a scan session or using TOPUP on a pair of initial EPI images to generate a distortion map (and hence B_0 map) then images can be corrected directly using FSL's geometric unwarping of EPI data (FUGUE) [55]. This method converts the B_0 map into pixel shift values using,

$$\Delta R = \gamma \Delta B_0(x, y, z) N \tau_{esp}. \tag{6.1}$$

where ΔR is the distance a pixel is shifted in the phase encoding direction, ΔB_0 is the field perturbation is Tesla, τ_{esp} is the echo spacing, N is the matrix size and γ is the gyromagnetic ratio. Pixels can then be shifted based on the local ΔB_0. In a typical fMRI experiment, a B_0 map is acquired at the beginning of the scan session which can then be used to unwarp later EPI images.

Compared to TOPUP, FUGUE can be relatively quick once a B_0 map is known. This assumes however, there is no movement between the B_0 map and the EPI images one is correcting.

Aim: To reduce artefacts induced by changes in susceptibility artefacts caused by inter-volume movements during EPI scans using prospective motion correction to preserve the phase encoding direction relative to the head. The relationship between B_0 and head orientation was modelled and used to predict B_0 maps for given head positions. Geometric distortions due to B_0 inhomogeneity were then corrected for using a predicted B_0 map for given head nod and shake positions.

Aspects of the work described in this chapter formed the basis of the following abstract: *James Smith, Olivier Mougin, Kingkarn Aphiwatthanasumet, Matthew Clemence, Andrew Peters, Richard Bowtell, and Paul Gloverand Penny Gowland. Combined prospective motion correction and retrospective B0 correction for EPI using optical tracking and predictive B0 maps. In Proceedings of the International Society for Magnetic Resonance in Medicine, Number 4089, 2018.* [97]

6.2 Theory

6.2.1 Magnetic Susceptibility

The susceptibility (χ) of material is a measurement of how magnetized an object becomes when placed in a magnetic field,

$$\chi = J/B_0, \tag{6.2}$$

The resulting polarisation (J) (magnetisation) of an object either opposes or aligns with the external magnetic field (B_0) depending on whether the material is diamagnetic or paramagnetic respectively. In either case the external field is perturbed, either it is reduced and field lines dispersed in the case of diamagnetic materials, or is increased and field lines concentrated if the material is paramagnetic. Most biological materials are diamagnetic but oxygen is paramagnetic. This means in areas where there are tissue boundaries between the two material types there are high levels of field perturbation due variation in field over a small volume. These field perturbations alter the local magnetic field experienced by a spin depending on its location within the head and hence cause localised changes in the phase of the spin over time.

6.2.2 Geometric Distortion

As seen in Equation 2.28, the MR signal (without relaxation) depends on the spin precession rate, which can be written in terms of the local magnetic field,

$$S(\mathbf{r},t) \propto \rho(\mathbf{r}) \exp^{i\gamma \int_0^t B_0(\mathbf{r},t')dt'}. \tag{6.3}$$

In a 2D imaging sequence with readout gradients (G_x), phase encoding gradients (G_y) and field inhomogeneity (ΔB_0) Equation 6.3 can be written as,

$$S(\mathbf{r},t) \propto \rho(\mathbf{r}) \exp^{i\gamma \int_0^t G_x(t')xdt' + \int_0^t G_y(t')ydt' + \int_0^t \Delta B_0(x,y,z)dt'}. \tag{6.4}$$

Equation 6.4 shows that the phase of the signal is dependent upon the gradient history of the spins as well as the field inhomogeneity. The phase difference between adjacent k-space data points in the readout ($\Delta\Phi_r$) and phase encoding (Φ_{pe}) directions for an EPI scan with trapezoidal readout gradients, and triangular phase encoding blips, can be written as,

$$\Delta\Phi_r = \gamma(G_x x t_d + \Delta B_0(x,y,z)t_d) \tag{6.5}$$

$$\Delta\Phi_{pe} = \gamma(G_y y t_{ramp} + \Delta B_0(x,y,z)(t_{ramp} + N t_d)) \tag{6.6}$$

where t_d is the dwell time which is defined as the interval between digitized samples, t_{ramp} is the ramp time of the switched gradients and N is the matrix size for an $N \times N$ matrix. A figure showing the contribution to the phase during each part of the EPI sequence can be seen in Figure 6.2.

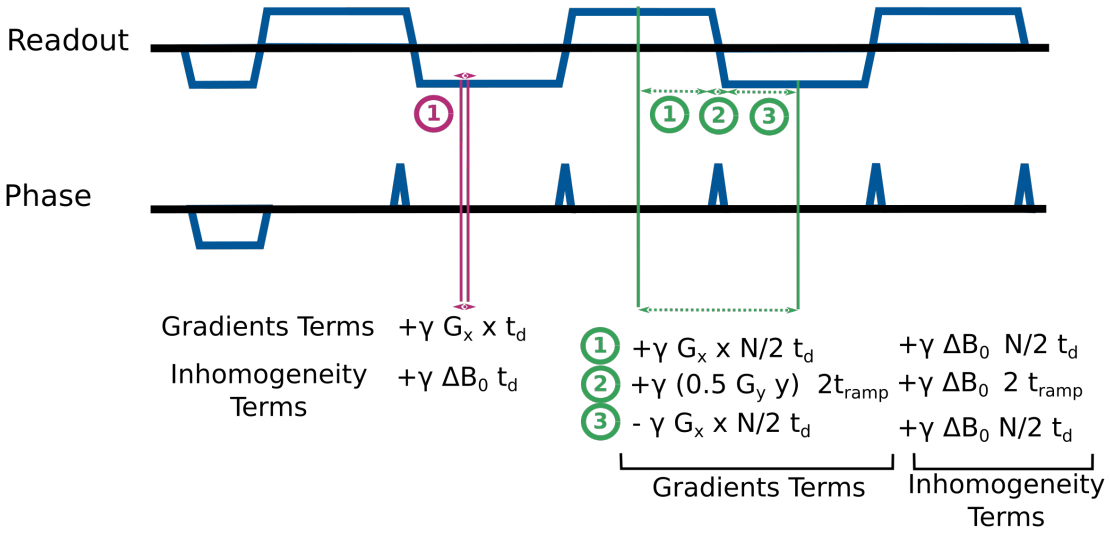


Figure 6.2: The phase acquired between two adjacent points in k-sapce in the readout (pink) and phase encoding directions (green) due to the applied gradients and the inhomogeneity.

The added phase due to field inhomogeneity (ΔB_0) during the acquisition results in geometric warping because the phase determines the position of a pixel along the phase encoding direction of the image. This geometric warping is dependent upon the inhomogeneity of the static field at that location. From Equa-

tions 6.5 and 6.6 one can see that B_0 inhomogeneity affects both the readout and phase encoding directions. Because of the $(t_{ramp} + Nt_d)$ factor in the phase encoding direction, the geometric distortions in this direction are much greater than those seen in the readout direction. It is common to simplify the geometric distortions to the phase encoding direction only, hence Equation 6.1, which describes the pixel shift along this direction alone.

6.3 Method

To correct the geometric distortions that arise from B_0 inhomogeneity and head orientation, we first verified the linear relationship between head orientation and B_0. Using this relationship, predicted field maps were then estimated and used to unwarp EPI images. The quality of the distortion correction was accessed by measuring the accuracy of registration between the EPI images in different head poses, and the EPI images and an anatomical image.

6.3.1 Modelling B_0

B_0 has been shown to vary linearly over small head movements [96, 98]. We verified this by scanning a subject at multiple head orientations (12 total). The subject wore a bite bar, which allowed head orientation to be tracked using a Moiré Phase Tracking (MPT) camera. The subject was asked to make a series of small nods and shakes and hold each position for ≈1 minute, while the actual B_0 map was acquired (Figure 6.3).

B_0 maps were acquired using a dual echo sequence with a TE and ΔTE = 5.6/1.0 ms; TR = 25 ms. B_0 maps had a FOV = 256 $\times$ 192 $\times$ 160 mm and a matrix size of 64 $\times$ 64 $\times$ 40. Maps were measured as a change in frequency and measured in Hz.

The actual B_0 maps were registered to the initial position retrospectively using pixel intensity based registration of the magnitude component of the actual B_0 map in MATLAB. The same transform was then applied to the corresponding

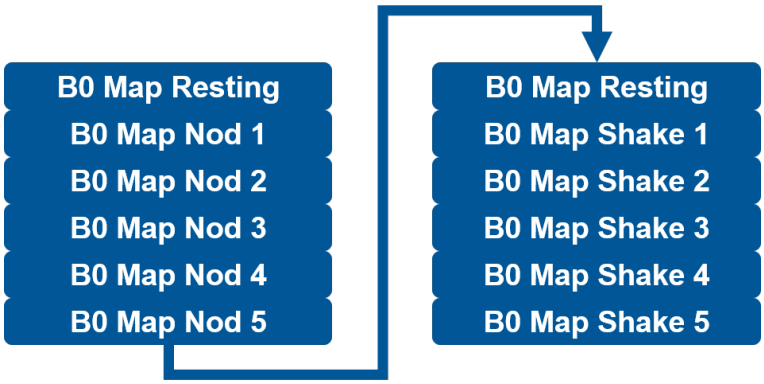


Figure 6.3: The order by which B_0 maps were acquired to demonstrate the linear relationship between B_0 and head orientation.

actual B_0 map. The difference in B_0 (ΔB_0) between the resting pose and each other pose was determined by subtracting each registered B_0 map from the initial actual B_0 map. For each voxel ΔB_0 was then plotted as a function of pitch (ϕ), roll (θ) and yaw (ψ). ΔB_0 was then fitted using multi linear regression based on Equation 6.7 to generate 4 fitting parameter maps (u, v, w and C).

$$\Delta B_0(r) = u(r)\phi + v(r)\theta + w(r)\psi + C \tag{6.7}$$

Using the orientation of the subject, parameter maps were then used to generate predicted B_0 values which were compared to the actual B_0 values.

6.3.2 Geometric Distortion Correction

To combine prospective motion correction and geometric distortion correction, 6 subjects (4 male, 2 female with mean age 25.6) were scanned using a 7 T Philips Achieva MR scanner with subject motion captured using Metria Moiré phase tracking (MPT) camera which was calibrated using the methods described in Chapter 4. Each subject wore a custom-made bite bar onto which the Moiré phase tracking marker was fixed.

A total of 12 B_0 maps were acquired for each subject, as well as 4 $\times$10 volume EPI scans and an anatomical scan for registration. At the start of each scan session a fast low angle shot (FLASH) anatomical image was obtained. The anatomical image had an isotropic field of view (FOV) of 192 mm and an isotropic

matrix size of 192 with a TE/TR of 5.94/3610 ms. EPI scans were interleaved between B_0 scans so that a B_0 map was acquired before and after movement (see Figure 6.4).

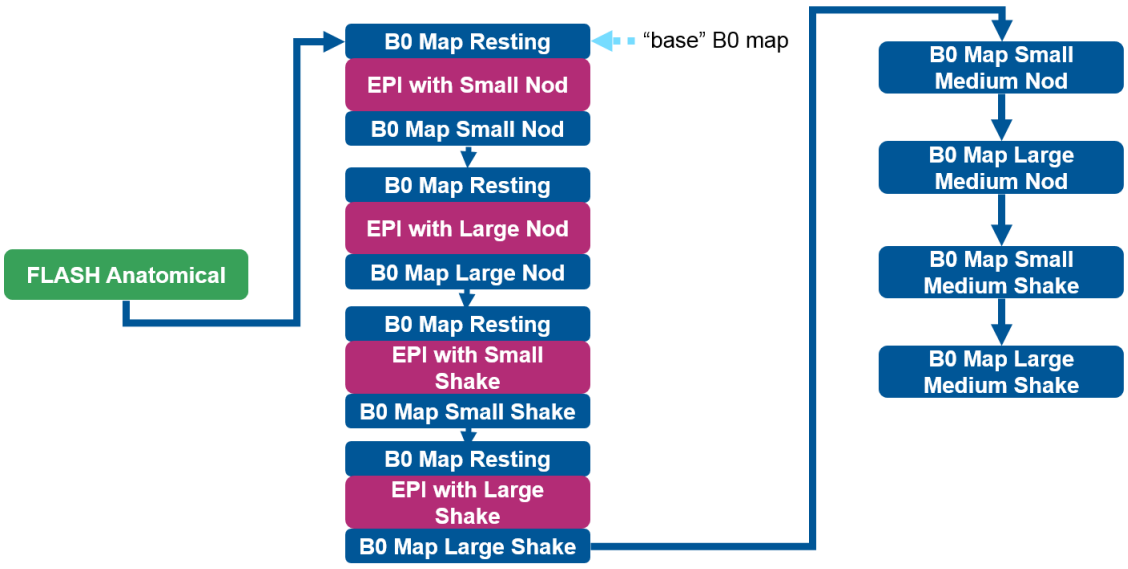


Figure 6.4: The sequence of scans used to investigate combined prospective motion correction and retrospective geometric distortion correction.

EPI images contained 10 volumes with an isotropic FOV of 192 mm, a matrix size of 96 $\times$ 96 $\times$ 96 and a TE/TR = 26/5056. Five of the subjects were acquired using AP phase encoding while 1 subject was acquired with RL phase encoding for comparison. Between the 5th and 6th EPI volume, scanning was paused and the subject was asked to move from a resting position to a new head position, either a small or large, shake or nod. They then remained in this position for the remaining 5 volumes and the following B_0 map (see Figure 6.5).

EPI scans had volume-by-volume prospective motion correct enabled using the Geocam patch (Chapter 5), which changed the geometry of the scan in real time so that volumes were aligned to the first volume of the scan (with a 2 volume delay before correction), the first EPI image of the 10 dynamics EPI was used as the resting volume, and the 10th dynamic was used as the rotated volume.

At all times the orientation of the subject was measured using the MPT camera and stored as a quaternion on an external log. These quaternions were converted

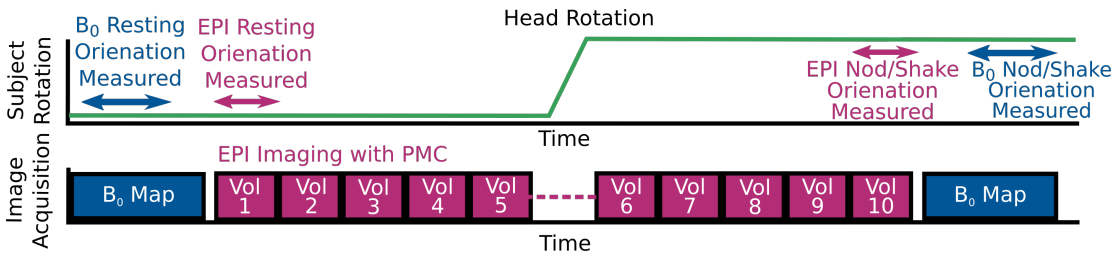


Figure 6.5: The relationship between when the actual B_0 maps images and EPI images were acquired compared to the subject movement. This sequence was repeated for a small and large, nod and shake.

to extrinsic Euler angles (zyx) in MATLAB after the scan. The subject orientation was averaged over a 5 second window ($\approx$ 400 data points) during each actual B_0 map acquisition, and over a 1 second window ($\approx$ 80 data points) for each EPI volume.

The actual B_0 maps and the subject orientation were used to generate the parameter maps (u, v, w and C) per subject. Using these parameter maps and Equation 6.7, predicted B_0 maps were generated for each pose in the EPI scan by using the orientation of the head measured during the scan.

All B_0 maps were then converted from Hz to radians. EPI distortion correction was performed using FSL's FUGUE which uses the pixel shift method as outlined in Section 6.2. In addition to the predicted B_0 maps, images were also corrected using a resting B_0 map (base B_0 map) gathered at the beginning of the scan to represent the current method of correcting geometric distortions using FUGUE.

To measure the accuracy of the geometric distortion correction, the corrected EPI images were registered to the anatomical images using the FSL linear image registration tool (FLIRT). Images were registered using a 9 parameter fit (3D rigid body and scale) to compensate for the different voxel sizes. The accuracy of the registration was measured using mutual information to avoid contributions from differences in image contrast. Mutual information was measured over a rectangular region of interest which contained the cerebrum (Figure 6.6).

In addition, to measure the accuracy of the combined prospective motion correction and the geometric correction together, the correlation between each EPI

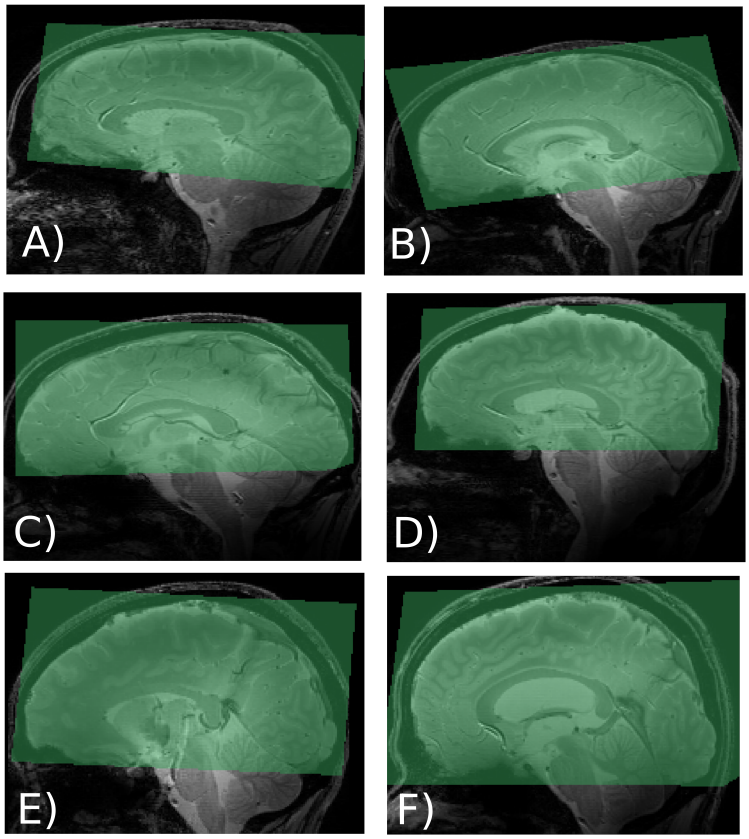


Figure 6.6: The typical region of interest the mutual information between the EPI and anatomical images was measured over.

image and the initial resting pose was measured using a Pearson correlation coefficient.

EPI Phase Artefact

During the acquisition of the EPI images, if significant phase artefacts (like those shown in Chapter 5) were observed, the image was discarded and the scan repeated and the subject was asked to make a smaller movement. Small phase artefacts should not impact the comparison between the correction method (base vs predicted B_0 map) because the each method is applied on the same EPI image, and so the artefact should affect each correction method equally. However large phase artefacts could cause reduced correlation between the anatomical and EPI images.

6.4 Results

6.4.1 Modelling B_0

B_0 maps were plotted as a function of nod movement (Rx) and shake movement (Rz) and are shown in Figure 6.7A and Figure 6.7B respectively. Figure 6.7 shows the ΔB_0 for a voxel located towards the front of the brain.

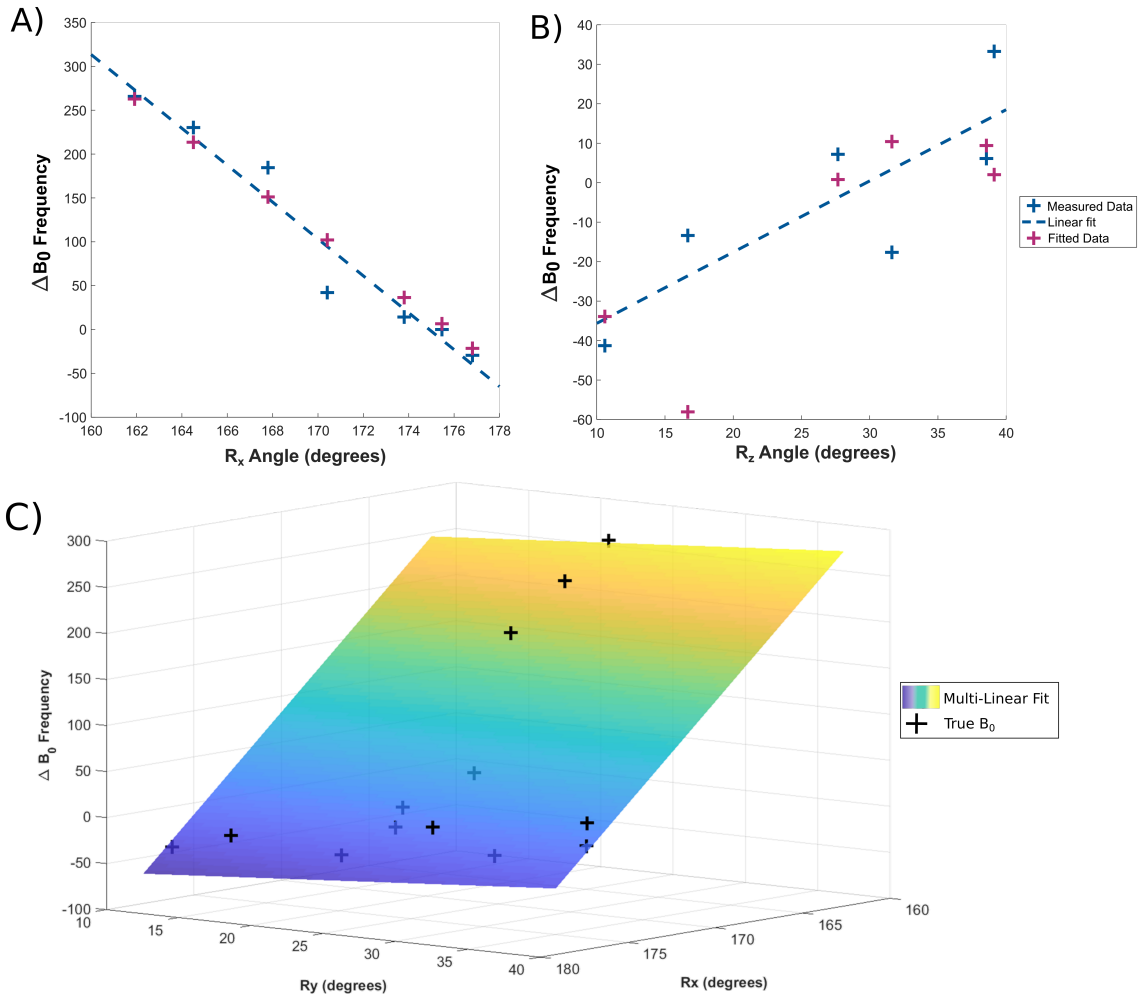


Figure 6.7: The measured B_0 for a voxel close to the sinuses plotted as a function of Rx (A) and Yaw (B). Also plotted is the predicted B_0 value based on Pitch Roll (Ry) and Yaw of the subject in each head position, the scatter is due to small head rotations about orthogonal axes not shown in these plots. The predicted B_0 as a function of Rx and Ry is shown in C).

Figure 6.8 shows the actual B_0 map for the small nod position. From this

one can see the high levels of inhomogeneity near the sinuses. Also shown is the predicted field map from the multilinear fit for the same head position which shows similar structure to the actual field map. The differences between the field maps were very low, with typical differences of approximately 10 Hz. Data were masked to contain only regions inside the head.

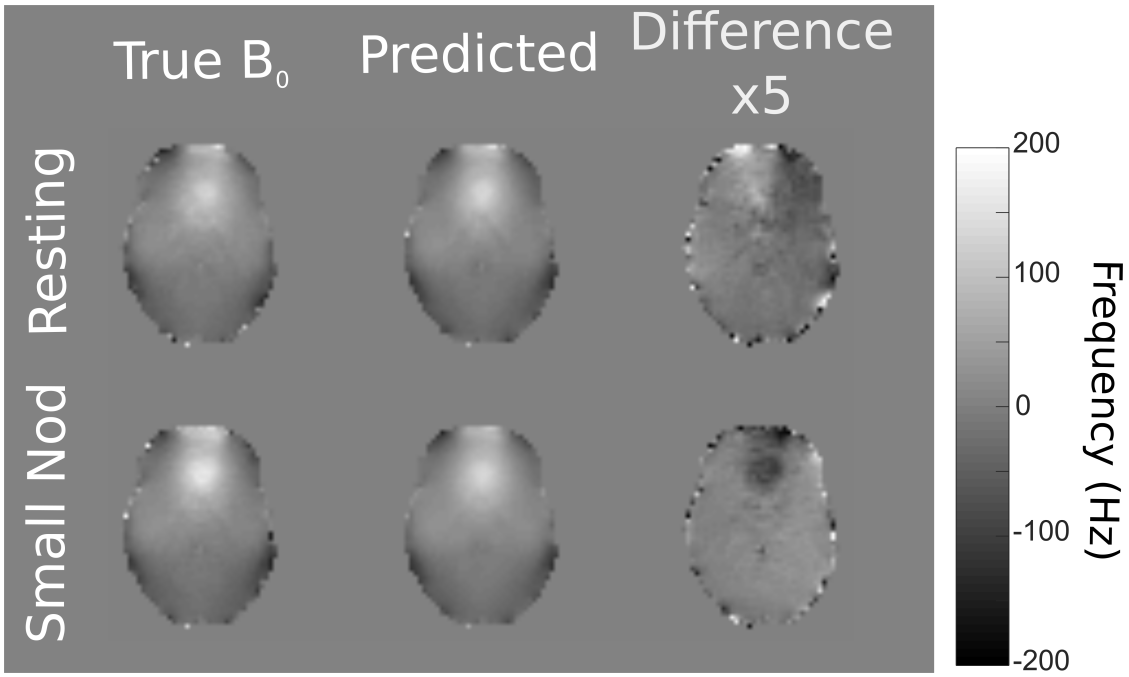


Figure 6.8: The actual B_0 map acquired in the small nod position (A) as well as the predicted B_0 map for that position (B). (C) shows the difference between the two B_0 maps, differences between the true B_0 map and the predicted B_0 were of the order of 10hz.

6.4.2 Distortion Correction

The original EPI images with AP phase encoding show considerable deformation towards the front of the head (Figure 6.9), which matches with the expected results (as this is where B_0 is most inhomogeneous, Figure 6.8). A delay in the prospective motion correction allowed the new head position to be seen before the geometry was updated. Thus, Figure 6.9 shows the head before the motion, after the motion and after the motion has been corrected for using prospective motion correction. It is important to remember that despite Figure 6.9C appear-

ing very similar to 6.9A, it was acquired in the new small shake position with respect to the magnetic field.

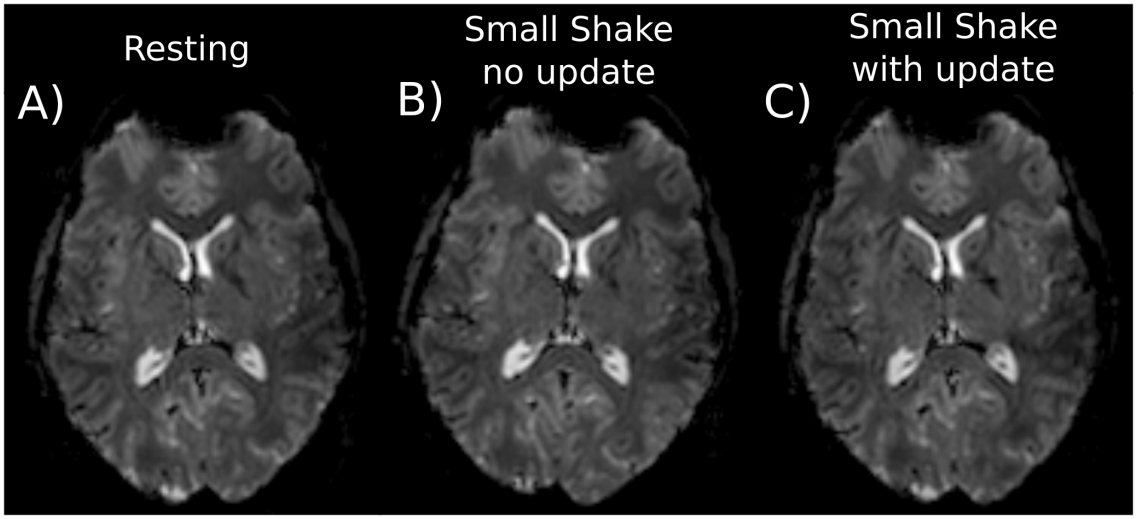


Figure 6.9: The EPI images acquired with subject movement. The Subject was asked to move from resting position (A) to a shake position (B and C). A delay in geometry update shows how the image would look without prospective motion correction (B) and with prospective motion correction (C).

EPI distortion correction was able to greatly reduce the geometric distortions caused by B_0 inhomogeneity. This is most evident towards the front of the head which has been unwarped into its correct shape. Both the base field map and the predicted field map improved the image quality in Figure 6.10, however there are some subtle differences between the two images.

Subjects were instructed to perform a small or large, nod or shake. The amplitude of these movements varied greatly between subjects as shown in Figure 6.11. All but one subject nodded in the same direction (chin to chest) however only half the subjects shook their head left to right.

After geometric distortion correction the EPI images were registered and overlaid on top of the anatomical image. Areas of poor registration can be seen in Figures 6.12, 6.13 and 6.14.

The mutual information between the anatomical and each of the corrected EPI images is shown in Figure 6.15. Since mutual information is effectively a measure of entropy, different subjects had different base levels. However in almost all

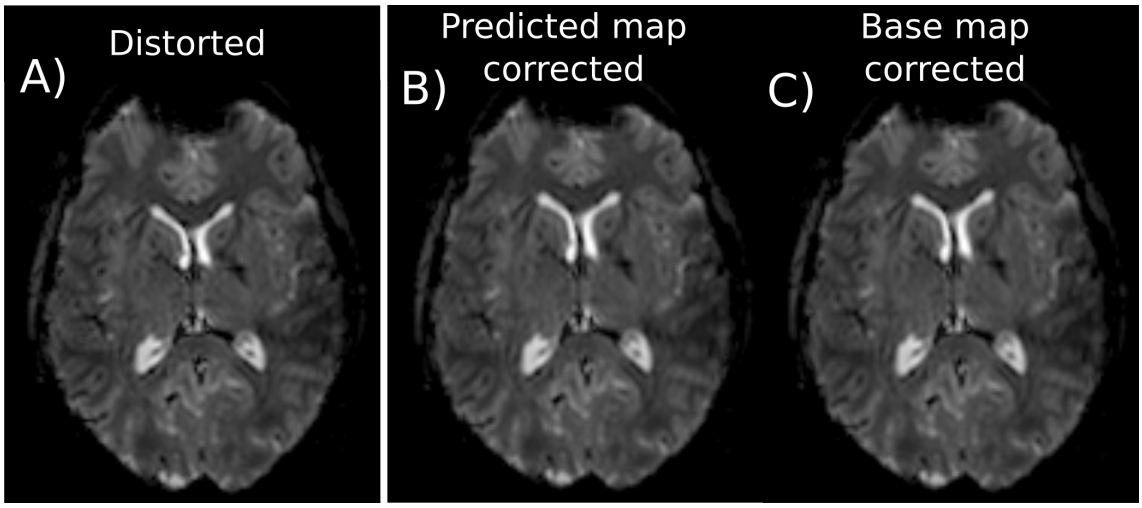


Figure 6.10: The EPI images after geometric distortion correction. (A) shows the EPI image after correction with the base B_0 map, while (B) shows the EPI image after correction with the predicted B_0 map

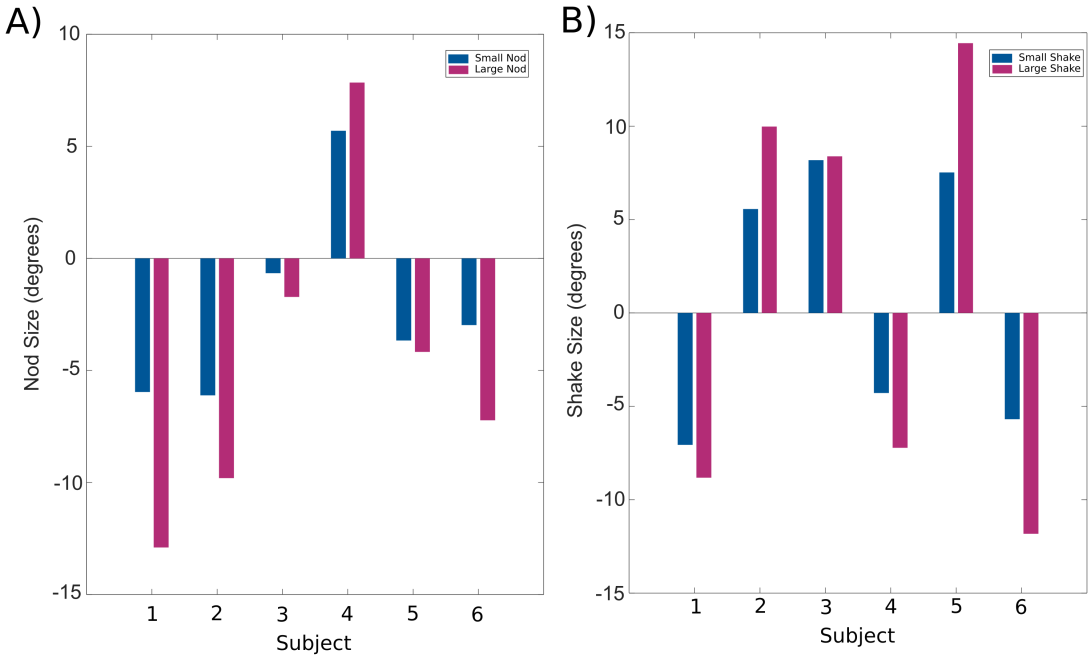


Figure 6.11: The size of the subject rotations for small and large nods (A) and shakes (B).

cases the image registration was improved by using the predicted B_0 maps rather than the base map, with the exception of the “resting” head pose. This is likely because the base B_0 map used to correct the “resting” EPI image was the actual B_0 map for that head pose.

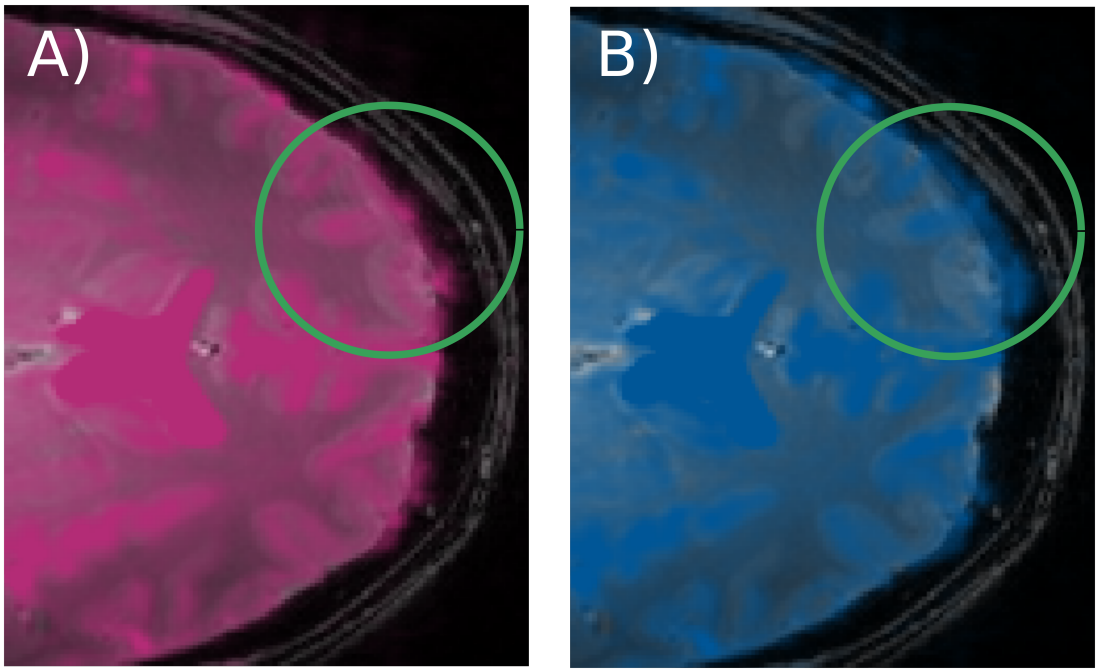


Figure 6.12: The EPI images overlaid on top of the anatomical in the transverse view. (A) shows the EPI image corrected with the predicted field map in pink while (B) shows the EPI corrected with the base field map in blue. The subject is in the small nod position.

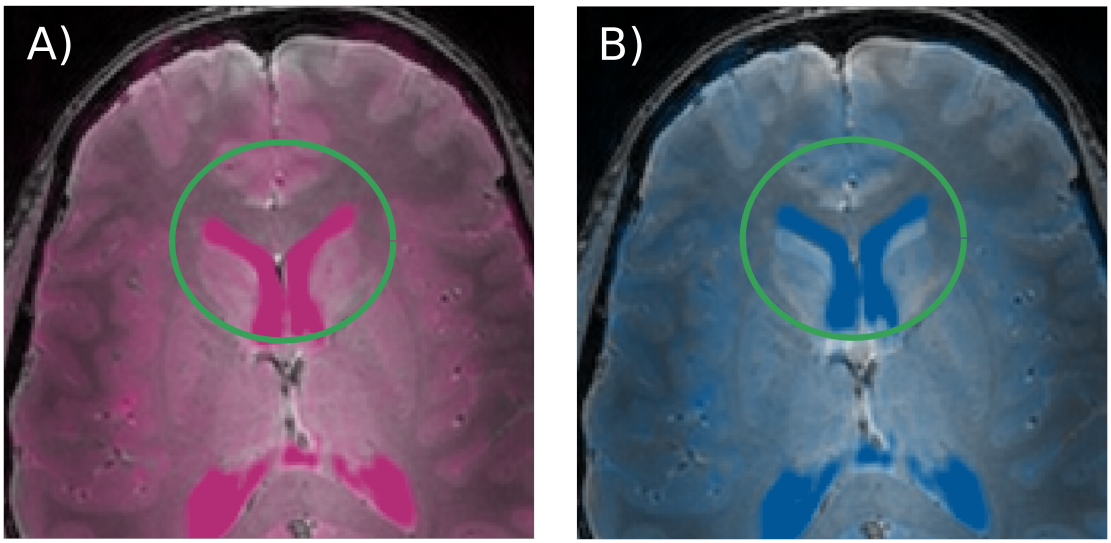


Figure 6.13: The EPI images overlaid on top of the anatomical in the transverse view. (A) shows the EPI image corrected with the predicted field map in pink while (B) shows the EPI corrected with the base field map in blue. The subject is in the large nod position.

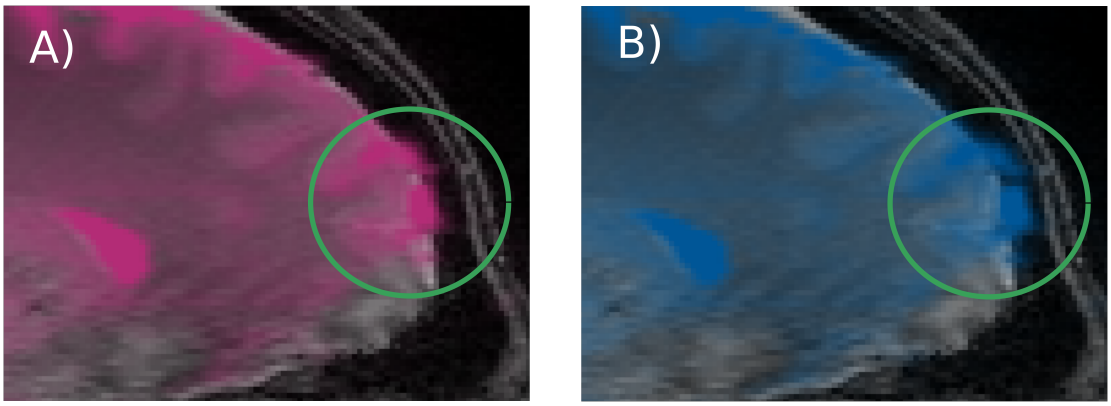


Figure 6.14: The EPI images overlaid on top of the anatomical in the transverse view. (A) shows the EPI image corrected with the predicted field map in pink while (B) shows the EPI corrected with the base field map in blue. The subject is in the small shake position.

Using predicted B_0 maps to correct EPI images resulted in more consistent imaging between head poses. The correlation between the EPI image in the resting position compared to any other head position was higher when the EPI images were corrected with predicted B_0 maps rather than with the base B_0 map as shown in Figure 6.16.

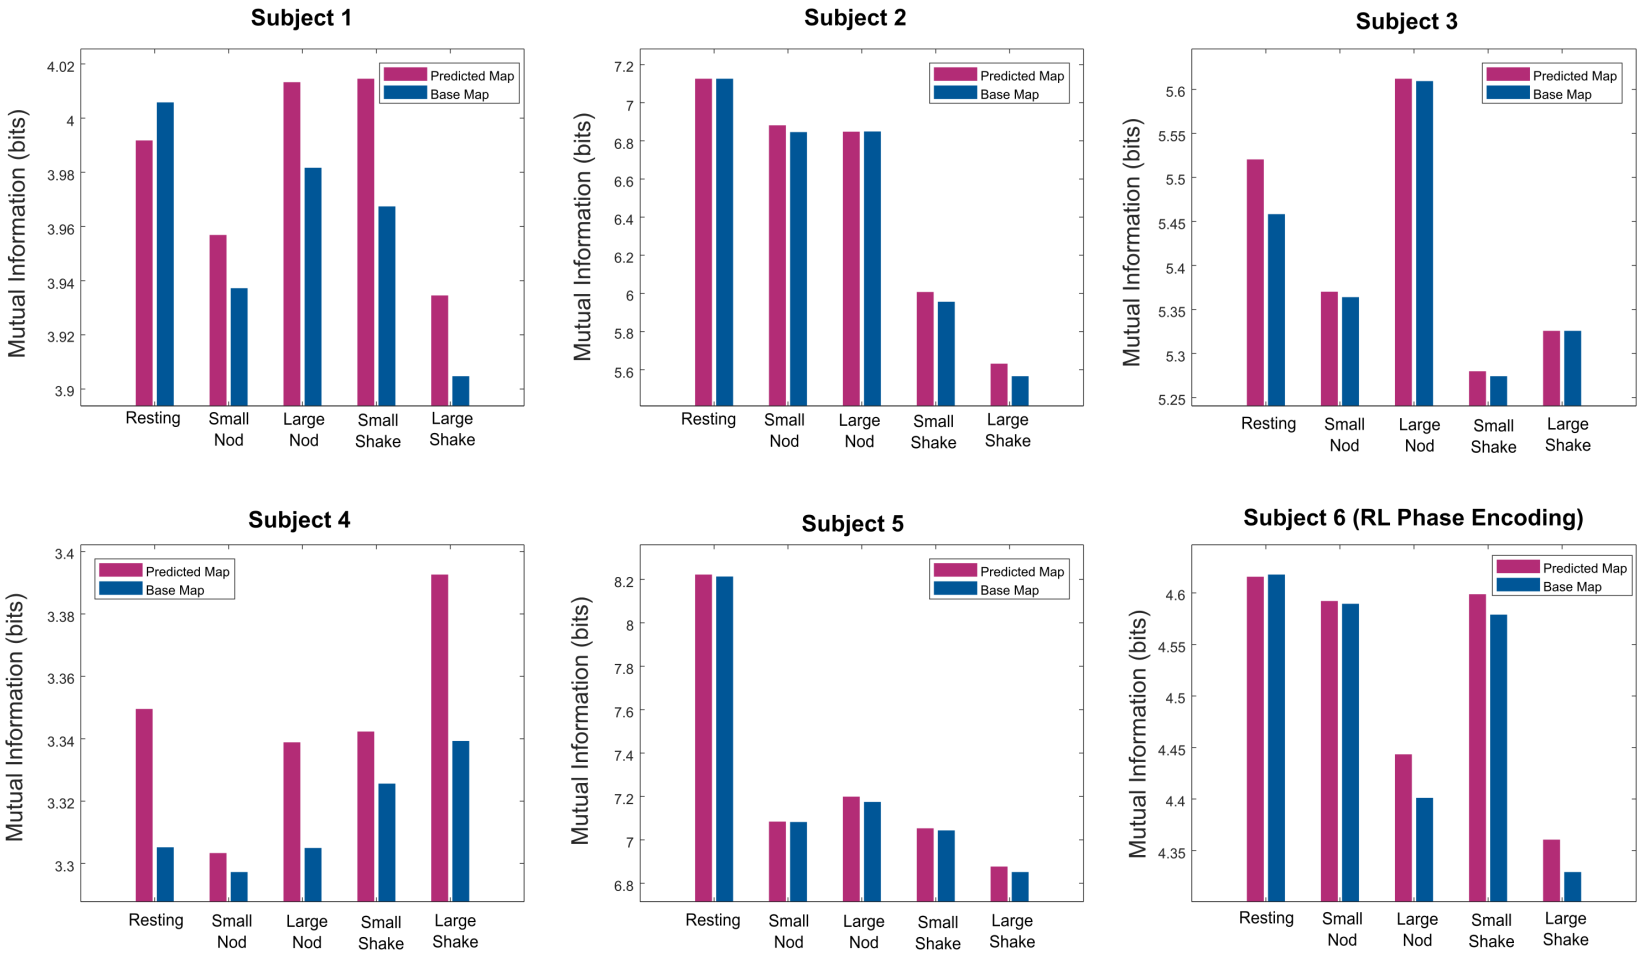


Figure 6.15: The mutual information between the corrected EPI images and the anatomical image for each subject. EPI images were corrected with the "base" B_0 map (blue) and the predicted B_0 map (pink)

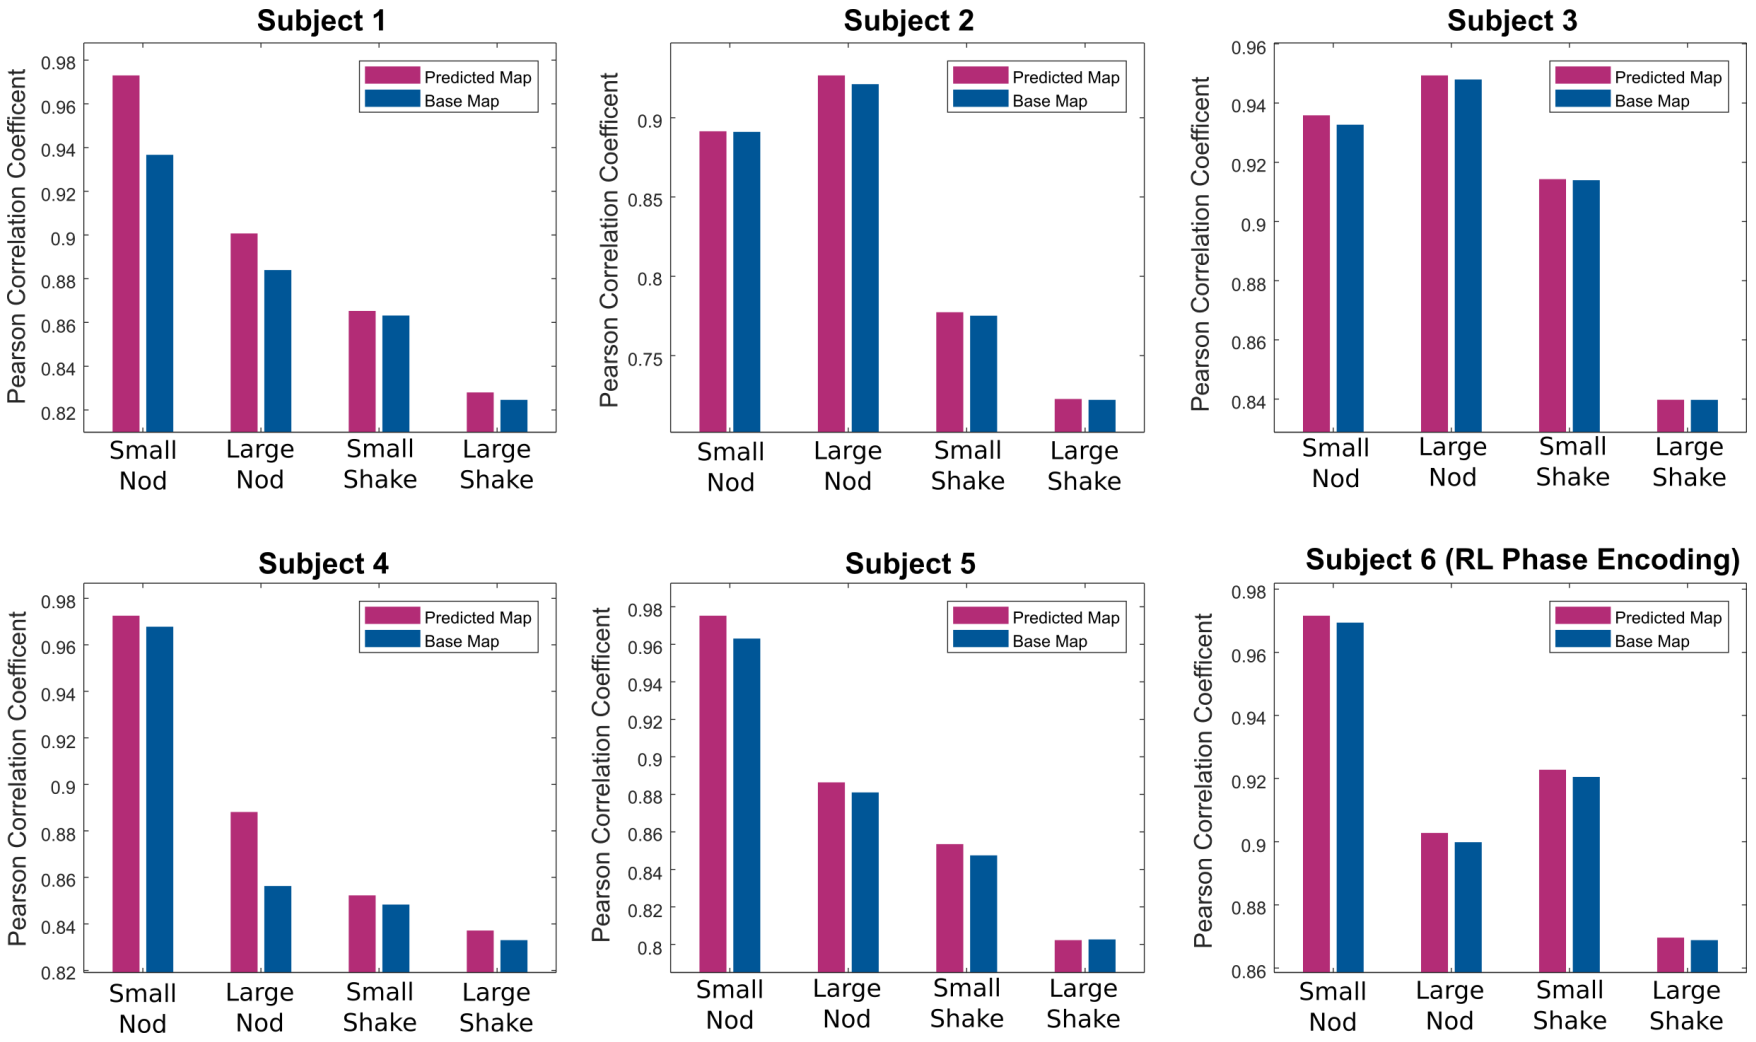


Figure 6.16: The correlation coefficent of the corrected EPI images in any given head position and the original resting position.

6.5 Discussion

6.5.1 Modelling B_0

Figure 6.7 shows that for all head nods and shakes there was a linear relationship between B_0 and head orientation. Subjects were able to nod their heads (Rx) up to approximately 20° and shake their head (Rz) up to 30°. These movements represent the full extent the subject was able to move within the coil. An approximately linear relationship between head orientation and B_0 holds over all reasonable head movements, supporting previous findings [96,98]. Unfortunately co-ordinated movements in Ry were not possible within the coil because of the limited head space available to the left and right of the subject's head due to the coil. In practice there were small variations in Ry when the subject moved to a new head position, but these were convoluted with head movements in Rx or Rz and it is difficult to directly show the relationship between Ry and B_0. The reduced range of movements in Rz mean that the yaw factor (w) is likely to be the most poorly fitted component of Equation 6.7.It should be noted that the positions presented represent when a subject was asked to move specifically in a given direction (either Rx or Rz) where the changes in Rx or Rz were much greater than the accompanying motions in other axes. Fitted points do not appear to be exactly linear because they are the function of Rx or Rz and unseen parameters such as Ry.

As one might expect, the field changes were more dependent on Rx than Rz. Field changes of approximately 80 Hz over 30° of head shakes and 250 Hz over 20° of head nods were observed. This was expected based on theory; head nodding causes the head to rotate around a direction that is orthogonal to the magnetic field changing the incident angle between the field and the head, head shakes are a rotation about the magnetic field direction and so the angle of incidence is the same. This theory is supported by the higher u and v values compared to w in Equation 6.7.

The predicted B_0 field maps showed strong similarity to the actual B_0 field maps. Both showed high levels of inhomogeneity towards the front of the head

of similar orders of magnitude. The largest variations with a B_0 map were seen around the outside of the head; where the actual B_0 maps is subject to spikes where the mask was not sufficiently eroded, the predicted B_0 maps were considerably more resistant to anomalous B_0 spikes because they are produced using a number of field maps where spiking is inconsistent. This spiking is common and hence when a B_0 map was used to unwarp the EPI images they were smoothed and de-spiked. Differences between the predicted B_0 map and the actual B_0 maps were of the order of 10 Hz, this corresponds to a pixel shift of approximately 0.5 pixels in these EPI images.

6.5.2 Distortion Correction

It is clear from comparison of Figure 6.9A to Figure 6.9B that the change in head orientation has altered the geometric warping at the front of the head, this is due to the combined effect of the change in B_0 and change in phase encoding direction (relative to the head) by motion correction. Figure 6.9C shows that prospective motion correction alone offers considerable improvements in distortion correction by maintaining the phase encoding direction. That being said some more subtle differences between A and C still exist due to this difference in B_0 between the head positions.

By eye, the correction with both the base B_0 map and the predicted B_0 map result in similar levels of improvement (Figure 6.10). The differences become much more obvious when the images are registered to an anatomical (Figures 6.12, 6.13, 6.14). Here you can see the original base corrected image results in deformation where the brain has been excessively unwarped, while the predicted field map aligns well with the anatomical. This finding is supported by the measured mutual information which was greater when EPI images were corrected with predicted B_0 map than with the base B_0 map, although varied largely between subjects due to the different base entropy of the subjects head. The improvement was seen in almost all head positions with the exception of the resting position. The resting image is better corrected with the base image because in

this case the base B_0 map is the true B_0 map for that position. It is also noted that the greatest improvements are seen with the head nodding movements rather than the head shake movements. Figure 6.7 showed that the largest changes in B_0 occurred with head nodding, which means that the base map is likely to be further from the actual B_0 when the when the head moves in Rx than when it moves in Rz.

fMRI experiments are reliant on the alignment of volumes throughout a time course. Figure 6.16 shows that images that are corrected using a predicted B_0 field map are better registered to each other than those corrected with the base B_0 field map. To measure the goodness of the registration we were able to use a pearson correlation co-efficient since both images were EPI images with the same contrast.

6.5.3 Optimisation

Currently we acquire 12 B_0 maps from which we generate fitting parameters for equation 6.7. This is shown to be a suitable number of maps to acquire a good fit. However it has not been optimised, it is possible that fewer maps could be used, thus decreasing acquisition time. Under the current implementation 4 B_0 maps are acquired in the initial resting pose, 4 in various head nod poses and in various head shake poses. This number could be reduced by reducing the number of maps acquired in the resting pose and replacing pure head movements (exclusively either a head nod or head shake) with mixed head rotations (simultaneous head nodding and shaking). This may also have the added benefit of increasing subject rotation around the AP axis.

6.5.4 EPI Phase Artefact

EPI images with large phase artefacts were repeated with smaller movements, however Figure 6.15 shows that overall head shakes had lower mutual information with the anatomical images. This is probably due to smaller less obvious phase artefacts present in the head shake EPI images. The largest EPI phase artefact

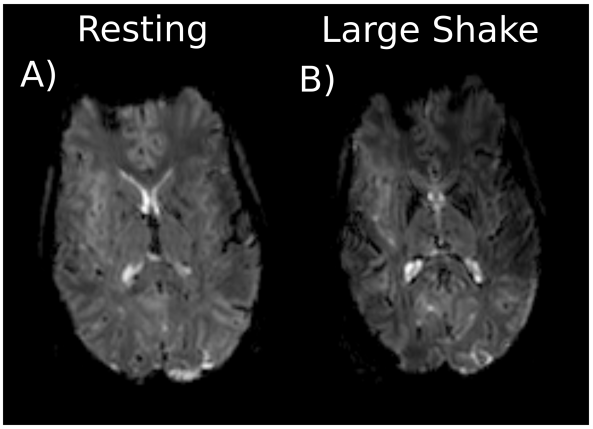


Figure 6.17: The original distorted EPI images for the resting (A) and large shake position (B) for the largest shake observed. There are no artefacts (besides the geometric distortion) present in (A) however (B) shows distortions and artefacts related to the phase artefacts.

is shown in Figure 6.17.

6.5.5 Alternative Correction Methods

In section 6.1 we discussed an alternative method of EPI distortion correction to FUGUE, known as TOPUP, however more recently alternative methods have been proposed for fMRI image correction including ICA-AROMA [99] and Eddy [100].

ICA-AROMA is a tool that uses independent component analysis to split data into components. It has been trained to identify components that are due to motion and remove them, such that when the data is reformed, motion artefacts are reduced. This method then is capable of correcting many motion artefacts simultaneously, however its ability to identify B_0 changes due to subject motion is unclear.

Eddy is a method of correcting diffusion weighted EPI images. This method uses an iterative approach to solve for eddy currents and subject motion. Initial estimates are made for each parameter and used in combination with a B_0 susceptibility map to unwarp each image. Each image is then passed into a "Gaussian Process" prediction algorithm described by *Andersson and Sotiropoulus* [101]

which assumes signals from each voxel vary smoothly from each image to the next. This generates a smoothed average of the input images, known as a prediction image. This predicted image is then deliberately warped using the B_0 map and estimates for eddy current and motion. The warped predicted image is compared to the original distorted image, and used to update the eddy current and motion parameters. This sequence then repeats for a number of iterations (usually 5) until the image distortions are sufficiently low. This method however, assumes separation of motion and B_0 which is not the case for large subject movements where head orientation changes B_0.

6.6 Conclusion

Here we have demonstrated the ability to prospectively correct inter-volume head motion using an optical tracking camera. This prospective motion correction was able to restore some of the image distortions associated with changes in B_0 and the phase encoding direction that retrospective correction could not preserve. We combined our prospective correction with retrospective geometric distortion correction to unwarp EPI images that had previously suffered from B_0 artefacts. We used predicted B_0 maps to estimate the pixel shift required to unwarp images. We tested these predicted B_0 maps in a series of different head positions and in nearly all cases they improved both the registration to anatomical images and resulted in more consistent EPI images between head poses.

Chapter 7

Electromagnetic Position Sensor for MRI

7.1 Introduction

In this thesis several external motion tracking devices used in MRI motion correction have been described, the most common of which are explained in Chapter 3.

External tracking devices, such as NMR probes [102] and optical tracking cameras [88, 89, 103, 104] offer the highest level of accuracy and temporal resolution compared to other motion detection technology. However they are also relatively expensive, with costs often in the region of £100,000 (*Kineticor, 1st Generation Moiré Phase Tracking Camera*). Because of this alternative tracking devices have been proposed.

Gradient pick-up coil systems [105–108] have been proposed to track subject motion. These systems use multiple pick-up coils affixed by some suitable method to a subject's head to localise position. Scanner gradient pulses are then applied which induce voltages in each of the coils, from which the position of each coil can be determined. However in order to determine head orientation, multiple rigidly connected coils must be used.

Recently *Bhuiyan et al.* [109] proposed using multiple head mounted, non-rigid coils. The relationship between the induced voltages and head movement was determined using known movements and decomposing the detected voltages with principle component analysis (PCA). A general linear model (GLM) was used to determine subsequent movements during the experiment. However, thus far all pick-up coil systems are based on a numerical fitting mechanism and require the use of multiple coils. If the orientation of the coil could be determined independently, only a single coil would be required and the position of the coil could be determined empirically or using similar methods to those described above.

Niekerk et al. [110] recently used a magnetometer and accelerometer to describe the orientation of an object within an MRI scanner. This method uses known field directions (gravity and B_0) to determine subject orientation, by measuring the angle between each field vector and the sensor. This method was shown to provide high accuracy using a relatively compact sensor. However while

this method was used to qualitatively describe subject movement, it did not provide a full quantitative description of subject movement because no information about subject translation was available and hence it could not be used for full 6 degrees of freedom motion correction.

Aim: To create a sensor capable of detecting its position and orientation within the scanner. We propose using a variant of the *Niekerk et al.* sensor which uses an Anisotropic Magneto Resistive (AMR) sensor instead of a magnetometer, combined with a *single* pick-up coil to measure the position and orientation of the sensor within a MRI scanner. Because the sensor derives its position from the scanner gradients, the sensor is self calibrating (i.e. already in the magnet coordinate scheme). The accuracy of this device was determined by comparing the sensor to existing motion tracking technology.

Aspects of the work described in this chapter formed the basis of the following abstract: *James Smith, Glyn Spencer, Richard Bowtell, Penny Gowland, and Paul Glover. A novel position and orientation sensor for MRI. In Proceedings of the International Society for Magnetic Resonance in Medicine, Number 3930, 2017.* [111]

7.2 Theory

The proposed electronic sensor utilises a combination of components to determine its orientation. This orientation information is then used in combination with a gradient pick-up coil to determine its position within the magnet.

7.2.1 Determining Orientation

A Micro-Electro-Mechanical accelerometer (MEMS Accelerometer) was used to determine the direction of gravity. For relatively slow movements the acceleration due to movement is negligible and the dominant acceleration is due to gravity. The MEMS accelerometer contains a small “proof mass” suspended in a substrate by a number of anchor points. This proof mass acts as an insulating layer

between two conductors and hence forms a capacitor with variable capacitance. When the accelerometer is accelerated the proof mass is displaced altering the capacitance which in turn changes the output voltage. The proof mass has 3 degrees of freedom allowing for 3 output voltages corresponding to the acceleration in each orthogonal direction r. The voltage from each output channel is linear with acceleration and can be written as,

$$g_r = 5\left(\frac{V_r}{V_s} - \frac{1}{2}\right), \tag{7.1}$$

where V_r is the output voltage for the, x, y, or z channel on the accelerometer, V_s is the supply voltage and g_r is the acceleration in a particular axis measured in g's [112]. This provides an acceleration in each axis of the sensor. Orientation about the scanner's x (LR) and z (FH) axes can then be determined from the acceleration vectors. The inclination about the scanner's z axis (ψ) can be found using,

$$\psi = \arctan\left(\frac{g_x}{g_y}\right). \tag{7.2}$$

Knowing ψ, the angle around x (ϕ) can then be determined,

$$\phi = \arctan\left(\frac{-g_z}{(g_x \sin\psi) - (g_y \cos\psi)}\right). \tag{7.3}$$

However rotations about the scanner's y (AP) axis (θ) are independent of the direction of gravity and so cannot be determined using the accelerometer alone.

To determine the orientation about the AP axis the magnetic field vector from the static B_0 field is determined. To measure this we used an AMR sensor. Magnetoresistance is a property of some materials to change their electrical resistance based on an applied magnetic field. Anistropic magnetoresistance is the property of a material to change its electrical resistance based on the angle at which a magnetic field is applied to the material. This effect is caused by the magnetisation of the material and the spin-orbit interaction. The direction of magnetisation is dependent on the direction of the applied field. In anisotropic magnetoresistive materials increased scattering of s-d electrons occurs in the direction

of this magnetisation [113]. This causes the resistance of the material to increase when the applied field is aligned with the direction of current.

The AMR sensor contains 8 strips of AMR permalloy to form two wheatstone bridge sensors [114] (see Figure 7.1). The second wheatstone bridge (bridge β)is rotated in-plane by 45° relative to the first bridge (bridge α).

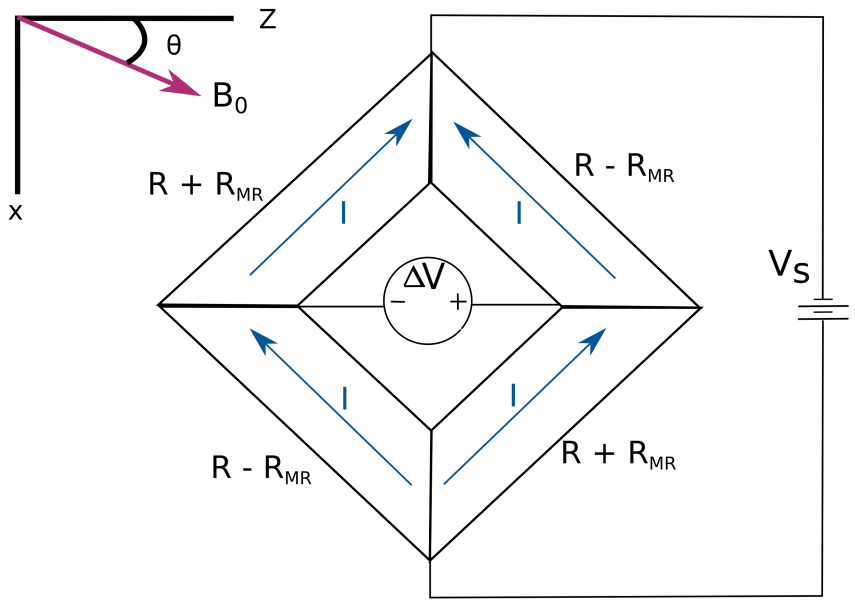


Figure 7.1: An AMR bridge used in the AMR sensor. Each bridge contains 4 strips of anisotropic magneto resistive material with a base resistance R and a magnetic field dependent resistance R_{MR}. A supply voltage (V_s) is passed over the bridge and the output voltage measured (ΔV). The measured voltage is orientation dependent.

The change in resistance of an AMR strip is related to the angle of field by,

$$R_{MR} = \cos^2(\theta), \tag{7.4}$$

where θ is the angle between the current flow ($\mathbf{I}$) and the magnetic moment $\mathbf{M}$. Assuming the applied magnetic field is large enough to saturate the AMR (> 60 Gauss) [114] the voltage output from the first bridge shown in Figure 7.1 is equal to,

$$\Delta V_\alpha = V_s S \sin(2\theta), \tag{7.5}$$

where S is a material constant. The second bridge which is rotated by 45° gives a voltage of,

$$\Delta V_\beta = V_s S \cos(2\theta). \tag{7.6}$$

The angle between the current and the magnetic field can then be obtained by measuring the ratio of these two output voltages.

$$\psi = \frac{1}{2} \arctan \frac{\Delta V_A}{\Delta V_B}, \tag{7.7}$$

however at $\psi = \pm 45$° arctan becomes poorly defined, so additional rules apply:

$$\begin{aligned}
&\text{if} \Delta V_\alpha = 0 \text{ then, } \theta = 0 \\
&\text{if} \Delta V_\beta = 0 \text{ and } \Delta V_\alpha < 0 \text{ then, } \theta = -45 \\
&\text{if} \Delta V_\beta = 0 \text{ and } \Delta V_\alpha > 0 \text{ then, } \theta = +45 \\
&\text{if} \Delta V_\alpha < 0 \text{ and } \Delta V_\beta < 0 \text{ then, } \theta = \theta - 90 \\
&\text{if} \Delta V_\beta > 0 \text{ and } \Delta V_\beta < 0 \text{ then, } \theta = \theta + 90
\end{aligned} \tag{7.8}$$

7.2.2 Determining Position

When determining the orientation of the sensor it is assumed that subject movement is small and that the measured acceleration is due to gravity entirely. In theory if the acceleration due to movement could be separated from gravity, the acceleration could be integrated with respect to time to give the velocity of the subject, which could be integrated again to give the change in position of a person over that time. However as shown in Chapter 5 a *change* in position of a marker (or in this case sensor) is not sufficient information to use for prospective motion correction. This is because the transform between the marker (or sensor) and the imaging volume must be known on initialisation (see Chapter 5). Because of this, the *absolute* position of the sensor must be known; which can't be obtained from the accelerometers alone, hence a coil is used to detect gradients pulses.

Faraday's law states that a coil placed within a changing magnetic flux (Φ) will

experience a electromotive force (η);

$$\eta = -\frac{d\Phi}{dt}. \tag{7.9}$$

This force causes a detectable voltage in the coil by the same principle of induction (as used to detect the MR signal using coils tuned to detect the RF signal). By placing the coil in the scanner and using 3 gradient pulses in each of the scanner directions, known positions in x, y, z can be determined relative to gradient iso-centre based on the voltages induced in the coil from each gradient pulse.

For a gradient in the scanner z direction, the projection of that field onto the device frame ($\hat{x}, \hat{y}, \hat{z}$) can be written as:

$$\hat{\mathbf{B}}_\mathbf{z} = \mathbf{R}_z(\psi)\mathbf{R}_y(\theta)\mathbf{R}_x(\phi)\mathbf{B}_\mathbf{z} \tag{7.10}$$

where,

$$\mathbf{R}_z(\psi)\mathbf{R}_y(\theta)\mathbf{R}_x(\phi) = \begin{bmatrix} \cos\psi\cos\theta & \cos\phi\sin\psi + \cos\psi\sin\theta\sin\phi & \sin\psi\sin\theta - \cos\psi\cos\phi\sin\theta \\ -\cos\theta\sin\psi & \cos\psi\cos\phi - \sin\psi\sin\theta\sin\phi & \cos\psi\sin\phi + \cos\phi\sin\psi\sin\theta \\ \sin\theta & -\cos\theta\sin\phi & \cos\theta\cos\phi \end{bmatrix} \tag{7.11}$$

Assuming a circular coil of size A with N turns aligned so that the $\hat{x}$ component is normal to A, the voltage recorded due to the applied fields would be.

$$\begin{aligned} V_z &= -AN\frac{dB_{zx}}{dt} \\ V_y &= -AN\frac{dB_{yx}}{dt} \\ V_x &= -AN\frac{dB_{xx}}{dt} \end{aligned} \tag{7.12}$$

where B_{zx} is the x component of $\mathbf{B_z}$:

$$\begin{aligned}\hat{B}_{zx} = G_z(z\cos\psi\cos\theta\ldots \\ -\frac{1}{2}x(\cos\phi\sin\psi+\cos\psi\sin\theta\sin\phi)\ldots \\ +\frac{1}{2}y(\sin\psi\sin\theta-\cos\psi\cos\phi\sin\theta))\end{aligned} \tag{7.13}$$

similarly $\hat{B}_{xx}$ and $\hat{B}_{yx}$ can be written as :

$$\hat{B}_{xx} = G_x(x\cos\psi\cos\theta + z(\cos\phi\sin\psi+\cos\psi\sin\theta\sin\phi)) \tag{7.14}$$

$$\hat{B}_{yx} = G_y(y\cos\psi\cos\theta - z(\sin\psi\sin\phi-\cos\psi\cos\psi\sin\phi)) \tag{7.15}$$

Equations 7.13, 7.14, 7.15 and 7.12 can then rewritten as a matrix:

$$\begin{bmatrix}V_x\\V_y\\V_z\end{bmatrix} = -\Gamma ANG\dot{\Theta}\begin{bmatrix}x\\y\\z\end{bmatrix} \tag{7.16}$$

where the additional term Γ represents any signal amplification applied during signal acquisition, $\dot{G}$ is the derivative of gradient strength with respect to time assuming that $\dot{G}_x = \dot{G}_y = \dot{G}_z = \dot{G}$, and,

$$\Theta = \begin{bmatrix}\cos\psi\cos\theta & 0 & \cos\phi\sin\psi+\cos\psi\sin\theta\sin\phi\\ 0 & \cos\psi\cos\theta & \sin\psi\sin\phi-\cos\psi\cos\psi\sin\phi\\ \cos\phi\sin\psi+\cos\psi\sin\theta\sin\phi & \sin\psi\sin\theta-\cos\psi\cos\phi\sin\theta & \cos\psi\cos\theta\end{bmatrix} \tag{7.17}$$

Equation 7.16 can then be solved by matrix inversion to give the position of the coil using,

$$\mathbf{r} = -\frac{1}{\Gamma ANG\dot{}}\mathbf{\Theta}^{-1}\mathbf{V} \tag{7.18}$$

where $\mathbf{r}$ is a 1×3 vector containing the x, y, z co-ordinates and and $\mathbf{V}$ is a 1×3 vector of the voltages.

7.2.3 Determining Sensor Error

Orientation

The error in the accelerometer is quoted as 50μg Hz$^{-\frac{1}{2}}$ [112]. For prospective motion correction, geometry data is required frequently, to maximise temporal resolution of the detected subject movement. For commercial devices this frequency is of the order of 100 Hz. Using 100 Hz as the estimated bandwidth of the device, g_x, g_y and g_z can then be measured with an accuracy of $\pm$ 0.5 mg. This error propagates into the angles measured by the accelerometer. The error was calculated by determining the additional area swept out by the error in acceleration as seen in the example Figure 7.2.

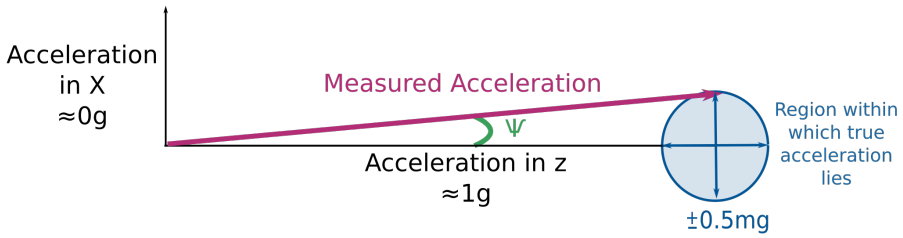


Figure 7.2: The propagation of errors from the accelerometer into ψ. The noise in the accelerometer causes the fluctuations in gradient vectors which creates an area of uncertainty, the angle swept out by this area describes the error in the measured angle

The sensor is positioned such that the z-axis of the accelerometer is approximately aligned with gravity. ($g_z \approx 1, g_y \approx 0, g_x \approx 0$, the errors from the accelerometer cause an estimated error of $\pm$0.03° in ψ. The errors in ϕ are significantly larger due ϕ's dependency on g_x, g_y and g_z. Assuming the error in ψ, is small the error in ϕ can be calculated using the same approach as was used for ψ, which gives an error of $\pm$0.06°. The error in θ from the AMR is quoted as 0.05° [114]. However, the angle measured by the AMR is averaged over a same time window as the accelerometer (0.01 s) at a sampling rate of 5 kHz and then averaged. The actual error in θ is then 0.007°.

Position

The estimated error in position was determined numerically. A typical position of the sensor was defined (20 mm away from the isocentre in each axis and orientated along the z-axis of the scanner), and corresponding voltages calculated using Equation 7.16, assuming no error.

Errors were then added to each component of Equation 7.16 separately to determine the maximum error in the calculated position. ψ had an error of $\pm$ 0.03°, ϕ had an error of $\pm$ 0.06°, θ had an error of $\pm$ 0.007°, the error in the measured voltage had an error of 1 μV as determined by the EEG amplifier(in practise errors in the measured voltage may be larger than this if Γ is large). The position of the sensor was then recalculated in the presence of each of these errors and the difference in position determined.

Negligible errors were observed in the presence of errors in measured voltage or angle ϕ. Errors of 0.1 mm were observed in the presence of 0.007° error in θ. The largest errors observed were due to the errors in ψ which cause a 0.5 mm error along the scanner x direction. These errors were added in quadrature to give an estimated error in position of $\pm$ 0.54 mm.

7.3 Method

7.3.1 Sensor Construction

The sensor was constructed on a printed circuit board measuring 21 mm by 21 mm (built by *Andrew Stewart, University of Nottingham, School of Physics and Astronomy*). The MEMS accelerometer [112] (LIS344) and the AMR sensor (APS00B) [114] were mounted parallel to each other. The supply voltage onto the board was 5 V, which was used by the AMR sensor. However this was regulated using a voltage regulator (LP2985) [115], to maintain the supply voltage at 3.3V for the accelerometer. An additional temperature sensor (TC1047) [116] was mounted on to the circuit board. The primary use of the temperature sensor was to monitor (and if needed compensate for) heating in the accelerometer. A

secondary advantage was the ability to measure possible heating in the coil due to RF interaction. A picture of the sensor can be seen in Figure 7.3 and a circuit diagram can be found in appendix A.

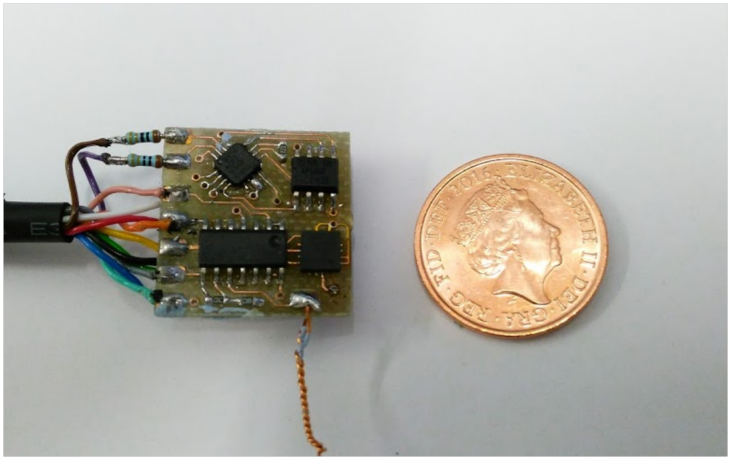


Figure 7.3: An image of the sensor board that contains the components required to determine the sensor orientation within a magnetic field. A British penny is added for comparison.

A circular copper coil was wound around a plastic spool of width 25 mm for 20 turns. Care was taken to ensure turns were wound tightly with no overlap. The sensor board was then mounted securely within the plastic spool to prevent movement occurring between the sensor board and the coil (Figure 7.4).

The sensor outputs a total of 8 voltage signals; 3 signals from the accelerometer corresponding to the accelerometer x, y, z channels, 2 signals from the AMR corresponding bridges α and β, 1 from the coil, 1 from the temperature sensor and 1 from the voltage regulator.

The voltages from the sensor were read in using a *Brain Vision* EEG Bi-polar amplifier and USB 2 Adapter with a read rate of 5000 Hz. Before being passed to the amplifier, the sensor outputs were passed to a separate circuit outside of the magnet. The output voltages from the components were stepped down to be between $\pm$ 32 mV to match the range of the maximum range of the amplifier. This transform was performed using a potential divider circuit ($R1 = 330\Omega$ and $R2 = 4700\Omega$). A constant power of 5 V was supplied to the sensor using the amplified battery pack.

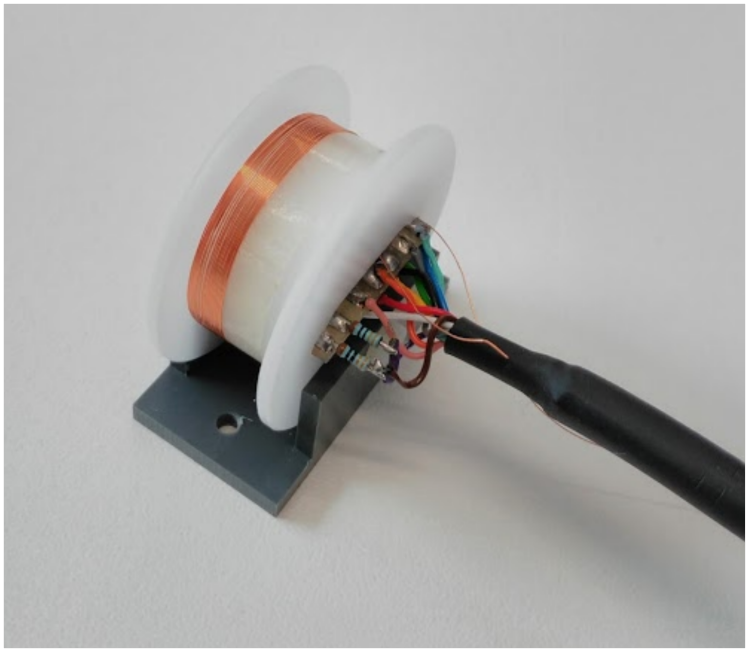


Figure 7.4: An image of the sensor coil. The sensor board was mounted inside of the copper coil housing to ensure no movement could occur between the coil and board.

7.3.2 Controlling Movement

To test the accuracy of the sensor, the sensor was mounted on to a “rig” (Figure 7.5) (Built by *Alan Dorkes, University of Nottingham, School of Physics and Astronomy*), onto which a Moirè phase tracking (MPT) marker was also placed. This rig enabled precise rotations and translations to be made, accurate to approximately 0.25 mm and 0.5°. Rig translations were performed by turning dials which were mounted to screws. This screw motion resulted in a translations of 1 mm per revolution. Rotations were achieved by exploiting the lever action effect. The rig contains a pivot point in each rotational axis, by displacing the rig from one side of the pivot, a rotation about this pivot axis can be performed.

However the pivot line for each axis did not intersect, for example a rotation in x rotated about a different point compared to a rotation about y. The rig was mounted into a Philips 3 T Achieva scanner on a beam that was suspended through the bore. This enabled the sensor to be decoupled from scanner vibrations.

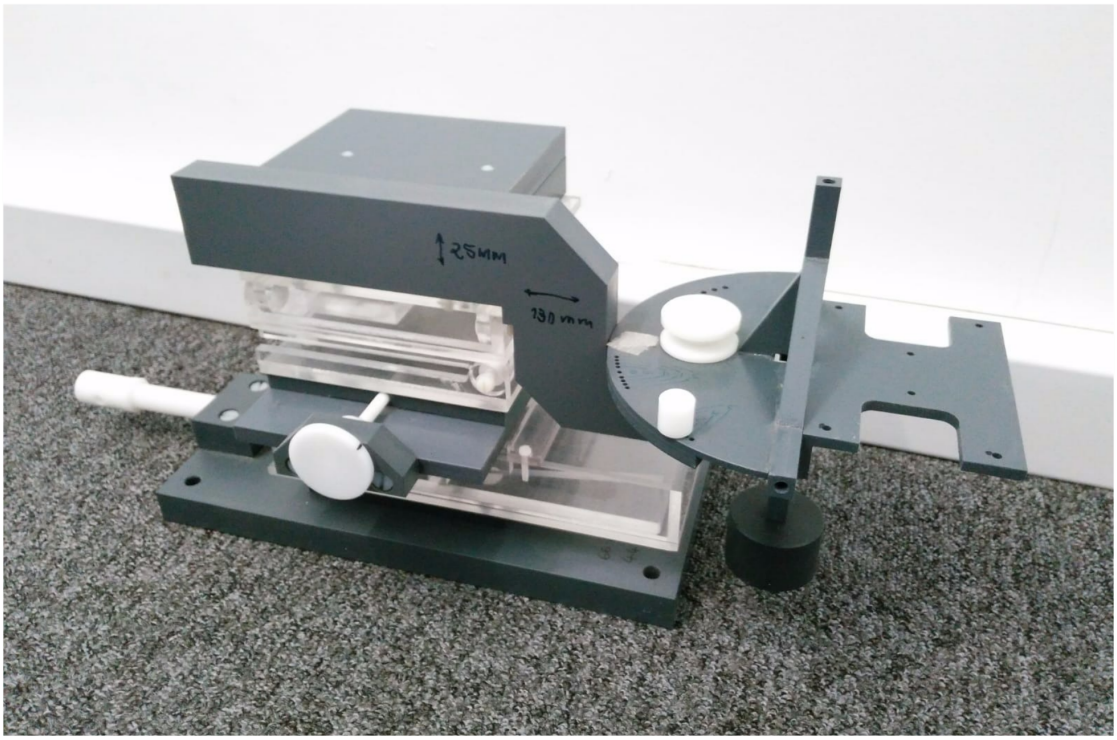


Figure 7.5: The rig used to control the motions of the sensor

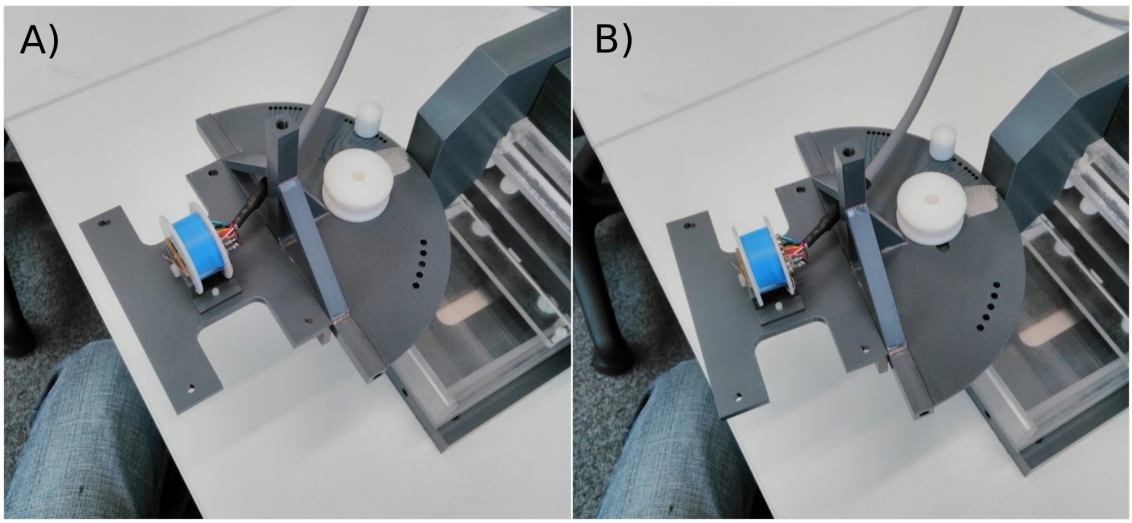


Figure 7.6: The rig with the sensor mounted to it. The images show how the sensor can be rotated by known amounts, in this case a rotation of 8 degrees between images A and B.

7.3.3 Calibrating the Sensor Coil

In theory all parameters in Equation 7.18 can be either measured ($\mathbf{V}$), known by design (Γ, A, N) or are selectable parameters on the scanner (G). However due

to construction the exact response of the coil to a gradient may vary. Because of this it was necessary to determine the constant k experimentally, where:

$$k = \Gamma ANG. \tag{7.19}$$

The rig and sensor were aligned with the scanner so that the sensor could be translated purely along the z axis. When the sensor is aligned in this way $\psi, \theta, \phi = 0$ and Θ is an identity matrix (Equation 7.2.2).

The sensor was discretely translated along the z-axis in 1 mm increments up to 5 mm. In each position gradient pre-pulses were applied. Gradient pulses were set to have a strength of 20 mTm^{-1}ms^{-1}, a slope duration of 4 ms and a peak duration of 5 ms with 25 ms between each pre-pulse. Each set of gradient pre-pulses (G_x, G_y, Gz) was repeated 15 times (45 gradient pre-pulses per pose) with 0.4 s between pulse sets. The gradients induced position dependent voltages in the coil.

The induced voltages were then integrated over the pulse duration and the maximum integrated voltage taken. The average over the 15 repeats was taken to give the mean-max-integrated voltage per pose.

The mean-max-integrated voltage was plotted as a function of translation in z and a first order polynomial line fitted where the gradient of the line is equal to k.

7.3.4 Sensor vs Optical Tracking

To validate the accuracy of the sensor for motion correction, it was compared to a MPT camera.

The camera was mounted into the bore of a 3 T Philips Acheiva MRI scanner and calibrated using the methods described in Chapter 4. A MPT marker was placed on to the rig using pressure sensitive double sided adhesive tape and the sensor which was screwed on to the rig so that they were rigidly coupled (Figure 7.7).

The rig was placed into the bore using the same method as described in Section 7.3.3, whereby the rig was placed on a beam through the scanner so that the

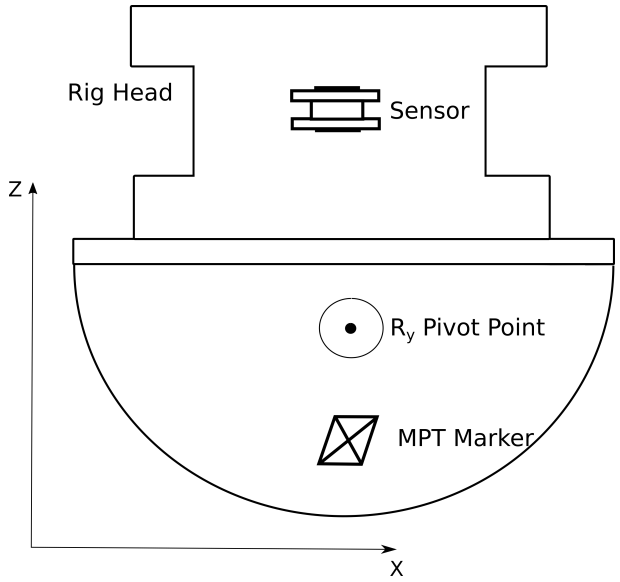


Figure 7.7: A diagram of the head of the rig where the sensor and MPT marker were mounted.

sensor was decoupled from the scanner vibrations. A series of discrete movements were performed using the rig to control the motion size. The sensor was first moved using a series of pure movements: translated along the z axis in 1 mm increments up to 5 mm and rotated around the x axis in 1° increments up to 5°. The sensor then underwent a series of mixed movements: rotated by 3 degrees in the z axis and translated along x in 1 mm increments up to 3 mm, the sensor was returned to it's initial position and then rotated by 4° in y and then translated along y in 1 mm increments up to 3 mm.

In each pose a total of 45 gradient pre-pulses were performed (15 repeats of each gradient axis pulse) to determine the position of the coil within the magnet and analysed using the same methods described in Section 7.3.3 using the k value derived from the calibration experiment. The orientation of the sensor was measured over a 1 second window (5000 data points), after the phantom had assumed its new pose, but before the gradient pre-pulses. The average voltage over the 5000 data points was used to compute the orientation using the methods described in Section 7.2.

The pose of the sensor as described by itself was compared to the pose of the sensor as described by the camera. The position of the sensor was calculated

from the position of the marker using the same methods described in Chapter 5. Unlike other external tracking devices, the sensor records its position in the magnet reference frame inherently due to its position being defined by the gradient coils. The position of the sensor relative to the marker ($\mathbf{f}$) could be defined at the start of the session when both the sensor and marker positions are known. Using $\mathbf{f}$ the position of the sensor can be determined using the camera. The sensor's position and orientation measured by the camera was then compared to its position and orientation measured by itself and the root-mean-square error was measured.

7.4 Results

Trapezoidal gradient pre-pulses were used to determine the position of the sensor within the coil. Each gradient pre-pulse was delivered in series (G_x then G_y then G_z) before the beginning of an imaging sequence. The focus of this experiment was determining the position and orientation of the sensor rather than obtaining an image and so the RF amplitude was turned down to minimum (1 °). The induced voltages for one set of pre-pulses can be seen in Figure 7.8.

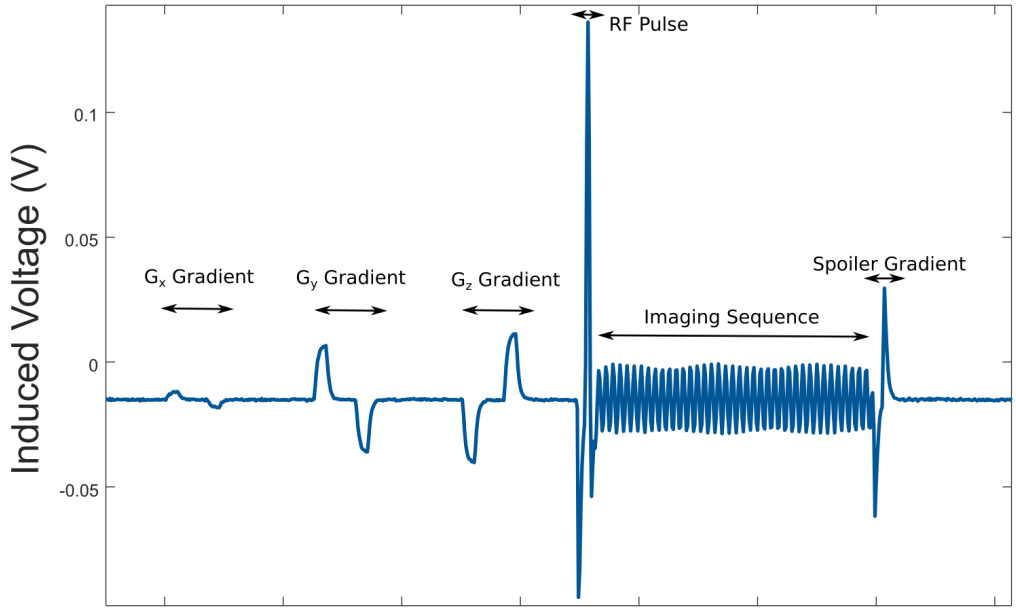


Figure 7.8: The voltages induced in the coils during an imaging sequence with gradient pre-pulses.

7.4.1 Calibrating the Sensor Coil

The sensor was calibrated by performing translations along the z axis of the scanner and measuring the induced voltages in the coil in response to a series of gradient pre-pulses. The voltages induced by the G_x and G_y gradients should be invariant to translations along z if the coil is oriented correctly. Figure 7.9A and 7.9B show that the induced voltages changed very little in response to the z translation. From Figure 7.9C we were able to determine the relationship between induced voltage and translations as k = 11.5 Vsm^{-1}.

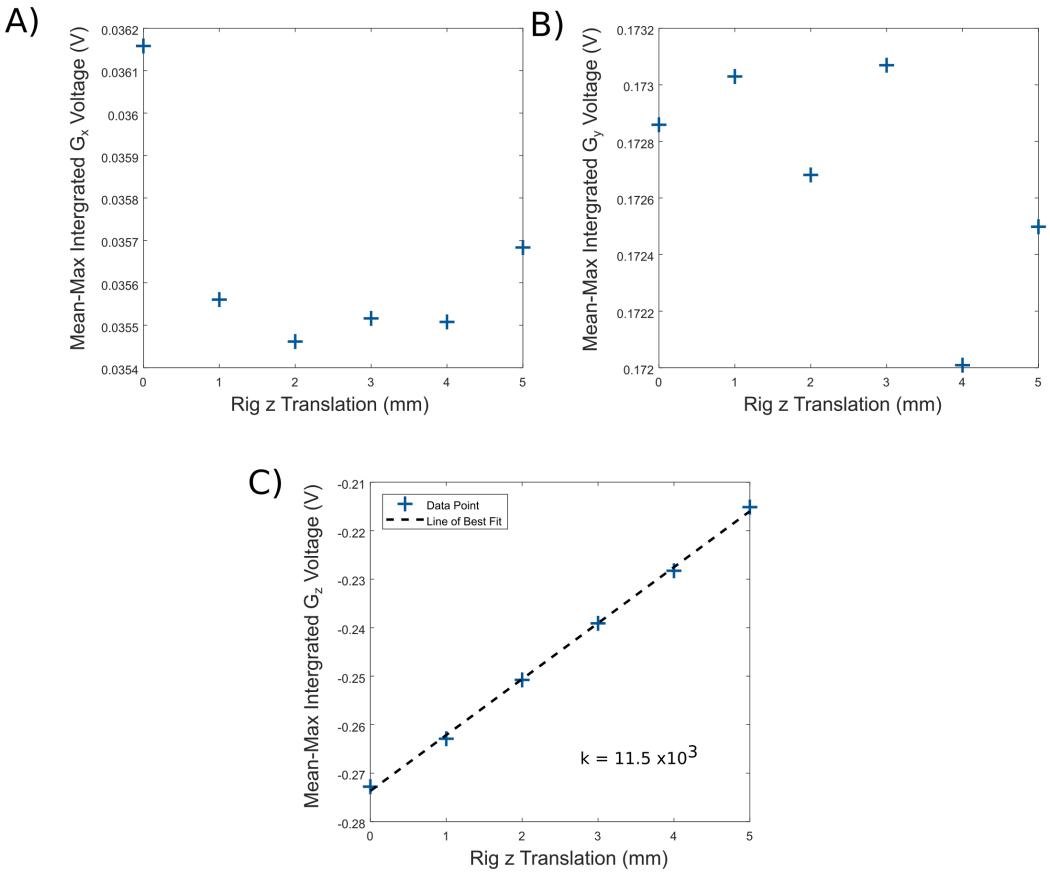


Figure 7.9: The mean-max integrated induced voltages from each gradient pre-pulse (G_x, G_y, G_z) in response to a translation in z. k was estimated from Figure C).

7.4.2 Sensor Vs Optical Tracking

The sensor pose was measured by both the camera and the sensor simultaneously. The pose, as measured by the camera, was used as the ground truth. The RMSE was calculated for each pose component ($x, y, z, \phi, \theta, \psi$). The sensor position was accurate to sub 0.5 mm and orientation accurate to 0.2° as seen in Figure 7.10.

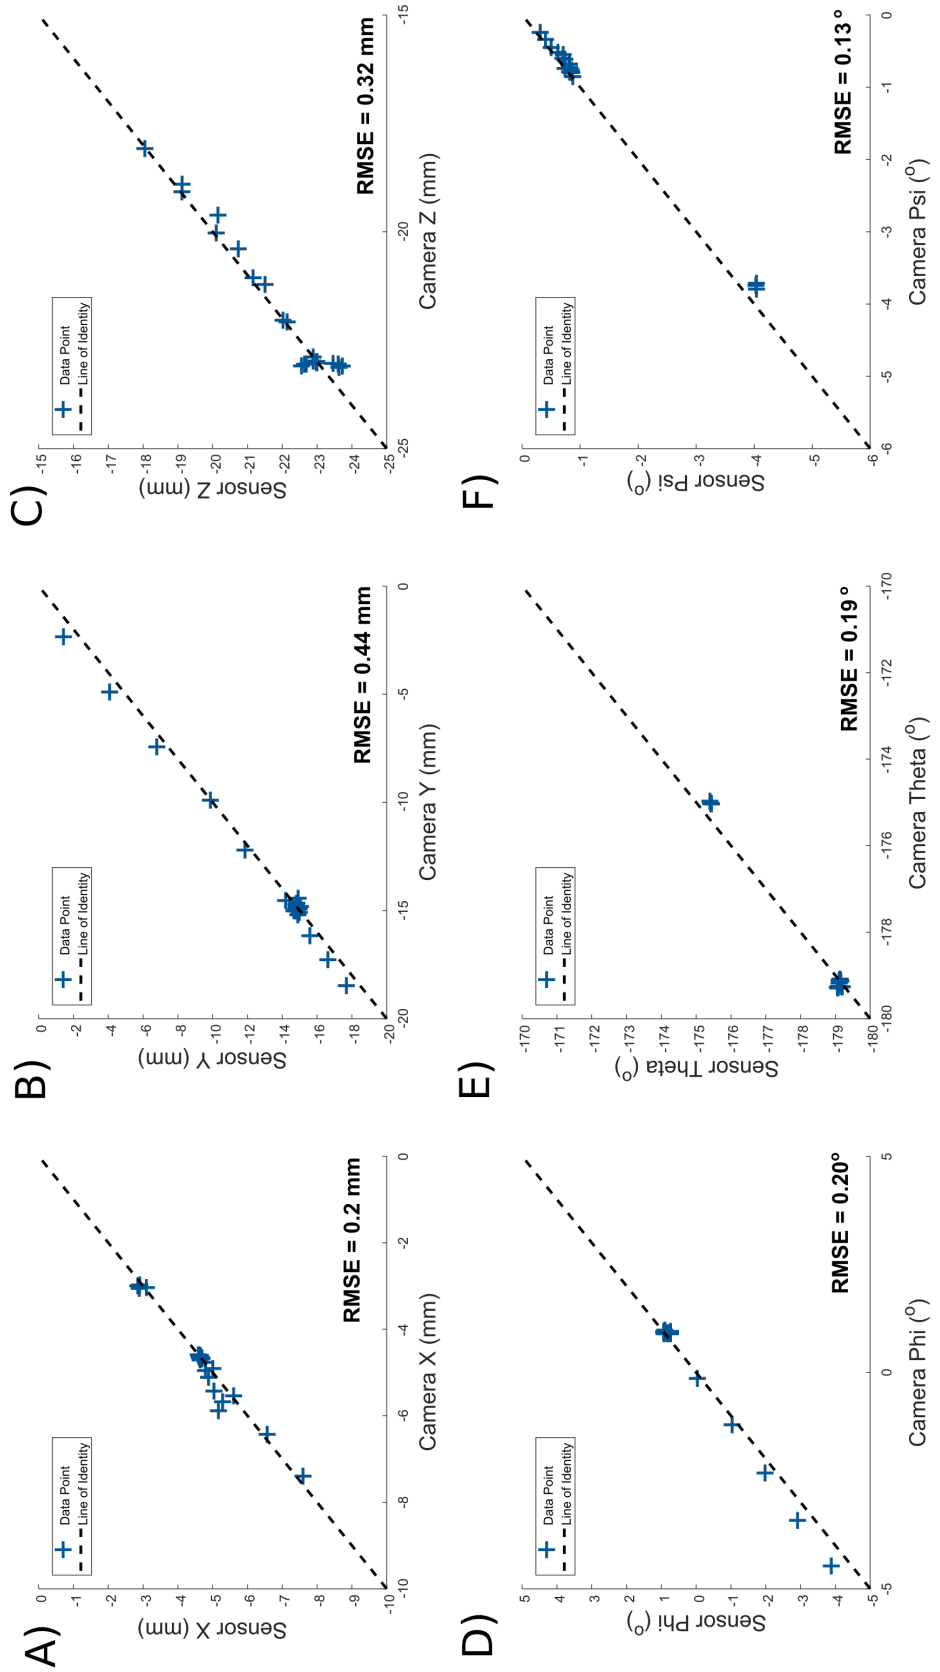


Figure 7.10: A comparison of the sensor position (A-C) and orientation (D-F) measured by the camera and the sensor. The RMSE of the sensor measurement is shown in the lower right of each subfigure.

7.5 Discussion

7.5.1 Calibrating the Sensor Coil

The measured value for k was 11.5 ± 0.5 Vsm^{-1} as seen in Figure 7.9C. This value however, is dependent on not only the sensor but also the gradient pulse characteristics. If the gradient slope rate were to change, for example, k would also change. However, because the parameters in k are linear, k can be adjusted for pulse parameters without having to be experimentally derived each time. In principle the device can be calibrated without experimentation assuming that the parameters in Equation 7.19 are known and the coil is constructed as specified.

The change in mean-max induced voltage due to a z translation was very small for the G_x and G_y pulses with voltages only varying by 0.001 Vs compared to 0.06 for the G_z pulse as seen in Figure 7.9. This shows that the rig was well aligned with the scanner and that only pure translations along z were made.

The accuracy of k could be improved by translating the phantom over a larger distance. However care should be taken to remain within the linear range of the gradients. In regions far away from the isocentre (approximately 200 mm) the gradients become non-linear. Without characterising the gradients at this range the position measured by the sensor would be incorrect. This represents a potential limitation for the sensor. While subject movements of $\pm$ 200 mm are unlikely to occur naturally, the initial placement of the sensor could be up to 200 mm away from isocentre if the sensor was mounted to a bite bar like the MPT marker. The advantage of the sensor is that it does not need to be visible to a camera and so there is greater flexibility as to where the sensor can be placed.

7.5.2 Sensor vs Optical Tracking

The sensor was able to detect its own orientation accurate to 0.2°. This is less accurate than would have been predicted by the theory (section 7.2.3), this could be caused by additional sources of noise that were not previously considered, such as scanner vibration or variations in sensor mounting between the AMR and

accelerometer (if the AMR and accelerometer are not mounted exactly parallel to each other). In addition Figure 7.10E shows that there may be some drift as the size of the rotations increase, which may mean the sensor is less accurate over larger rotations. Typically however a reasonably compliant subject is unlikely to move more than a few degrees during a scan and 0.2° is comparable with existing motion tracking methods.

The sensor was able to determine its own position with an accuracy of ± 0.44 mm. This is a higher accuracy than predicted based the error analysis. The max integrated voltage from each gradient pulse over the 15 repetitions varied very little, suggesting that the signal-to-noise ratio is sufficient to recover the position on the sensor. The reduced accuracy in position could be caused by small errors in the rotation as it was shown in the theory (7.2.3) that these cause the largest source of error when determining the position of the sensor.

Care should be taken to ensure that the sensor is not directly connected to the subject. The sensor temperature remained at a steady 20 ± 1 ℃ throughout the experiment as measured by the temperature sensor. However this temperature may not be representative of the temperature of the coil, or the temperature during an imaging sequence. The voltages induced from the gradients are relatively low, however those induced from the RF pulse are much higher and could cause larger increases in temperature in the coil. This may be an issue for long imaging sequences with many excitation pulses in a short time span [117]. Another reason to limit direct contact with the subject in the risk of surface electric fields from the subject being detected in the coil (similar to EEG methods), which could be mistaken for geometry information. The shape of the plastic spool helps to prevent direct contact to the subject and is a factor is future designs.

7.5.3 Further Work

To determine the position we used trapezoidal gradients pulses because they were readily available on the scanner. However these could be refined by using sinusoidal gradient wave forms instead. The waveforms would then induce

sinusoidal voltages in the coil, where the amplitude of the wave form would be dependent on the position of the sensor within the gradient field. This offers advantages over the current method because all 3 gradients could be performed simultaneously by encoding each gradient with a different frequency. The induced voltages could be Fourier transformed to the give the amplitude of each gradient wave form. This means that the position of the coil could be determined very quickly, using a very short gradient sequence, which would allow for more regular motion updates or increased bandwidth to reduce errors. *Brunner et al.* [118] demonstrated this technique to calibrate NMR field probes using gradient tones applied *during* image acquisition.

Under the current implementation the sensor is not suitable for prospective motion correction because the geometry data is only available after it has been taken offline.

Currently data is acquired using a MRI compatible EEG amplifier kit *Brain Products, Germany*, this was chosen because it is known to be MR safe and resistant to RF pulses. For an optimised setup the sensor data would be read in using a dedicated analogue to digital converter, this would allow data to be processed directly on the acquisition system. The input however must be able to measure the the low voltage signal from the sensor and also withstand the high voltage spikes from an RF pulse. The RF pulse generates significantly higher voltages in the coil even when the RF amplitude is set to minimum (1 ° flip angle) as seen in Figure 7.8. Alternatively data could be transmitted to a processing unit optically instead of electronically along a cable. To do this however the signal would have to be converted from an electronic signal to an optical signal on the sensor board requring significant modification.

The system could be designed so that it has enough dynamic range to ensure it can accurately measure the relatively low voltages of the sensor, as well as the high voltages that are induced in the presence of an RF pulse similar to how the EEG amplifier works. Because of the distinct “spike-like” nature of the RF pulse it may also be possible to filter out the RF pulse from the signal by circuit design. Alternatively the system could be gated so that no data is acquired during an

RF pulse, provided that the timing of the sequence is known, similar to how the receive coils function on the scanner.

Currently data is manually segmented to determine the location of the gradient pre-pulses within the sequence. This is not practical for prospective motion correction and a method would have to be used to extract the timing of the pulses. One method would be to synchronise sensor acquisition system and scanner using transistor-transistor-logic (TTL) triggering. Provided the timing of the sequence is known, gradient pulses can be automatically recorded at set times after the receiving the TTL trigger, however this would have to be altered between scan sequences to ensure the timing between trigger and pre-pulse was correct. Alternatively if the tone gradients (as described above) were used, the coil signal could be windowed and Fourier transformed. Provided the frequencies of the tones used are not within the imaging sequence (low frequency tones of 10's Hz) the gradients could be extracted from the signal directly.

Lastly, data would have to be read in, processed and exported in a timely fashion on a dedicated system. The output format could contain a similar structure to that seen in Table 5.1 to ensure compatibility. Processing time for a single pose (averaging over 15 gradient repetitions) is approximately 0.04 s. This processing time is more than sufficient for motion correction. The largest source of latency would be frequency of gradient pre-pulses in the imaging sequence. However a single set of gradient pre-pulses could be applied in 60 ms making it suitable for slice-by-slice or shot-by-shot correction, although the optimum number of sets to perform would have to be determined to provide sufficent signal-to-noise ratio. Alternatively a tone burst of sinusoidal gradients could be used which could reduce the time required to $\approx$ 5ms. *Aranovitch et al.* [119] recently demonstrated prospective motion correction by localising NMR probes affixed to the subject head using native sequence elements. This method could easily be adapted for use with an RF coil and would allow geometry updates to be determined without additional gradient pre-pulses to the sequence.

7.6 Conclusion

In this chapter the design for an electronic sensor capable of determining its own position and orientation within an MRI scanner is presented. This sensor has been prototyped and validated against existing technology. The sensor was able to determine its own position accurate to better than 0.5 mm and 0.2°. This accuracy is similar to accuracies seen by optical tracking cameras. The sensor is not subject to the limitations of other techniques: it does not require line of sight like optical tracking, it is considerably cheaper than NMR probes and can be used on multiple scanners with very little calibration required.

The modifications and developments that would allow for such a device to be used for prospective motion correction are outlined. They contain a description of the required hardware to ensure quick read and processing time of geometry data as well as proposed sequence changes to minimise impact to scan times.

Chapter 8

Summary

8.1 Concluding Remarks

The focus of this thesis has been on the use of motion tracking technology in MRI. Despite the wide range of motion tracking methods available, no method thus far has been shown to be ultimately superior, although some methods may be more suitable than others depending on the user requirements. External tracking devices offer an advantage over MR tracking methods in that they are minimally disruptive to the imaging sequence and often more accurate.

Despite the apparent accuracy of optical tracking cameras, errors quoted by vendors may be misleading. While the error of the camera measurement may be small, errors in the calibration transform may cause large errors in the apparent subject position. The non-iterative calibration approach proposed by *Daniilidis et al.* [73] was shown to be extremely sensitive to errors in the orientation of the calibration poses as shown in Chapter 4. Often when the error was high the calibration algorithm failed to determine a reasonable calibration. However the bootstrapping method shown in Chapter 4 presents a fast, easy method of identifying unstable calibrations. We now commonly use this method to verify the calibration transform before any motion correction experiment.

The Geocam patch demonstrated in Chapter 5 was shown to be comparable in realignment capability to the existing prospective image registration method available on Philips 7 T scanners for inter-volume motion correction. Optically tracked motion correction is not new to the MRI community, however, the Geocam patch offers new functionality to the Philips 7 T scanner by correcting for movement within a volume acquisition, which reduces blurring and ghosting from movement as well limiting contrast changes caused by spin history effects. Although its use clinically would depend on the subject's disposition towards wearing a bite bar (see Section 8.2.2 for alternatives).

Recently, *Kineticor* have released a new model of the Moiré phase tracking camera that addresses this problem. In the new model they use a multi-camera system to increase the field of view, and use disposable nose clips to mount the markers to the subject. This is a considerable step forward for motion correc-

tion, but is not suitable for the 7 T environment where the coil obscures the face, preventing direct line of sight between the camera and marker.

Despite the current implementation of the Moiré phase tracking camera at 7 T, the camera may be better suited for lower field strengths. At 3 T the transmit coil is fitted into the bore of the magnet, and the receive coils for neuroimaging are typically more open allowing for increased marker visibility which may eliminate the need for bite bars. The shorter bore of the 3 T scanner would also make camera calibration easier as the experimenter would not have to remain in the bore during calibration and could instead move the phantom by leaning into the bore.

The additional artefacts in EPI arising from orientation dependent changes in B_0 can be corrected for using the methods described in Chapter 6. This method requires prospective motion correction for the EPI images so that the phase encoding direction is maintained relative to the head. We demonstrated this using the Geocam patch, however any method of motion correction can be used provided that it is performed prospectively, and that the change in head orientation can be measured quantitatively so that it can be used to generate the predicted B_0 maps.

The sensor shown in Chapter 7 represents a potential replacement for the camera. The main advantages of this sensor over the optical tracking camera are twofold: Firstly, it is significantly cheaper than the optical tracking camera (with sensor construction costs of the order of £20, although this could be significantly more for a commerical product) and secondly it does not require line of sight to an external device. This device is an early prototype and has potential to be further developed. Similar devices such as that demonstrated by *Niekerk et al.* [120] which also uses known field directions and RF coils to determine position, are growing in popularity.

8.2 Future Work

The work in this thesis demonstrates the practical applications and limitations of using external tracking devices for motion correction, some of which could be addressed in future work.

8.2.1 Geocam Patch Improvements

In Chapter 5 we discussed potential improvements upon the existing Geocam framework. These corrections relate to the frequency of geometry updates and the methods by which the geometry is updated. These improvements could further increase image quality in the presence of subject motion by correcting for subject motion more frequently.

8.2.2 Marker Fixation

The use of a bite bar for marker fixation is a limiting factor for clinical adaptation of optically tracked motion correction. The bite bars can cause excess salivation and can be uncomfortable for the subject depending on mouth shape. However, they offer the most rigid connection to the skull through fixation with the upper teeth. The strong fixation means that the bar onto which the marker is placed can be sufficiently long to be seen outside of the coil.

The nose clip method used by the new version of the Kineticor camera is far more suitable for the clinical environment, but at 7 T the marker would still be obscured by the coil. Given the fixation between the nose and the clip is relatively weak compared to the bite bar, they can not be easily adapted to include a bar to extend out of the coil because the weight of the bar is likely to reduce the coupling between the marker and the nose. Furthermore, there is limited clearance between the face and the inside of the head coil for such modifications to be placed.

Alternatively a facemask could be used to provide a larger anchor point and hence potentially stronger fixation. The mask could have a bar fixed to it which

extends outside of the coil to allow the marker to be seen. The masks would have to be 3D printed to each subject's face to ensure that they fit tightly to a subject, with face models obtained from prior MR scans. However 3D printing facemasks for each subject would substantially increase the cost of fixation and is unlikely to couple to the skull as well as a bite bar.

Modifications could be made to the camera-scanner set up, by using a small mirror to increase visibility into the coil. However this method would increase the complexity of the motion tracking given that the camera would track a reflection of the marker rather than the marker itself. Given that the camera obtains information about the marker from properties on the marker face (Moiré phase patterns, starburst, marker ID) these would have to be considered by the tracking algorithm. Furthermore the mirror could potentially increase the error of the motion tracking if the mirror is moved after camera calibration. The current version of the camera has limited tracking range, and so the mirror would have to be placed close to the coil to reduce the increased distance between the camera and marker. A light source strong enough to illuminate the inside of the coil, but weak enough not to dazzle the subject would also have to be used.

8.2.3 Alternative Calibration Methods

The *Daniilidis* method remains the current method by which we determine the calibration transform. As discussed in Chapter 4 using the Geocam patch, we may be able to begin to use iterative methods to determine the calibration transform. These methods would allow the error in imaging volume to be estimated directly during the calibration experiment. With optimisation, iterative approaches may be faster than current techniques to determine the calibration.

Alternative methods for re-determining the calibration could also be used. The standard MPT markers for motion tracking are 15 mm wide and can be tracked by the camera from up to 400 mm, however we also have access to a large marker that is 70 mm wide and can be tracked from up to 700 mm away. If this was placed within the bore permanently, then the full calibration experiment would only

have to be performed once. Subsequent calibrations could be estimated from the relative position of this permanent marker in the camera reference frame. Since in the magnet reference frame the marker does not move, so any apparent change in the position of the marker is due to camera movement alone. The siting of the marker would have to be tested to ensure that it was stable over long periods of time. Any movement of the permanent marker could potentially cause large errors in the calibration. On the 7 T scanner at Nottingham the permanent marker was not able to remain in the bore as no stable place with line of sight between marker and camera was available. The clearance between the bed and the beam (the structure in the bore that the bed slides over) was not sufficient for the marker to be placed without being moved by the bed as it moved. There is also a caterpillar track connecting the bed to the rear of the scanner, which also prevents line of sight between the camera and any possible marker placement.

8.2.4 EPI Distortion Correction Applications

The geometric distortion correction methods described in Chapter 6 represent a new method of correcting fMRI data, demonstrated for single large movements. The next step would be to use these methods in a typical fMRI experiment. Typical fMRI experiments acquire significantly more volumes, which presents new problems for the speed in which this method can be performed. Currently each EPI image is corrected using its own individually generated B_0 map, the method described in Chapter 6 will allow for the generation of these B_0 maps. However, correcting each EPI image with each B_0 map individually would increase post-processing time significantly unless further optimisation and streamlining was performed.

8.2.5 Sensor Development

The electronic sensor demonstrated in Chapter 7 requires optimisation before it could be used for motion correction. The further work required to use this for prospective motion correction is outlined in Section 7.5.3. These improvements

involve alterations to the scanner by changing the gradient pre-pulses from trapezoidal gradients to gradient tones. Additional changes to the hardware of the sensor would also be required to allow data to be processed in a timely fashion.

This thesis represents significant advances in practical motion correction for MRI. Using the Geocam patch and distortion correction methods we were able to improve image quality drastically despite the presence of large subject movement. The proposed sensor in Chapter 7 represents a potential replacement for optical tracking technology which is not limited by line-of-sight and is significantly cheaper. The methods shown in this thesis would be well suited for MR imaging of subjects where motion is likely and unpreventable.

Bibliography

[1] F. Bloch. Nuclear induction. *Physical Review*, 70(7-8):460–474, oct 1946.

[2] Nicolaas Bloembergen, Edward Mills Purcell, and Robert V Pound. Relaxation effects in nuclear magnetic resonance absorption. *Physical review*, 73(7):679, 1948.

[3] Peter Mansfield and Peter K Grannell. NMR diffraction in solids. *Journal of Physics C: solid state physics*, 6(22):L422, 1973.

[4] Paul C Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 1973.

[5] RC Weast and MJ Astle. *Handbook of Chemistry and Physics*. Number ISBN 1-4822-6096-4. CRC Press, 96 edition, 2015.

[6] P. Zeeman. The effect of magnetisation on the nature of light emitted by a substance. *Nature*, 55(1424):347–347, feb 1897.

[7] P. J. Wright, O. E. Mougin, J. J. Totman, A. M. Peters, M. J. Brookes, R. Coxon, P. E. Morris, M. Clemence, S. T. Francis, R. W. Bowtell, and P. A. Gowland. Water proton T1 measurements in brain tissue at 7, 3, and 1.5T using IR-EPI, IR-TSE, and MPRAGE: results and optimization. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 21(1-2):121–130, feb 2008.

[8] Essa Yacoub, Amir Shmuel, Josef Pfeuffer, Pierre-Francois Van De Moortele, Gregor Adriany, Peter Andersen, J. Thomas Vaughan, Hellmut Merkle, Kamil Ugurbil, and Xiaoping Hu. Imaging brain function in humans at 7 tesla. *Magnetic Resonance in Medicine*, 45(4):588–594, 2001.

[9] Andrew M. Peters, Matthew J. Brookes, Frank G. Hoogenraad, Penny A. Gowland, Susan T. Francis, Peter G. Morris, and Richard Bowtell. T2 mea-

surements in human brain at 1.5, 3 and 7 T. *Magnetic Resonance Imaging*, 25(6):748–753, jul 2007.

[10] E. L. Hahn. Spin echoes. *Physical Review*, 80(4):580–594, nov 1950.

[11] A D Elster. Gradient-echo MR imaging: techniques and acronyms. *Radiology*, 186(1):1–8, jan 1993.

[12] Allen N Garroway, Peter K Grannell, and Peter Mansfield. Image formation in nmr by a selective irradiative process. *Journal of Physics C: Solid State Physics*, 7(24):L457, 1974.

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

[14] William A Edelstein, James MS Hutchison, Glyn Johnson, and Thomas Redpath. Spin warp nmr imaging and applications to human whole-body imaging. *Physics in medicine & biology*, 25(4):751, 1980.

[15] T. R. Brown, B. M. Kincaid, and K. Ugurbil. NMR chemical shift imaging in three dimensions. *Proceedings of the National Academy of Sciences*, 79(11):3523–3526, jun 1982.

[16] Stig Ljunggren. Imaging methods. *Journal of Magnetic Resonance*, 54:338–343, 1983.

[17] Donald B. Twieg. The k -trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods. *Medical Physics*, 10(5):610–621, sep 1983.

[18] Y. Zur, M. L. Wood, and L. J. Neuringer. Spoiling of transverse magnetization in steady-state sequences. *Magnetic Resonance in Medicine*, 21(2):251–263, oct 1991.

[19] Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, and Peter Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic resonance in medicine*, 42(5):952–962, 1999.

[20] 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). *Magnetic Resonance in Medicine*, 47(6):1202–1210, jun 2002.

[21] Ben Parkinson. Design considerations and experimental results for MRI systems using HTS magnets. *Superconductor Science and Technology*, 30(1):014009, 2016.

[22] A. Ishiyama and H. Hirooka. Magnetic shielding for MRI superconducting magnets. *IEEE Transactions on Magnetics*, 27(2):1692–1695, mar 1991.

[23] David G. Norris and James M.S. Hutchison. Concomitant magnetic field gradients and their effects on imaging at low magnetic field strengths. *Magnetic Resonance Imaging*, 8(1):33–37, jan 1990.

[24] Mat A. Bernstein, Xiaohong Joe Zhou, Jason A. Polzin, Kevin F. King, Alexander Ganin, Norbert J. Pelc, and Gary H. Glover. Concomitant gradient terms in phase contrast MR: Analysis and correction. *Magnetic Resonance in Medicine*, 39(2):300–308, feb 1998.

[25] C. Meier, M. Zwanger, T. Feiweier, and D. Porter. Concomitant field terms for asymmetric gradient coils: Consequences for diffusion, flow, and echo-planar imaging. *Magnetic Resonance in Medicine*, 60(1):128–134, jul 2008.

[26] Ulrich Katscher and Peter Börnert. Parallel RF transmission in MRI. *NMR in Biomedicine*, 19(3):393–400, 2006.

[27] Jason Filos, Jo Hajnal, and David Larkman. *Parallel Transmission MRI.* PhD thesis, Imperial College, London, 2008.

[28] Martijn de Greef, Ozlem Ipek, Alexander J. E. Raaijmakers, Johannes Crezee, and Cornelis A. T. van den Berg. Specific absorption rate intersubject variability in 7T parallel transmit MRI of the head. *Magnetic Resonance in Medicine*, 69(5):1476–1485, jul 2012.

[29] Maxim Zaitsev, Julian Maclaren, and Michael Herbst. Motion artifacts in MRI: a complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging*, 42(4):887–901, 2015.

[30] Max C Keuken, Pierre-Louis Bazin, Andreas Schäfer, Jane Neumann, Robert Turner, and Birte U Forstmann. Ultra-high 7T MRI of structural age-related changes of the subthalamic nucleus. *Journal of Neuroscience*, 33(11):4896–4900, 2013.

[31] Jalal B Andre, Brian W Bresnahan, Mahmud Mossa-Basha, Michael N Hoff, C Patrick Smith, Yoshimi Anzai, and Wendy A Cohen. Toward quantifying the prevalence, severity, and cost associated with patient motion during clinical MR examinations. *Journal of the American College of Radiology*, 12(7):689–695, 2015.

[32] F Godenschweger, U Kägebein, D Stucht, U Yarach, A Sciarra, R Yakupov, F Lüsebrink, P Schulze, and O Speck. Motion correction in MRI of the brain. *Physics in Medicine & Biology*, 61(5):R32, 2016.

[33] Douglas N Greve and Bruce Fischl. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage*, 48(1):63–72, 2009.

[34] P Lanzer, C Barta, E H Botvinick, H U Wiesendanger, G Modin, and C B Higgins. ECG-synchronized cardiac MR imaging: method and evaluation. *Radiology*, 155(3):681–686, jun 1985.

[35] Jon M. Chia, Stefan E. Fischer, Samuel A. Wickline, and Christine H. Lorenz. Performance of QRS detection for cardiac magnetic resonance imaging with a novel vectorcardiographic triggering method. *Journal of Magnetic Resonance Imaging*, 12(5):678–688, 2000.

[36] RL Ehman, MT McNamara, M Pallack, H Hricak, and CB Higgins. Magnetic resonance imaging with respiratory gating: techniques and advantages. *American Journal of Roentgenology*, 143(6):1175–1182, dec 1984.

[37] Claudio Santelli, Reza Nezafat, Beth Goddu, Warren J. Manning, Jouke Smink, Sebastian Kozerke, and Dana C. Peters. Respiratory bellows revisited for motion compensation: Preliminary experience for cardiovascular MR. *Magnetic Resonance in Medicine*, 65(4):1097–1102, nov 2010.

[38] D J Atkinson and R R Edelman. Cineangiography of the heart in a single breath hold with a segmented turboFLASH sequence. *Radiology*, 178(2):357–360, feb 1991.

[39] James C. Carr, Orlando Simonetti, Jeff Bundy, Debiao Li, Scott Pereles, and J. Paul Finn. Cine MR angiography of the heart with segmented true fast imaging with steady-state precession. *Radiology*, 219(3):828–834, jun 2001.

[40] Richard L Ehman and Joel P Felmlee. Adaptive technique for high-definition MR imaging of moving structures. *Radiology*, 173(1):255–263, 1989.

[41] Edward Brian Welch, Armando Manduca, Roger C. Grimm, Heidi A. Ward, and Clifford R. Jack Jr. Spherical navigator echoes for full 3D rigid body motion measurement in MRI. *Magnetic Resonance in Medicine*, 47(1):32–41, dec 2001.

[42] André J.W. van der Kouwe, Thomas Benner, and Anders M. Dale. Real-time rigid body motion correction and shimming using cloverleaf navigators. *Magnetic Resonance in Medicine*, 56(5):1019–1032, nov 2006.

[43] Joel P Felmlee and RL Ehman. Spatial presaturation: a method for suppressing flow artifacts and improving depiction of vascular anatomy in mr imaging. *Radiology*, 164(2):559–564, 1987.

[44] Andrew A. Badachhape, Ruth J. Okamoto, Ramona S. Durham, Brent D. Efron, Sam J. Nadell, Curtis L. Johnson, and Philip V. Bayly. The relationship of three-dimensional human skull motion to brain tissue deformation in magnetic resonance elastography studies. *Journal of Biomechanical Engineering*, 139(5):051002, mar 2017.

[45] P.V. Bayly, T.S. Cohen, E.P. Leister, D. Ajo, E.C. Leuthardt, and G.M. Genin. Deformation of the human brain induced by mild acceleration. *Journal of Neurotrauma*, 22(8):845–856, aug 2005.

[46] Jonathan Law and Richard Rennie. *A Dictionary of Physics*. OUP Oxford, 2015.

[47] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[48] S. C. Cramer, G. Nelles, R. R. Benson, J. D. Kaplan, R. A. Parker, K. K. Kwong, D. N. Kennedy, S. P. Finklestein, and B. R. Rosen. A functional MRI study of subjects recovered from hemiparetic stroke. *Stroke*, 28(12):2518–2527, dec 1997.

[49] Ronald S Litman, Komal Soin, and Abdul Salam. Chloral hydrate sedation in term and preterm infants: an analysis of efficacy and complications. *Anesthesia & Analgesia*, 110(3):739–746, 2010.

[50] Simon C Kao, Susan D Adamson, Leann H Tatman, and Kevin S Berbaum. A survey of post-discharge side effects of conscious sedation using chloral hydrate in pediatric CT and MR imaging. *Pediatric radiology*, 29(4):287–290, 1999.

[51] KJS Anand and RW Hall. Pharmacological therapy for analgesia and sedation in the newborn. *Archives of Disease in Childhood-Fetal and Neonatal Edition*, 91(6):F448–F453, 2006.

[52] Arno Klein, Jesper Andersson, Babak A Ardekani, John Ashburner, Brian Avants, Ming-Chang Chiang, Gary E Christensen, D Louis Collins, James Gee, Pierre Hellier, et al. Evaluation of 14 nonlinear deformation algorithms applied to human brain mri registration. *Neuroimage*, 46(3):786–802, 2009.

[53] JL Duerk and PM Pattany. Analysis of imaging axes significance in motion artifact suppression technique (MAST™): MRI of turbulent flow and motion. *Magnetic resonance imaging*, 7(3):251–263, 1989.

[54] Mark Jenkinson and Stephen Smith. A global optimisation method for robust affine registration of brain images. *Medical image analysis*, 5(2):143–156, 2001.

[55] Peter Jezzard. Correction of geometric distortion in fMRI data. *Neuroimage*, 62(2):648–651, 2012.

[56] Michael S Hansen, Thomas S Sørensen, Andrew E Arai, and Peter Kellman. Retrospective reconstruction of high temporal resolution cine images from real-time MRI using iterative motion correction. *Magnetic resonance in medicine*, 68(3):741–750, 2012.

[57] M Dylan Tisdall, Aaron T Hess, Martin Reuter, Ernesta M Meintjes, Bruce Fischl, and André JW van der Kouwe. Volumetric navigators for prospective motion correction and selective reacquisition in neuroanatomical mri. *Magnetic resonance in medicine*, 68(2):389–399, 2012.

[58] Nathan White, Cooper Roddey, Ajit Shankaranarayanan, Eric Han, Dan Rettmann, Juan Santos, Josh Kuperman, and Anders Dale. PROMO: real-time prospective motion correction in mri using image-based tracking. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 63(1):91–105, 2010.

[59] Daniel Gallichan, José P. Marques, and Rolf Gruetter. Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magnetic Resonance in Medicine*, 75(3):1030–1039, apr 2015.

[60] James G Pipe. Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 42(5):963–969, 1999.

[61] Erez Nevo. Method and apparatus to estimate location and orientation of objects during magnetic resonance imaging, February 4 2003. US Patent 6,516,213.

[62] Julian Maclaren, Michael Herbst, Oliver Speck, and Maxim Zaitsev. Prospective motion correction in brain imaging: a review. *Magnetic resonance in medicine*, 69(3):621–636, 2013.

[63] Lei Qin, Peter van Gelderen, John Andrew Derbyshire, Fenghua Jin, Jongho Lee, Jacco A de Zwart, Yang Tao, and Jeff H Duyn. Prospective head-movement correction for high-resolution MRI using an in-bore optical tracking system. *Magnetic resonance in medicine*, 62(4):924–934, 2009.

[64] Rasool Khadem, Clement C Yeh, Mohammad Sadeghi-Tehrani, Michael R Bax, Jeremy A Johnson, Jacqueline Nerney Welch, Eric P Wilkinson, and Ramin Shahidi. Comparative tracking error analysis of five different optical tracking systems. *Computer Aided Surgery*, 5(2):98–107, 2000.

[65] Joshua T Weinhandl, Brian SR Armstrong, Todd P Kusik, Robb T Barrows, and Kristian M O'Connor. Validation of a single camera three-dimensional motion tracking system. *Journal of biomechanics*, 43(7):1437–1440, 2010.

[66] Brian Armstrong, Thomas Verron, Lee Heppe, Jim Reynolds, and Karl Schmidt. RGR-3D: simple, cheap detection of 6-dof pose for teleoperation, and robot programming and calibration. In *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, volume 3, pages 2938–2943. IEEE, 2002.

[67] Roger Y Tsai and Reimar K Lenz. A new technique for fully autonomous and efficient 3D robotics hand/eye calibration. *IEEE Transactions on robotics and automation*, 5(3):345–358, 1989.

[68] R.L. Hirsh, G.N. DeSouza, and A.C. Kak. An iterative approach to the hand-eye and base-world calibration problem. In *International Conference on Robotics and Automation (Cat. No.01CH37164)*. IEEE, 2001.

[69] Hanqi Zuang and Yiu Cheung Shiu. A noise-tolerant algorithm for robotic hand-eye calibration with or without sensor orientation measurement. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(4):1168–1175, 1993.

[70] Y. Motai and A. Kosaka. Hand–eye calibration applied to viewpoint selection for robotic vision. *IEEE Transactions on Industrial Electronics*, 55(10):3731–3741, oct 2008.

[71] F. Dornaika and R. Horaud. Simultaneous robot-world and hand-eye calibration. *IEEE Transactions on Robotics and Automation*, 14(4):617–622, 1998.

[72] Nicolas Andreff, Radu Horaud, and Bernard Espiau. Robot hand-eye calibration using structure-from-motion. *The International Journal of Robotics Research*, 20(3):228–248, mar 2001.

[73] Konstantinos Daniilidis. Hand-eye calibration using dual quaternions. *The International Journal of Robotics Research*, 18(3):286–298, 1999.

[74] James Smith, Olivier Mougin, Maxim Zaitsev, Benjamin Knowles, Richard Bowtell, Paul Glover, and Penny Gowland. From hand-eye calibration of optical tracking cameras to motion correction. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, number 1296, 2017.

[75] Jack C. K. Chou and M. Kamel. Finding the position and orientation of a sensor on a robot manipulator using quaternions. *The International Journal of Robotics Research*, 10(3):240–254, jun 1991.

[76] Ben Kenwright. A beginners guide to dual-quaternions. *WSCG'2012*, 2012.

[77] William Rowan Hamilton. *Elements of quaternions*, volume 1. Longmans, Green, and Company, 1899.

[78] M.A. Clifford, editor. *Preliminary Sketch of bi-quaternions*, volume 1. Proceedings of the London Mathematical Society, November 1873.

[79] Michael W. Walker, Lejun Shao, and Richard A. Volz. Estimating 3-d location parameters using dual number quaternions. *CVGIP: Image Understanding*, 54(3):358–367, nov 1991.

[80] Homer H Chen. A screw motion approach to uniqueness analysis of head-eye geometry. In *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR'91., IEEE Computer Society Conference on*, pages 145–151. IEEE, 1991.

[81] Benjamin Zahneisen, Chris Lovell-Smith, Michael Herbst, Maxim Zaitsev, Oliver Speck, Brian Armstrong, and Thomas Ernst. Fast noniterative calibration of an external motion tracking device. *Magnetic resonance in medicine*, 71(4):1489–1500, 2014.

[82] P.M. Glover, I. Cavin, W. Qian, R. Bowtell, and P.A. Gowland. Magnetic-field-induced vertigo: A theoretical and experimental investigation. *Bioelectromagnetics*, 28(5):349–361, jul 2007.

[83] Donald W. Chakeres and Frank de Vocht. Static magnetic field effects on human subjects related to magnetic resonance imaging systems. *Progress in Biophysics and Molecular Biology*, 87(2-3):255–265, feb 2005.

[84] Zhiqiang Zhang, Lin Zhang, and Guang-Zhong Yang. A computationally efficient method for hand–eye calibration. *International Journal of Computer Assisted Radiology and Surgery*, 12(10):1775–1787, jul 2017.

[85] Klaus Strobl and Gerd Hirzinger. Optimal hand-eye calibration. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, oct 2006.

[86] Stefan Thesen, Oliver Heid, Edgar Mueller, and Lothar R. Schad. Prospective acquisition correction for head motion with image-based tracking for real-time fMRI. *Magnetic Resonance in Medicine*, 44(3):457–465, 2000.

[87] Maxim Zaitsev, Christian Dold, Georgios Sakas, Jürgen Hennig, and Oliver Speck. Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. *Neuroimage*, 31(3):1038–1050, 2006.

[88] Aditya Singh, Benjamin Zahneisen, Brian Keating, Michael Herbst, Linda Chang, Maxim Zaitsev, and Thomas Ernst. Optical tracking with two markers for robust prospective motion correction for brain imaging. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 28(6):523–534, 2015.

[89] Daniel Stucht, K Appu Danishad, Peter Schulze, Frank Godenschweger, Maxim Zaitsev, and Oliver Speck. Highest resolution in vivo human brain MRI using prospective motion correction. *PloS one*, 10(7):e0133921, 2015.

[90] Kay Nehrke and Peter Börnert. Prospective correction of affine motion for arbitrary MR sequences on a clinical scanner. *Magnetic Resonance in Medicine*, 54(5):1130–1138, 2005.

[91] Calvin R Maurer Jr, Georges B Aboutanos, Benoit M Dawant, Srikanth Gadamsetty, Richard A Margolin, Robert J Maciunas, and J Michael Fitzpatrick. Effect of geometrical distortion correction in MR on image registration accuracy. *Journal of computer assisted tomography*, 20(4):666–679, 1996.

[92] Chen Hua-Hsuan, Boykin Rex D., Clarke Geoffrey D., Gao Jia-Hong T., and Roby John W. Routine testing of magnetic field homogeneity on clinical mri systems. *Medical Physics*, 33(11):4299–4306, 2006.

[93] Weston A Anderson. Electrical current shims for correcting magnetic fields. *Review of Scientific Instruments*, 32(3):241–250, 1961.

[94] Stephen M. Smith, Mark Jenkinson, Mark W. Woolrich, Christian F. Beckmann, Timothy E.J. Behrens, Heidi Johansen-Berg, Peter R. Bannister, Marilena De Luca, Ivana Drobnjak, David E. Flitney, Rami K. Niazy, James Saunders, John Vickers, Yongyue Zhang, Nicola De Stefano, J. Michael Brady, and Paul M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23:S208–S219, jan 2004.

[95] 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.

[96] Aleksandra Sulikowska. *Motion correction in high-field MRI*. PhD thesis, University of Nottingham, 2016.

[97] James Smith, Olivier Mougin, Kingkarn Aphiwatthanasumet, Matthew Clemence, Andrew Peters, Richard Bowtell, and Paul Gloverand Penny Gowland. Combined prospective motion correction and retrospective B0 correction for EPI using optical tracking and predictive B0 maps. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, number 4089, 2018.

[98] Jesper L.R. Andersson, Mark S. Graham, Ivana Drobnjak, Hui Zhang, and Jon Campbell. Susceptibility-induced distortion that varies due to motion: Correction in diffusion MR without acquiring additional data. *NeuroImage*, 171:277–295, may 2018.

[99] Raimon H.R. Pruim, Maarten Mennes, Daan van Rooij, Alberto Llera, Jan K. Buitelaar, and Christian F. Beckmann. ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *NeuroImage*, 112:267–277, may 2015.

[100] Jesper L.R. Andersson and Stamatios N. Sotiropoulos. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage*, 125:1063–1078, jan 2016.

[101] Jesper L.R. Andersson and Stamatios N. Sotiropoulos. Non-parametric representation and prediction of single- and multi-shell diffusion-weighted MRI data using gaussian processes. *NeuroImage*, 122:166–176, nov 2015.

[102] Maximilian Haeberlin, Lars Kasper, Christoph Barmet, David O. Brunner, Benjamin E. Dietrich, Simon Gross, Bertram J. Wilm, Sebastian Kozerke, and Klaas P. Pruessmann. Real-time motion correction using gradient tones and head-mounted NMR field probes. *Magnetic Resonance in Medicine*, 74(3):647–660, sep 2014.

[103] Nikolaus Weiskopf, Martina F. Callaghan, Oliver Josephs, Antoine Lutti, and Siawoosh Mohammadi. Estimating the apparent transverse relaxation time (r2) from images with different contrasts (ESTATICS) reduces motion artifacts. *Frontiers in Neuroscience*, 8, sep 2014.

[104] Julian Maclaren, Brian SR Armstrong, Robert T Barrows, KA Danishad, Thomas Ernst, Colin L Foster, Kazim Gumus, Michael Herbst, Ilja Y Kadashevich, Todd P Kusik, et al. Measurement and correction of microscopic head motion during magnetic resonance imaging of the brain. *PLOS one*, 7(11):e48088, 2012.

[105] J. Andrew Derbyshire, Graham A. Wright, R. Mark Henkelman, and R. Scott Hinks. Dynamic scan-plane tracking using MR position monitoring. *Journal of Magnetic Resonance Imaging*, 8(4):924–932, jul 1998.

[106] JL Ackerman, MC Offutt, RB Buxton, TJ Brady, et al. Rapid 3D tracking of small rf coils. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, page 1131, 1986.

[107] Melvyn B. Ooi, Sascha Krueger, William J. Thomas, Srirama V. Swaminathan, and Truman R. Brown. Prospective real-time correction for arbitrary head motion using active markers. *Magnetic Resonance in Medicine*, 62(4):943–954, oct 2009.

[108] Melvyn B. Ooi, Sascha Krueger, Jordan Muraskin, William J. Thomas, and Truman R. Brown. Echo-planar imaging with prospective slice-by-slice motion correction using active markers. *Magnetic Resonance in Medicine*, 66(1):73–81, feb 2011.

[109] E. H. Bhuiyan. *Monitoring movement in MRI by measuring changes in the EMF induced in head-mounted coils*. PhD thesis, University of Nottingham, 2017.

[110] Adam van Niekerk, Andre van der Kouwe, and Ernesta Meintjes. A method for measuring orientation within a magnetic resonance imaging scanner using gravity and the static magnetic field (VectOrient). *IEEE Transactions on Medical Imaging*, 36(5):1129–1139, 2017.

[111] James Smith, Glyn Spencer, Richard Bowtell, Penny Gowland, and Paul Glover. A novel position and orientation sensor for MRI. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, number 3930, 2017.

[112] STMicroelectronics. LIS344ALH, MEMS inertial sensor, high performance 3-axis ±2/±6g ultracompact linear accelerometer. Data Sheet, April 2008.

[113] Alfred Brian Pippard. *Magnetoresistance in metals*, volume 2. Cambridge University Press, 1989.

[114] Honeywell International Inc. APS00B high resolution magnetic angular position sensor ic. Data Sheet, August 2011.

[115] Texas Instruments. LP2985 150-mA low-noise low-dropout regulator with shutdown. Data Sheet, January 2015.

[116] Microschip Technology Inc. TC1046 precision temperature-to-voltage converter. Data Sheet, December 2012.

[117] Jonathan S. Friedstat, Molly E. Moore, Jeremy Goverman, and Shawn P. Fagan. An unusual burn during routine magnetic resonance imaging. *Journal of Burn Care & Research*, 34(2):e110–e111, 2013.

[118] DO Brunner, C Barmet, M Haeberlin, BJ Wilm, and KP Pruessmann. Autocalibration of field monitoring arrays by reference tones. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 19, page 1841, 2011.

[119] Alexander Aranovitch, Maximilian Haeberlin, Simon Gross, Benjamin E Dietrich, Bertram J Wilm, David O Brunner, Thomas Schmid, Roger Luechinger, and Klaas P Pruessmann. Prospective motion correction with nmr markers using only native sequence elements. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, 2018.

[120] Adam M.J. van Niekerk, Andre J. W. van der Kouwe, and Ernesta M. Meintjes1. Hybrid motion sensing with a wireless device, self-synchronized to the imaging pulse sequence. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, number 1003, June 2018.

Appendix A

Sensor Circuit Schematics

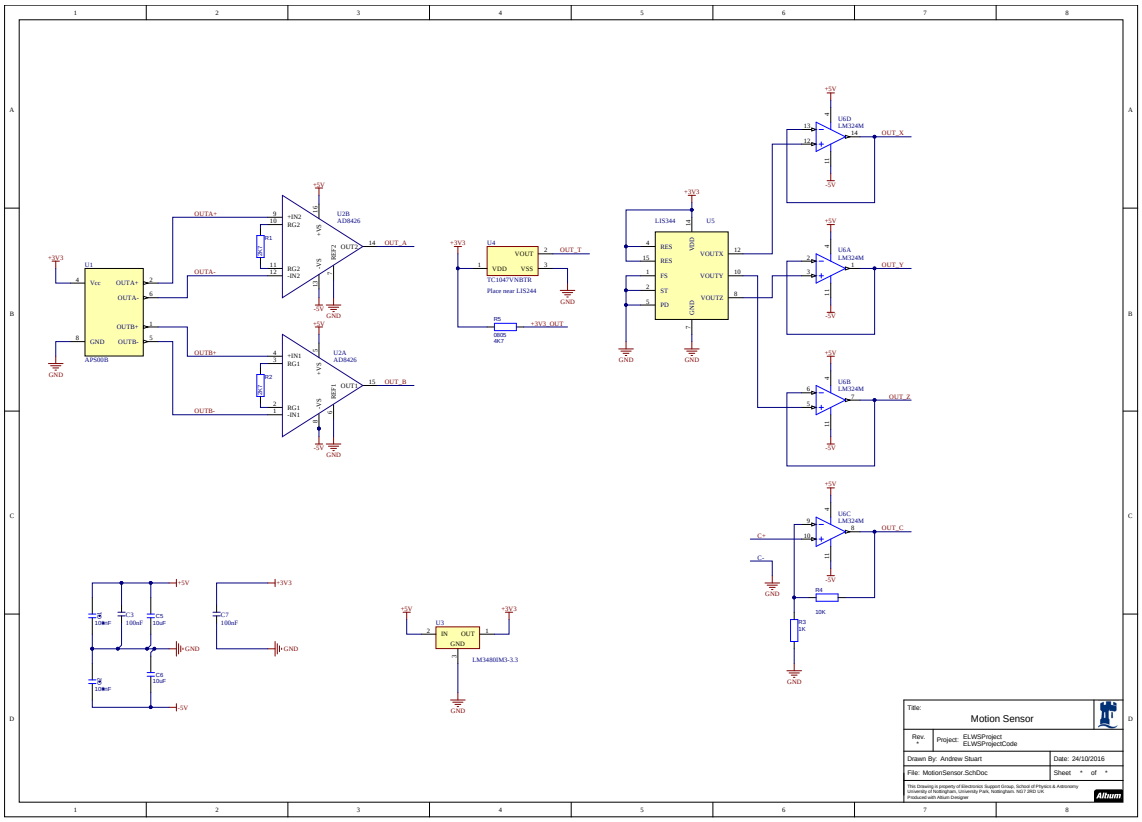


Figure A.1: The circuit schematics for the electronic motion sensor. Shown are the circuits containing the AMR, temperature sensor, accelerometer and voltage regulator.