Imperial College of Science, Technology and Medicine
Department of Computing

Motion robust acquisition and reconstruction of quantitative T2* maps in the developing brain

Serhiy (Serge) Vasylechko

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Computing at Imperial College London,
September 2019

Abstract

The goal of the research presented in this thesis was to develop methods for quantitative T2* mapping of the developing brain. Brain maturation in the early period of life involves complex structural and physiological changes caused by synaptogenesis, myelination and growth of cells. Molecular structures and biological processes give rise to varying levels of T2* relaxation time, which is an inherent contrast mechanism in magnetic resonance imaging. The knowledge of T2* relaxation times in the brain can thus help with evaluation of pathology by establishing its normative values in the key areas of the brain. T2* relaxation values are a valuable biomarker for myelin microstructure and iron concentration, as well as an important guide towards achievement of optimal fMRI contrast. However, fetal MR imaging is a significant step up from neonatal or adult MR imaging due to the complexity of the acquisition and reconstruction techniques that are required to provide high quality artifact-free images in the presence of maternal respiration and unpredictable fetal motion. The first contribution of this thesis, described in Chapter 4, presents a novel acquisition method for measurement of fetal brain T2* values. At the time of publication, this was the first study of fetal brain T2* values. Single shot multi-echo gradient echo EPI was proposed as a rapid method for measuring fetal T2* values by effectively freezing intra-slice motion. The study concluded that fetal T2* values are higher than those previously reported for pre-term neonates and decline with a consistent trend across gestational age. The data also suggested that longer than usual echo times or direct T2* measurement should be considered when performing fetal fMRI in order to reach optimal BOLD sensitivity. For the second contribution, described in Chapter 5, measurements were extended to a higher field strength of 3T and reported, for the first time, both for fetal and neonatal subjects at this field strength. The technical contribution of this work is a fully automatic segmentation framework that propagates brain tissue labels onto the acquired T2* maps without the need for manual intervention. The third contribution, described in Chapter 6, proposed a new method for performing 3D fetal brain reconstruction where the available data is sparse and is therefore limited in the use of current state of the art techniques for 3D brain reconstruction in the presence of motion. To enable a high resolution reconstruction, a generative adversarial network was trained to perform image to image translation between T2 weighted and T2* weighted data. Translated images could then be served as a prior for slice alignment and super resolution reconstruction of 3D brain image.

Acknowledgements

I would like to express a sincere thank you to my supervisors, Daniel Rueckert and Jo Hajnal, for the incredible support and encouragement that they have provided me with during this PhD journey. It was a genuine pleasure to work with both of them. I could never be grateful enough for the opportunity that they have provided me with. I hope that I will be able to keep crossing paths with both of them for a long time to come and not lose the connection.

I would also like to thank all of the members at the Center of the Developing Brain, with whom I had worked with. In particular - Joanna Allsop, Matthew Fox, Christina Malamateniou, Rita Nunes, Maryam Abael, Tomoki Arichi and Josef Habib. Their help had been invaluable. Moreover, those were some of the happiest days of my research. The working atmosphere that they had cultivated, together with the rest of the CDB staff, had been the most enjoyable.

Finally, there are some people who had made a big difference to my research outside of academia. Omer El Fakir - a friend, a gym buddy, a business partner and a true confidante for any trouble - thank you for your advice throughout. Iryna Lurye - a very special person whose love and care had helped me to keep going. And, most importantly, my mother - who had always been there for me since the beginning of this degree. It is to her that I'm indebted the most.

Copyright Declaration

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives license. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the license terms of this work.

Declaration of Originality

I hereby declare that the work described in this thesis is my own, except where specifically acknowledged.

Contents

Abstract **i**

Acknowledgements **iii**

1 Introduction **1**

1.1 Preface 1

1.2 Precession 2

1.3 Relaxation 5

1.4 Molecular basis of relaxation 6

1.5 Relaxometry 8

1.6 Biological basis of T2* relaxation 9

1.7 Developing brain 10

1.8 T2* in the newborn brain 11

1.9 Fetal Imaging 11

2 Acquisition **13**

2.1 Spatial Encoding 14

2.2 K-space 15

2.3 Gradient Echo 16

2.4 Sequences for T2* mapping 17

2.4.1 Multi-echo GRE 17

2.4.2 GRE EPI at varying echo times 19

2.4.3 Multi-echo GRE EPI 21

2.4.4 Comparison of approaches 22

2.5 Controlling T2* signal loss 24

2.5.1 Shimming 25

2.5.2 Voxel size 25

2.5.3 Receiver Bandwidth 27

2.5.4 Slew Rate 27

2.5.5 Susceptibility induced signal loss 28

2.5.6 Phase Encode Direction 32

2.6 Motion Compensated Imaging 32

2.6.1 Parallel Imaging 32

2.6.2 Partial Fourier 35

2.6.3 Simultaneous Multi-Slice 35

2.6.4 Navigators 36

2.6.5 Non-cartesian sequences 37

2.7 Artifacts 37

2.7.1 Chemical shift 37

2.7.2 Distortion 41

2.7.3 Nyquist ghosts 42

2.7.4 Aliasing 43

2.7.5 Bias Field 44

2.7.6 Gibbs' artifact 45

2.8 Patient comfort 46

3 Reconstruction 48

3.1 Optimal sampling rate and range 48

3.2 T2* estimation algorithms 50

3.3 Truncation, baseline and offset models 51

3.4 Spatial smoothing models 52

3.5 Region based relaxometry 53

3.6 Multi-exponential models 53

3.7 Slice to volume reconstruction 54

3.7.1 Slice alignment 55

3.7.2 Volume reconstruction 56

3.7.3 Evolution of SVR methods 57

4 Fast motion robust fetal T2* relaxometry at 1.5T 59

4.1 Introduction 59

4.2 Method 60
4.2.1 Acquisition 60
4.2.2 Subjects 63
4.2.3 Regions of Interest 64
4.2.4 T2* Estimation 66
4.2.5 Protocol Verification 67
4.3 Results 68
4.3.1 Protocol Verification 68
4.3.2 T2* in Fetal Subjects 69
4.4 Discussions and Conclusions 75

5 Fetal and neonatal whole brain T2* mapping at 3T 78

5.1 Introduction 78
5.1.1 Prior literature 78
5.1.2 Contributions 81
5.2 Method 82
5.2.1 Acquisition 82
5.2.2 Subjects 84
5.2.3 T2* estimation 85
5.2.4 Automatic segmentation of T2* weighted slices 86
5.3 Results 94
5.3.1 nRMSE threshold 95

5.3.2 General Observations 96
5.3.3 Fetal vs Preterm 97
5.3.4 Preterm vs Preterm At Term 99
5.3.5 Coefficient of Variation 99
5.3.6 Correlation with age 100
5.4 Conclusions and Discussion 101
5.4.1 Sensitivity to iron concentration 103
5.4.2 Sensitivity to BOLD 105
5.4.3 Variation 106
5.4.4 Simultaneous multi slice imaging 107
5.4.5 Segmentation of subjects below 30 GA 107

6 Slice to volume reconstruction for sparsely sampled image space 109

6.1 Abstract 109
6.2 Introduction 109
6.3 Literature Review 111
6.3.1 Generative Algorithms 111
6.3.2 Conditional networks 113
6.3.3 Preserving high and low frequency information 115
6.4 Method 116
6.4.1 Data 117
6.4.2 GAN Architecture 119

6.4.3 Motion simulation . . . 120
6.4.4 Reconstruction . . . 123
6.5 Results . . . 124
6.5.1 Network parameters . . . 124
6.5.2 Target Registration test . . . 126
6.5.3 Simulated Acquisition test . . . 128
6.5.4 Reconstruction . . . 128
6.6 Discussion and Conclusion . . . 131
6.6.1 3D GANs . . . 132
6.6.2 High resolution synthesis . . . 133
6.6.3 Spatially constrained GANs . . . 134

7 Conclusion **136**

7.1 Summary of Thesis Achievements . . . 136
7.1.1 Acquisition . . . 136
7.1.2 Reconstruction . . . 138
7.1.3 Clinical Significance . . . 139
7.2 Applications . . . 140
7.3 Future Work . . . 141
7.3.1 SMS . . . 141
7.3.2 T2 relaxometry . . . 142
7.3.3 Multi-contrast SVR for sparse data . . . 143

List of Tables

4.1 Diagnoses of the clinical group of subjects 63

4.2 A conventional gradient echo sequence was used to measure T2* values in an adult volunteer and compared against an accelerated single shot EPI readout on the same subject. 67

4.3 Fetal brain T2* values (mean and standard deviation) for normal and clinical subjects. Values from previous investigations into neonatal T2* are also listed. Both neonatal studies reported values for medial and lateral occipital lobes separately. For comparison these are displayed here as an average aggregate for occipital white matter. There is no significant difference between lateral occipital lobe (LOL) and medial occipital lobe (MOL) at 95% confidence interval for all cases, given the variance and size of these samples. 71

4.4 Parameter estimates for the linear correlation of T2* with gestational age for normal and clinical fetal subject groups. Pearson correlation coefficient (PCC) and corresponding p-values indicate congruency and significance of these estimates. 72

5.1 A summary of key acquisition parameters used for fetal and neonatal sequence variants that were optimized for single shot multi-echo GRE EPI on a 3T system. 83

5.2 Distribution of the acquired subject data. Due to limitations imposed by acquired resolution, all patients with gestational age of less than 30 weeks of age were excluded from the study. 85

5.3 T2* estimates (in ms) in 14 regions of interest. All measurements were made in the acquired image space via an automatic brain reconstruction and segmentation framework. 97

5.4 Pairwise comparisons of T2* mean estimates between subject groups using two tailed t-test with p-values reported at 95% confidence interval. Significant differences are found in cortical regions, cerebellum and limbic system between fetal and preterm subjects, indicating non linear change in T2* values. No significant differences are detected in any areas between preterm and preterm at term subjects. 98

5.5 Coefficient of Variability has no apparent trend between regions or groups. The highest percentage change in mean T2* estimates is detected between fetal and preterm groups in the occipital lobe, temporal lobe and cerebellum, which is consistent with prior post-mortem studies of fetal development. 100

5.6 Evaluation of correlation of T2* with age. R^2 and its p-value estimate indicate the strength and significance of the estimated correlation. The comparisons are made for all subject groups combined (ALL groups) and for neonatal subjects separately (Neonatal only). Only basal ganglia and thalamus show strong linear correlation of T2* estimates with gestational age in both comparisons. 101

6.1 Motion simulation parameters for translation (in mm) and rotation (in degrees) for ten different scenarios. 122

6.2 Final hyperparameters of the proposed neural network for synthesizing T2* slices from T2 slices. 125

6.3 Target registration error (TRE) for motion simulated slices from 10 different subjects, across 3 locations in the brain and 10 different transformation parameters (denoted as dof). The first target volume for registration is the ground truth T2 volume. The second target volume is the synthesized T2* volume (sT2*). Mean and standard deviation are reported in mm. TRE was computed by comparing the estimated location of the slice with regards to its original location in the ground truth T2* volume. 127

6.4 Mean target registration accuracy (TRE) calculated over all motion corrupted slices that were registered back to the ground truth T2* volume via Slice to Volume Reconstruction (SVR) algorithm. The comparison is made between different types of input to the SVR algorithm. Mean and standard deviation are reported in mm. 129

List of Figures

1.1 Behaviour of the magnetization vector in the rotating frame of reference, which rotates around the axis of B_0 at the Larmor frequency. a) the spins are precessing around the main magnetic field B_0, resulting in the net magnetization M_0 b) 90° RF pulse tips M_0 into the transverse plane c) immediately after the RF pulse the spins start to lose coherence and thus cause phase dispersion, which leads to signal attenuation over time. 4

1.2 Free Induction Decay is an exponentially decreasing oscillation at resonant frequency due to transverse relaxation. It is recorded immediately after the RF pulse tips the magnetization vector into the transverse plane. The upper part of the envelope of the oscillating signal indicates the signal decay as it is conventionally displayed in MRI diagrams. 6

1.3 Signal decay over time in the transverse plane. T2 relaxation is always longer than T2* relaxation as the latter includes the effects of magnetic field inhomogeneities. 7

2.1 K-space is described by a matrix with dimensions N_x and N_y, which denote the localization in the frequency encoding and phase encoding directions respectively. Δk_x and Δk_y refer to the spacing of the recorded frequency and phase components in k-space. This example shows a simplified k-space with only 7x7 matrix for ease of reference. All the subsequent sequence diagrams would represent a readout for a 7x7 k-space as well. 15

2.2 Sequence diagram for gradient echo: RF pulse (red), gradients (blue) and signal readout (green). Echo time (TE) and repetition time (TR) are marked. The phase encoding gradient, G_y, is applied repeatedly after each RF pulse at different strength, as marked by the stacked diagramatic representation. This conventional gradient echo readout is often used as a reference in studies that explore accelerated imaging and related artifacts that may arise henceforth. 17

2.3 k-space acquisition for a conventional GRE readout with a single phase encoding step per individual RF pulse. 18

2.4 Sequence diagram for multi-echo gradient echo, which represents a 'golden reference' scan for T2* relaxometry. The signal is read by switching the G_x gradient backwards and forwards multiple times to acquire a single line of k-space at varying TE values, thus resulting in multiple signal readouts along the T2* decaying signal. 19

2.5 Sequence diagram for gradient echo EPI. The entire k-space matrix of 7x7 is filled following a single RF pulse by fast switching of G_x gradients alongside the signal readout. The phase encoding is incremented progressively by following each switched G_x gradient with a 'blip' in the G_y gradient. 20

2.6 k-space acquisition for an accelerated readout with EPI. G_x and G_y rewinder gradients are turned on initially to navigate to the bottom of the matrix and then acquire the whole k-space by fast switching of G_x gradient. To move along k_y, the phase is incremented via 'blips' in G_y. 21

2.7 Sequence diagram for multi echo gradient echo EPI acquisition. In this example, k-space is traversed twice for each RF pulse, resulting in two fully sampled images of the same slice at two different TE values. This acquisition is highly resistant to motion. 22

2.8 An example of signal susceptibility losses at air-tissue interfaces during EPI readout and their correction. a) conventional gradient echo with short TE of 10ms b) gradient echo EPI image at TE of 20ms with visible susceptibility signal losses c) a B_0 field map d) corrected gradient echo EPI image by the method presented in Liu et al. (2006). Image courtesy of Liu et al. (2006) 31

2.9 An example of parallel imaging reconstruction from the landmark paper that introduced the SENSE technique Pruessmann et al. (1999). Top left: conventional reconstruction from multiple coils by using sum of the squares of the signal of the full FOV acquisition. Bottom left: conventional reconstruction from multiple coils by using sum of the squares of the signal of the half of the FOV acquisition. Top right: SENSE reconstruction of the full FOV acquisition. Bottom right: SENSE reconstruction of the half of the FOV acquisition. The aliased data in the bottom left image is disentangled by the SENSE algorithm. 33

2.10 The conjugate symmetry of k-space makes P and Q equivalent. As a consequence, under the perfect field homogeneity conditions, it is possible to reconstruct the k-space from half of the data in either k_x or k_y axis. 34

2.11 Sequence diagram for gradient echo EPI with spectral presaturation with inversion recovery pre-pulse that will suppress all fat signal in the tissue. The RF pulse can vary between 100-180°, which will invert the magnetization along B_0 and, at the same time, tip it partially into the transverse plane. The transverse signal is spoiled by G_y and G_x. The main RF pulse is timed such that the fat signal reaches zero mark along the longitudonal plane as it recovers its magnetization. 39

2.12 Sequence diagram for gradient echo EPI with spatial saturation preparatory pulse that will suppress the signal from a selected cross-section in the through-slice direction. Large spoiler gradients are turned on following the 90° RF pulse in order to destroy all transverse magnetization. 40

2.13 Example of EPI distortion correction by a way of a field map as published in the classic paper by Jezzard et al. (1995). a) shows the image acquired with a conventional gradient echo b) a map of field inhomogeneity, as described in (2.31). The numbers alongside a scale indicate the pixel shift c) gradient echo EPI acquisition, which resulted in significant geometric distortion and susceptibility signal loss d) geometric distortion is corrected via a field map 42

2.14 An example of an EPI nyquist ghost that's manifested as a repeating pattern shifted by half the FOV. These artifacts are most commonly caused by eddy currents generated due to fast switching of the gradients in EPI sequences, which in turn causes inconsistencies between the odd and even k-space readouts. Courtesy of Chung (2019) 43

2.15 Example of a single slice from a fetal brain slice stack acquired in the transverse orientation - a) with bias field uncorrected b) bias field corrected and 3D reconstructed c) smoothly changing bias field calculated for the given slice a). Courtesy of Tourbier et al. (2018) 45

2.16 A phantom scanned at two different k-space matrix sizes resulting in varying Gibbs' artifact. Left: Substantial artifacts appear as a result of the coarse sampling of the signal (small matrix size). Right: improved version of the same scan due to a larger matrix size. Courtesy of Chung (2019) 45

3.1 Top row: an example of a single slice stack acquired in the coronal plane. Significant motion is visible in the transverse and sagittal planes. Bottom row: High resolution slice to volume reconstruction of the fetal brain in 3D. Reconstruction is made possible by combining multiple slice stacks in three orthogonal planes, which yields a dense oversampling of the image space. Courtesy of Kainz et al. (2015). 55

4.1 A schematic representation of a single-shot fast field echo multi-echo EPI sequence used in the study, which results in a set of spatially aligned T2*-weighted images at different TEs. SL - slice selective gradient, PE - phase encoding gradient, FE - frequency encoding gradient, RF - radio frequency pulse. 60

4.2 Distribution of scanned subjects across gestational age in whole weeks. 62

4.3 Measured regions of interest include: FWM - frontal white matter, THA - thalamic grey matter, OWM - occipital white matter. Cerebral appearance at gestational ages of a) 22 weeks and b) 32 weeks. Echo 4 (199.8ms) is shown in cases. 65

4.4 a) Definition of 6 ROIs measured in a healthy adult brain, marked on a single shot multi-echo GRE EPI scan b) Estimated T2* map (in ms) c) Estimated nRMSE (as percentage) 68

4.5 a) single voxel T2* fit in a multi-echo GRE scan. T2* = 66ms, nRMSE = 0.02, region = OWM b) single voxel T2* fit in a multi-echo single-shot GRE EPI scan. T2* = 70ms, nRMSE = 0.02, region = OWM c) T2* fit in a multi-echo single-shot GRE EPI scan with large fitting error. T2* = 56ms, nRMSE = 0.09, region = FWM. 69

4.6 Examples of fetal scans at 22, 28 and 34 weeks of gestation. T2* weighted images for TE range 27ms-199.8ms and corresponding reconstructions of T2* maps. . . 70

4.7 An example of a T2* decay fit for a voxel in the fetal frontal white matter. T2* = 220ms, nRMSE = 0.015 70

4.8 Scatter plot for Frontal White Matter displaying individual subject T2* values against gestational age. Trend lines are drawn for normal and clinical subjects. Previous measurements reported by Rivkin. Linear trend lines for clinical subjects in frontal and occipital white matter were removed as the correlation between gestational age and T2* was not significant, as reported in Table 4.4. 72

4.9 Scatter plot for Thalamus displaying individual subject T2* values against gestational age. Trend lines are drawn for normal and clinical subjects. Previous measurements reported by Rivkin. 73

4.10 Scatter plot for Occipital White Matter displaying individual subject T2* values against gestational age. Trend lines are drawn for normal and clinical subjects. Previous measurements reported by Rivkin and Lee are also displayed. 74

5.1 An overview of the proposed processing framework used for obtaining segmentations on the individual T2* slices. The steps are described in detail in each subsection of Section 5.2.4 87

5.2 Top and middle rows: Consecutively acquired T2* weighted scans of the same brain slice, with increased TE value. Top row - the signal intensity in all images is shown relative to the first image (TE1). Middle row - the signal intensity in each image is normalized with respect to its min-max range. Bottom row: For the same slice - estimated T2* map, through slice susceptibility map, nRMSE map, brain segmentation, corresponding T2 slice to which the T2* weighted data was fitted. 88

5.3 An example of a single stack of slices acquired in the transverse orientation (right column). When these slices are viewed in other planes, such as saggital or coronal planes (rows 1 and 2 in the right column), the anatomy of the fetal brain is not coherent. Slice to Volume Reconstruction (SVR) is used to align all of the acquired slices from multiple stacks into a coherent 3D representation of the fetal brain (left column). 90

5.4 An overview of the processing scheme for mapping individual T2* slices to 3D reconstructed T2 volumes. 91

5.5 Example of registration between a 3D reconstructed T2 image (left) and a stack of T2* brain slices (center). A visual error comparison is shown on the right. . 92

5.6 Example segmentation on T2* weighted slices (echo 2), propagated from a 3D reconstructed T2 image. The labels correspond to the following regions: green - frontal lobe, blue - parietal lobe, light blue - temporal lobe, purple - basal ganglia, red - limbic system. 93

5.7 a) T2 weighted image b) T2* weighted image (echo 2) c) T2* map d) normalized root mean squared error map for the T2* relaxation fit e) through slice susceptibility component estimate This example depicts three regions of interest in the neonatal brain – one in the thalamus and two in frontal lobe. One of the frontal lobe areas exhibits high field inhomogeneity effects and signal loss due to proximity of air-tissue interface near the sinuses. This region has a high nRMSE with respect to the other areas that is not recoverable via through-slice susceptibility correction method. 95

5.8 Individual voxel T2* signal decay fit with a mono-exponential model (green) and a through slice susceptibility corrected model (red) in the neonatal brain. a) signal in the deep gray matter in the thalamus. Both fitting methods achieve alignment as the signal is unaffected by through slice susceptibility gradient with T2* estimates of 97ms and 95ms, with nRMSE estimates of 0.019 and 0.013, respectively b) signal in the frontal gray matter with high through slice susceptibility signal loss. Through slice susceptibility correction adjusts the T2* estimate from 146ms to 64ms, and lowers nRMSE from 0.118 to 0.054. 96

6.1 An overview of the proposed approach that includes a combination of slice to volume from real acquired samples and synthesized samples from another image contrast with sparse coverage of the brain. 117

6.2 An overview of the pre-processing pipeline used to prepare the input images for ingestion by the image to image translation network. 118

6.3 An overview of the processing steps used in the Target Registration Error test. . 121

6.4 An overview of the processing steps in the Simulated Acquisition test. Simulated rigid transformations for each slice were not larger than the fetal motion parameters measured and reported by [1] Jiang et al. (2007) and [2] Ferrazzi et al. (2014) 123

6.5 Examples of synthesized T2* slices from 4 different subjects. Top row: original T2 weighted slice. Middle row: original T2* weighted slice. Bottom row: synthesized T2* weighted slice from the T2 weighted slice. 126

6.6 A visual comparison between T2* synthesized volume and the original ground truth T2* volume, shown in two planes that are orthogonal to the plane in which the T2* slices were synthesized from the T2 target slices. Top row: Original ground truth T2* volume. Yellow outline are the edges of the synthesized T2* volume. Bottom row: Synthesized T2* slices concatenated into a slice stack. Yellow outline are the edges of the ground truth T2* volume. 127

6.7 An example of a slice stack that has been motion corrupted with a simulation. The slices in this example were rotated by 1 degree and translated by 1mm in one orthogonal plane with a 3 slice interleaved slice acquisition model. 128

6.8 Slice to volume reconstruction results for a single subject with varying levels of image data available to each reconstruction. Results are presented as a view in 3 orthogonal planes. A difference image is generated by comparing a given reconstruction result with the original T2* volume prior to motion simulation. . 130

Chapter 1

Introduction

The aim of this chapter is to introduce the reader to the principles of T2* relaxation in nuclear magnetic resonance and relate it through the rising complexity of quantum mechanical, molecular and biological phenomena that affects it. A brief discussion on the clinical need for T2* relaxometry is presented. Following this, the reader is introduced to the key concepts of biological changes that occur in the developing brain. A review of previous literature that measured T2* relaxation rates in the adult and neonatal brain is also presented. Finally, the reader is introduced to the challenges of performing MR imaging in-utero on the fetal brain.

1.1 Preface

The invention of Magnetic Resonance Imaging (MRI) is attributed to the works of Sir Peter Mansfield and Paul Lauterbur (Lauterbur et al. 1973; Mansfield et al. 1977). Since that time, MRI technology has undergone significant transformations to become a widely used diagnostic tool for medical imaging around the world. For example, in 2019 the Organisation for Economic Co-operation and Development reported 67 MRI exams per 1,000 inhabitants performed annually in the UK (OECD 2019). Advancements in ultra fast motion-robust MR imaging protocols has made it possible to acquire artifact-free 2D images of the fetal anatomy (Ertl-Wagner et al. 2002). Introduction of retrospective 3D slice to volume reconstruction techniques have enabled

for the first time the detailed study of structure and function in the fetal brain(Ferrante et al. 2017). In 2009, the ambitious effort to map the neural pathways of the adult human brain was launched by National Institute of Health, the Human Connectome Project. In 2013, European Research Council awarded 4 universities in the UK to undertake the equivalent study of the fetal and neonatal brain, the Developing Human Connectome Project. Since 2016, the advent of deep learning has made a profound impact on medical image analysis and is bound to continue pushing the envelope, especially in the challenging domains such as the fetal MR imaging. There remain a number of great and exciting challenges in the domain of fetal and neonatal scanning. Fully automated acquisition, reconstruction, localization and segmentation have yet to be reached, but are arguably already feasible with the recent technological advancements. The works in this thesis consider the use of ultra fast acquisition methods, slice to volume reconstruction and deep learning techniques to develop quantitative T2* imaging methods in the developing brain within clinically acceptable time frames and constraints. In the next sections we introduce the reader step by step to the aspects of quantitative T2* imaging.

1.2 Precession

In quantum mechanics, protons, neutrons and electrons possess a fundamental characteristic property called spin, which is an intrinsic form of angular momentum. This property is quantized into discrete values and can interact with electromagnetic fields. Nuclei with odd number of protons and neutrons will exhibit a magnetic dipole moment. A magnetic dipole moment is a vector quantity that relates the torque that is experienced by an object that is acted upon by an external magnetic field:

$$\mathbf{T} = \mathbf{m} \times \mathbf{B} \tag{1.1}$$

where $\mathbf{T}$ is the torque vector, $\mathbf{B}$ is the external magnetic field and $\mathbf{m}$ is the magnetic moment.

Hydrogen atoms form 10% of all the human biological tissue and are highly abundant throughout the body since they make up the content of water and fat. Hydrogen atoms have one proton in their nuclei, resulting in a magnetic dipole moment. In the absence of an external magnetic

field, the spin states of hydrogen nuclei are not observable, and are said to be degenerate. In the presence of an external magnetic field, a small but significant number of hydrogen protons will split into two possible orientations of its spin state – spin up and spin down. Spin up state is considered to be the low energy state and spin down is the high energy state. The overall energy of nuclei in a magnetic field will be determined by thermodynamics of the Boltzmann distribution. At temperatures that are close to absolute zero, most of the hydrogen nuclei will be in the low energy state. However, at normal human body temperatures, a slightly larger proportion of nuclei will be in low energy state than the high energy state. The nuclei in low energy state can be excited to the high energy state with discrete levels of electromagnetic radiation.

When hydrogen nuclei are placed in an external magnetic field B_0, they will precess around this field (Brown et al. 2014). The frequency of their precession is termed Larmor frequency:

$$\omega = \gamma B_0 \tag{1.2}$$

where γ is the gyromagnetic ratio that is determined by the ratio of the magnetic moment to the angular momentum of the nuclei. In the case of hydrogen, its gyromagnetic ratio is 42.58 MHz/T.

If electromagnetic radiation is deposited to a set of hydrogen nuclei that are placed in an external magnetic field, the nuclei can flip to the higher energy spin state. In order for this to occur, the energy that is deposited to these nuclei has to be delivered at Larmor frequency. In an MR system, this is done via a radiofrequency (RF) pulse, which generates a magnetic field B_1.

The sum of the magnetic dipole moments of the nuclei that are placed in an external magnetic field can be measured via net magnetization $\mathbf{M}$, as a vector in 3D space. At equilibrium, the net magnetization is aligned along the axis of B_0. However, a RF pulse that generates B_1 for a short period of time τ will tip the net magnetization into the transverse plane, away from the B_0, by an angle θ (see Figure 1.1):

$$\theta = \gamma B_1 \tau \tag{1.3}$$

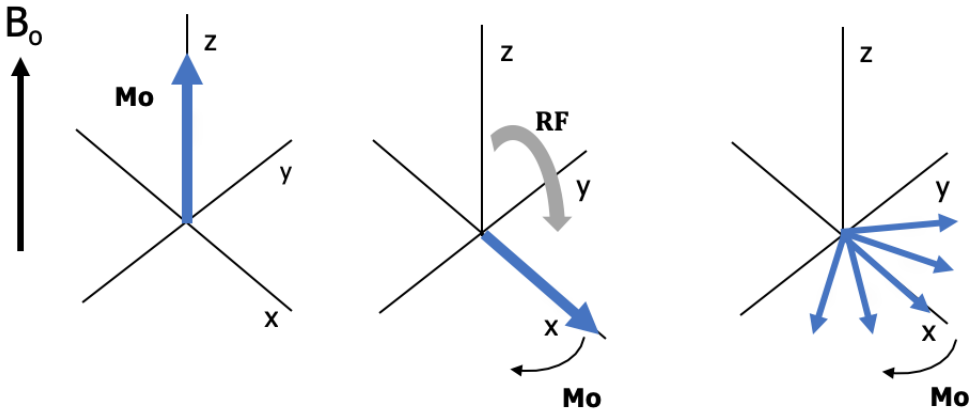

Figure 1.1: Behaviour of the magnetization vector in the rotating frame of reference, which rotates around the axis of B_0 at the Larmor frequency. a) the spins are precessing around the main magnetic field B_0, resulting in the net magnetization M_0 b) 90° RF pulse tips M_0 into the transverse plane c) immediately after the RF pulse the spins start to lose coherence and thus cause phase dispersion, which leads to signal attenuation over time.

When $\mathbf{M}$ is tipped into transverse plane with a finite amount of energy, which is deposited into the spin state of hydrogen nuclei, thereby pushing the low energy spins into high energy spins, the spin states will return to lower spin state after a certain amount of time. The return of the spin states towards the equilibrium can be recorded in the longitudinal and transverse planes via electromagnetic waves that are being emitted in the process. Such a process is known as relaxation. The evolution of $\mathbf{M}$ over time is modelled by Bloch equations (Bloch 1946):

$$\frac{d\mathbf{M}}{dt} = \gamma(\mathbf{M} \times \mathbf{B}) \tag{1.4}$$

where $\mathbf{B}$ is the magnetic field experienced by the nuclei and comprises from a combination of B_0 and B_1 fields. This formulation can be split into its component subparts as:

$$\begin{aligned}
\frac{\partial M_x(t)}{\partial t} &= \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_x - \frac{M_x(t)}{T_2} \\
\frac{\partial M_y(t)}{\partial t} &= \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_y - \frac{M_y(t)}{T_2} \\
\frac{\partial M_z(t)}{\partial t} &= \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_z - \frac{M_z(t) - M_0}{T_1}
\end{aligned} \tag{1.5}$$

where t is the time, $T1$ is the longitudonal relaxation time and $T2$ is the transverse relaxation time.

1.3 Relaxation

The return of the net magnetization to equilibrium is described by two components – longitudinal relaxation and transverse relaxation.

Longitudinal relaxation time, or T1 relaxation, is the time that is required for M_z, the measure of magnetization along the z-axis, to return to 63% of its original value from the initial magnetization state of $\mathbf{M}_0$:

$$\mathbf{M}_{\mathrm{z}}(\mathrm{t}) = \mathbf{M}_0 \left(1 - \mathrm{e}^{-\mathrm{t}/\mathrm{T1}}\right) \tag{1.6}$$

Longitudinal relaxation is governed by mobility of the nuclei in their surroundings, while the signal grows back to its initial state. The energy that is lost in the process is deposited into the surroundings as vibration and rotation of particles.

Transverse relaxation time, T2, is represented by the time that it takes for the signal that is measured in the transverse plane to decay to 37% of its maximum initial value:

$$\mathbf{M_{xy}}(t) = \mathbf{M_0} e^{-t/T2} \tag{1.7}$$

The main cause of transverse relaxation is a loss of phase coherence of spins over time in the transverse plane following a RF pulse, which results in an exponentially decaying sinusoidal wave, also known as Free Induction Decay (FID) (see Figure 1.2). In general, there are two types of interactions that spins experience during transverse relaxation. The first type arises from interactions with other spins, which is termed spin-spin relaxation, and is described by T2 time constant. The second type is caused by interactions of spins with the irregularities in the local static magnetic field. These irregularities can arise either from general large scale field inhomogeneities due to imperfections of the magnetic field, or the local susceptibility changes in the magnetic field caused by intra- and inter- molecular magnetic dipole moments. The second

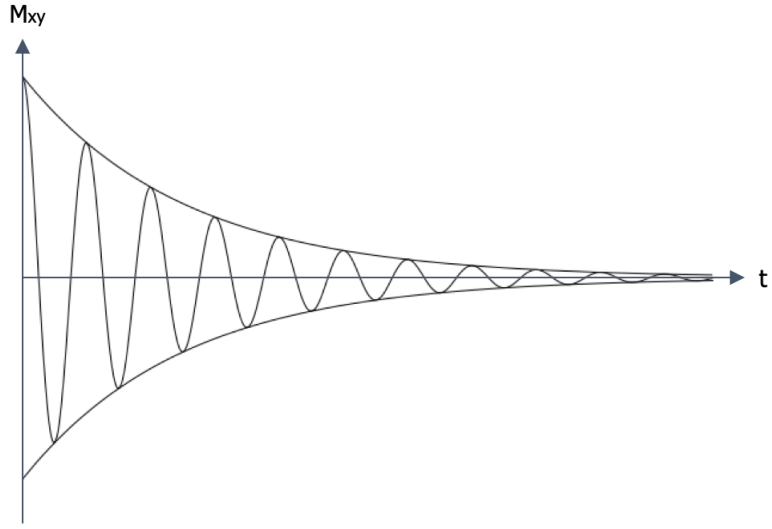

Figure 1.2: Free Induction Decay is an exponentially decreasing oscillation at resonant frequency due to transverse relaxation. It is recorded immediately after the RF pulse tips the magnetization vector into the transverse plane. The upper part of the envelope of the oscillating signal indicates the signal decay as it is conventionally displayed in MRI diagrams.

type of interactions is represented by:

$$\frac{1}{\mathrm{T}_2'} = \frac{\gamma}{2\pi}\Delta B_{\mathrm{inhom}} \tag{1.8}$$

where ΔB_{inhom} where is the local field inhomogeneity. The general combined model is described by the effective transverse relaxation rate, R2*:

$$R_2^* = \frac{1}{\mathrm{T}_2^*} = \frac{1}{\mathrm{T}_2} + \frac{\gamma}{2\pi}\Delta B_{\mathrm{inhom}} \tag{1.9}$$

where T2* is the effective transverse relaxation time (see Figure 1.3). Conventionally, it can also be referred to as the T2* decay constant and T2* relaxation time.

1.4 Molecular basis of relaxation

The early experiments on relaxation processes in nuclear magnetic resonance were conducted on homogenous samples such as certain solids with crystalline lattice structure. Incidentally, this

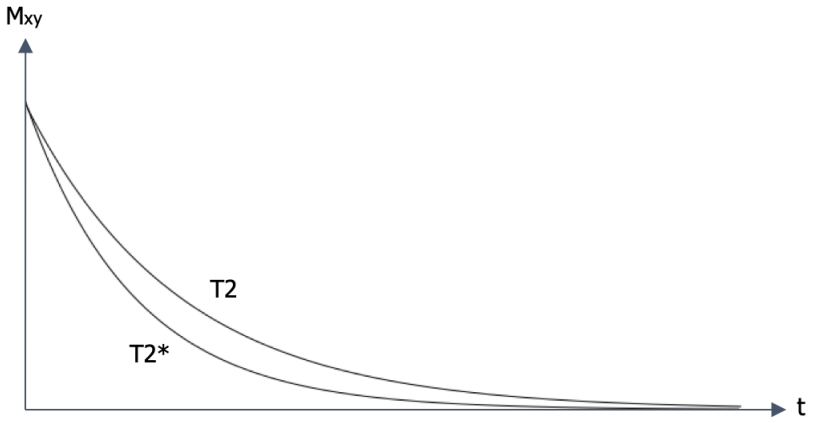

Figure 1.3: Signal decay over time in the transverse plane. T2 relaxation is always longer than T2* relaxation as the latter includes the effects of magnetic field inhomogeneities.

gave rise to the commonly used synonym for T1 relaxation - spin-lattice relaxation. However, relaxation processes in non-homogenous molecular environments are much more complex. Before we move on to discuss the basis of transverse relaxation in biological tissues, an outline of the basic principles of relaxation in complex molecular environments is provided. Spin interactions for hydrogen protons that cause relaxation in molecular environments can be described by four major mechanisms – dipole-dipole interactions, molecular motion, chemical shift and chemical exchange (Bloembergen et al. 1948).

Dipole-dipole interaction is the key mechanism for relaxation of spins for hydrogen nuclei and is caused by interactions between dipoles, such as protons or free electrons. The exchange of the magnetic fields between the dipoles leads to loss of phase coherence. Dipole-dipole interactions are mostly caused by intra-molecular effects as the strength of the exchanges decreases significantly over distance, and most significantly between proton-electron pairs, rather than proton-proton pairs.

Chemical shift is induced by electrons that shield the nuclei from the applied magnetic field, causing distortions in the local magnetic field. This process depends on orientation of particles with respect to B_0. As a result, it affects molecules that are restricted in motion. One example of this may be a water molecule that is bounded to a protein layer and is not able to move.

Such restrictions, cause the transverse relaxation rates to decrease much faster in comparison to theoretical computations which may have been calculated from independent components of molecular constituents alone.

Hydrogen atoms form part of many chemical exchanges that take place in biological tissues. Hydrogen proton exchanges, binding or separation within molecular structures results in faster relaxation times due to local changes in susceptibility effects.

Another mechanism of spin interactions is motion of molecules in non-homogenous environments. Molecular motion may be in bulk, such as flow of blood, or singular motion of individual molecules through a membrane, such as water diffusion.

In general, most chemical exchanges, molecular motion and complex structural environments of molecular composition result in shortening of relaxation time constants due to effects of irregularities in the local static magnetic field.

1.5 Relaxometry

Molecular structures and processes, and properties of the local magnetic field, give rise to varying relaxation times for different tissues. By controlling various MRI parameters measured signals can be weighted to achieve a certain type of relaxation that dominates in the response. For most clinical applications, MRI imaging takes use of this phenomena to reconstruct images that can be analysed for structural coherence to detect anatomical and pathological features. However, such images are qualitative, as the received signal is weighted and will change with varying MRI parameters. The signal intensity in such images cannot be compared if the MR parameters are changed or the patient is scanned at a different facility.

In contrast, MRI relaxometry is the study of relaxation values that are governed by the underlying tissue microstructure. The measure of relaxation values is replicable across different scanners and MRI parameters, as long as the biological properties of the tissue of interest do not change (Carneiro et al. 2006). This renders the signal recorded in such images as a quantitative

representation of the underlying biological tissue. The knowledge of relaxation times for specific tissue types across population groups allows estimation of normative values and their variations due to pathology (Deoni 2010).

1.6 Biological basis of T2* relaxation

The focus of this thesis is the study of T2* relaxation values and development of T2* relaxometry techniques for the developing brain. In this section, we discuss three areas where T2* relaxation values can provide valuable insights - as a clinical biomarker for myelination and iron concentration, as well as a tool for achieving optimal fMRI contrast.

Long stretches of myelin sheaths play a vital role in axonal signal propagation. Their lipid rich microstructure restricts the movement of water molecules within small reservoirs between myelin bilayers, separating it from intra/extraceullular stores of water (Laule et al. 2018). Prior studies have shown that T2* relaxometry can be used as a biomarker for myelin microstructure (Hwang et al. 2010b; T.-Q. Li et al. 2009) and is dependent on orientation and size of myelin fibers (Bender et al. 2010; Cherubini et al. 2009).

Proliferation of glial cells that support the function of neurons can have an effect on local T2* relaxation values (Cohen-Adad 2014). In particular, iron deposits that are not associated with haemoglobin molecules (non-heme iron), are predominantly stored as ferritin in the cortex and are used as a supply for oligodendrocytes, which in turn sustain the myelin sheath lining (Todorich et al. 2009). Iron stores are paramagnetic which manifests in strong susceptibility effects that distort the local magnetic field and therefore cause faster spin dephasing (Fukunaga et al. 2010). As a result, quantitative T2* mapping is considered to be a reliable marker for iron concentration in the brain (Langkammer et al. 2010).

Knowledge of T2* relaxation times is of substantial value for optimizing contrast in functional MRI studies. De-oxygenated hemoglobin contain unpaired iron molecules, which have at least one unpaired electron. This leads to the effect of nuclear shielding, causing the molecules to become paramagnetic (Ogawa et al. 1993). In contrast, oxygenated hemoglobin has no unpaired electron,

thereby making the molecule weakly diamagnetic. Paramagnetic deoxygenated hemoglobin in the blood creates susceptibility differences in the local magnetic field, yielding increased rate of T2* signal decay in the areas where blood is flowing in the brain. This phenomenon is the principal concept behind blood oxygen level dependeing imaging (BOLD) in functional MRI (Ogawa et al. 1990). Optimal contrast in fMRI is achieved when the echo time, which defines the time of signal formation and readout in the MR system, is set to the T2* relaxation time of the tissue of interest (Chavhan et al. 2009). This configuration yields the largest possible difference between the signal from deoxygenated blood and the surrounding tissue in the brain, thereby yielding the best possible contrast.

1.7 Developing brain

Brain maturation in the early period of life involves complex structural and physiological changes caused by synaptogenesis, myelination and growth of cells. The growth mechanisms are non linear and asynchronous (Dubois et al. 2014), particularly in the final stages of fetal gestation and in the first weeks of neonatal life. Post mortem studies showed that development of axonal pathways start with the regions that are proximal to the center of the brain and protract to the distal areas over time (Kinney et al. 1988). Myelination in white matter occurs much later in the somatosensory and visual systems than in other regions (Miller et al. 2012). Simultaneously, a large number of supporting cells are being established in the periphery of the neurons - astrocytes, oligodendricites and microglia, while neurons start to grow their dendritic trees (Dubois et al. 2014). These changes cause significant synaptic overproduction, soon followed by apoptosis and pruning (Counsell et al. 2003). All of these changes have a direct and substantial impact on longitudinal and transverse relaxation of the MR signal. The study of relaxation times can inform of such changes in the brain. Thus, the knowledge of relaxation times in the brain can help with evaluation of pathology, by establishing normative values of MR signal relaxation during fetal gestation and early neonatal development.

1.8 T2* in the newborn brain

At the time of this review, and with exception to the work that is presented hereon in this thesis, two prior studies of T2* relaxometry had been performed in the newborn brain. Rivkin et al. (2004) reported normative values of T2* in four regions of interest - frontal white matter, thalamus, medial occipital lobe and lateral occipital lobe. The measurements were studied in four subject groups: preterm born neonates scanned at 33 weeks of gestational age (GA), preterm born neonates scanned at 42 weeks of GA (also referenced as 'preterm neonates scanned at term' from herein), 9 months infants and adult subjects. The study concluded that the T2* relaxation values in the newborn subjects were 2-3 times higher than those of the adult subjects and significantly higher than those of 9 month old infants. Furthermore, the study observed substantial variation in T2* relaxation values in both neonatal subject groups, but a much tighter grouping in the infant and adult subject groups. W. Lee et al. (2012) measured T2* relaxation values in the medial and occipital cortex between 31 week GA preterm neonates and preterm neonates scanned at 44 week GA (also referred to as 'preterm neonates scanned at term'). The average T2* values were found to be similar to those measured by Rivkin et al. (2004). However, W. Lee et al. (2012) observed significant difference between T2* relaxation values of the the preterm and preterm at term groups in both regions of interest, which was in disagreement to the prior study of Rivkin et al. (2004). Notably, the infants scanned at preterm age in the W. Lee et al. (2012) study were significantly younger than those of Rivkin et al. (2004), which could explain this difference. A more detailed discussion on the results and methods of both studies are presented in the context of Chapters 4 and 5 of the thesis.

1.9 Fetal Imaging

Fetal MR is increasingly used as a complementary tool to ultrasound scanning to investigate inconclusive findings or to allow for screening when ultrasound imaging may be occluded. The key indications for fetal brain MR imaging include congenital abnormalities, screening for family risk of anomalies and vascular abnormalities (Saleem 2014). Fetal MR imaging is a significant

step up from neonatal or adult MR imaging in terms of complexity of the acquisition and reconstruction techniques that are required to provide artifact-free images with high signal to noise ratio. The biggest challenge is motion that manifests from maternal respiration and spontaneous fetal movements. Breath-hold scanning may be used for short sequences, although such sequences are limited in time and scope. Maternal or fetal sedation is generally not used. Consequently, most clinical scanning is focused on fast MRI sequences that can provide artifact-free 2D slices of the fetal anatomy in a breath-hold free environment (Malamateniou et al. 2013). To achieve this, highly accelerated scanning is performed by combining techniques such as single shot EPI (Poustchi-Amin et al. 2001a), partial Fourier (D. A. Feinberg et al. 1986) and parallel imaging (D. J. Larkman et al. 2007). Further acceleration can be achieved with the recently introduced simultaneous multi-slice imaging (Barth et al. 2016), which is quickly gaining momentum in clinical practice. When designing fast acquisition methods, careful considerations must be made with respect to the acoustic noise levels that arise from fast switching of MR gradients and heat deposition due to the multitude of radiofrequency pulses that are delivered into the fetal tissue (Saleem 2014).

Pervasive motion prevents acquisition of consistent 3D anatomy, which is a significant constraint on the analysis of the underlying pathology. Retrospective 3D reconstruction is an active area of research that has enabled quantitative study of the fetal brain development (Kainz et al. 2015). To date, it has enabled reconstruction of fetal brain T1 and T2 weighted structural images, diffusion tensor images and functional MRI images.

Chapter 2

Acquisition

The aim of this chapter is to introduce the reader to key acquisition sequences for performing T2* relaxometry and present a discussion of how these may be implemented within the constraints of fetal and neonatal MR imaging. At first, the reader is introduced to basic MR signal localization and T2* weighted sequence design. Next, three common methods for T2* relaxometry are presented. Methods are compared primarily by the speed of acquisition, which is the foremost requirement for avoiding motion related artifacts when scanning fetuses and newborns. Notably, fast acquisition sequences often present a significant challenge both in terms of imaging artifacts as well as the accelerated dephasing of T2* signals. The latter is particularly important as it limits the number of samples that may be used for robust estimation of underlying T2* relaxation times. Therefore, the next section discusses the key MR parameters that can be used to control the T2* signal decay. This is followed by an introduction to multiple methods that can further accelerate scanning, such as parallel imaging, partial Fourier acquisition, simultaneous multi-slice imaging and others. A brief overview of artifacts is followed with recipes for how they may be avoided or minimized. Finally, a short discussion on the importance of patient comfort is presented as a way to further improve the scanning outcomes.

2.1 Spatial Encoding

Consider a group of spins precessing in a static magnetic field B_0. In a perfectly homogeneous B_0 all the spins will all precess at Larmor frequency w_0 with identical phase ϕ_0. If a small magnetic field gradient $\mathbf{G}$ is super-imposed onto B_0, the Larmor frequency of the spins will vary spatially:

$$\omega(\mathbf{r}) = w_0 + \gamma \mathbf{G} \cdot \mathbf{r} \tag{2.1}$$

where r is a 3 dimensional position vector. If $\mathbf{G}$ is applied for a certain period of time, τ, the spins will accumulate phase, such that:

$$\phi(\mathbf{r}) = \phi_0 + \gamma \mathbf{G} \cdot \mathbf{r}\tau \tag{2.2}$$

The key principle for spatial localization of signal in MR imaging is based on variation of frequency and phase along a time-varying controlled magnetic field. If $\mathbf{G}$ is decomposed into 3 orthogonal linear gradients $\mathbf{G} = (G_x, G_y, G_z)$, the signal can be localized in three dimensions by applying the gradients in a sequence of steps. By convention, G_x, G_y, G_z are referred to as frequency encoding (FE), phase encoding (PE) and slice selective gradients respectively.

A simple conventional sequence starts with a RF pulse that is applied to the system together with a superimposed G_z gradient. The modulated profile of the RF pulse carries a range of frequencies that will only affect the spins that are precessing within a corresponding Larmor frequency range. The Larmor frequency range is controlled spatially by G_z. Therefore, a slice profile is given by:

$$\Delta z = \Delta w / \gamma G_z \tag{2.3}$$

where $\Delta w = 2\pi \Delta f$, and is the range of Larmor frequencies matching the bandwidth of the RF envelope. In the next step, G_y is applied for a certain period of time τ such that the spins will accumulate phase difference along the direction of the phase encoding gradient, which is typically assigned to y-axis:

$$\phi(y) = \gamma G_y y \tau \tag{2.4}$$

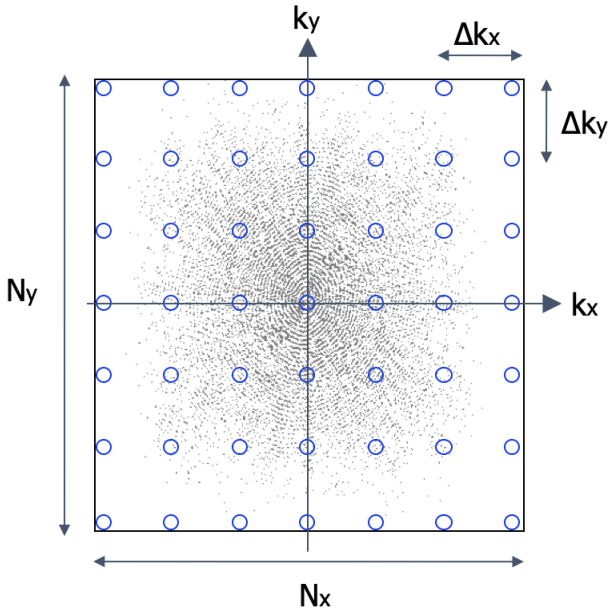

Figure 2.1: K-space is described by a matrix with dimensions N_x and N_y, which denote the localization in the frequency encoding and phase encoding directions respectively. Δk_x and Δk_y refer to the spacing of the recorded frequency and phase components in k-space. This example shows a simplified k-space with only 7x7 matrix for ease of reference. All the subsequent sequence diagrams would represent a readout for a 7x7 k-space as well.

After this, G_x is applied to vary the frequency of precession along the x-axis:

$$\omega(z) = w_0 + \gamma G_x x \tag{2.5}$$

While G_x is turned on, the signal is sampled by the receiver coil and digitized by the analogue to digital converter. In a simple conventional sequence, the process is repeated with varying levels of G_y on every iteration.

2.2 K-space

The recorded signals of discretized frequencies and phases are mapped to a 2D matrix. The matrix represents a theoretical concept known as k-space (see Figure 2.1). By convention, the columns, of which there are N_x, correspond to signals sampled while G_x was turned on, and the rows, of which there are N_y, correspond to signal sampled at varying levels of G_y. To formally

define the k-space matrix, we denote:

$$\begin{aligned} k_x(t) &= \int_0^t G_x(t')dt' \\ k_y(t) &= \int_0^t G_y(t')dt' \end{aligned} \tag{2.6}$$

where t is the time during which each of the gradients is turned on for each sampling period, also known as the dwell time. The resulting time-varying signal can therefore be written as:

$$S\left(k_x(t), k_y(t)\right) = \iint \rho\left(x, y, \Delta z\right) \mathrm{e}^{-i\gamma(k_x(t)x + k_y(t)y)} dxdy \tag{2.7}$$

where ρ is the spin density within the excited slice Δz. The above equation, (2.7), describes the Fourier transformation between the frequency encoded data and the image space. By computing its inverse, an image can be recovered to represent the spatial localization of spins.

2.3 Gradient Echo

T2* weighted images are obtained with a gradient echo (GRE) sequence, which is shown in Figure 2.2. In a conventional GRE sequence, a single RF pulse is applied in tandem with G_z to tip the net magnetization into the transverse plane. It is followed by a rephasing gradient in G_z that will reverse the accumulated phase along the z-axis. A phase encoding gradient, G_y, is applied for a short period of time to create a phase difference along y-axis, such that the signal can be localized along y-axis. Next, a dephasing gradient G_x is applied to delay the formation of an echo. The process is then reversed with the final gradient in G_x, which causes an echo to be formed and read by the receiver coil, which is referred to as an echo time (TE). The two step approach along x-axis, which involves a dephasing gradient, is necessary to ensure that the strongest signal from the echo is recorded at the center of k-space, which defines the contrast of the final image. At the next iteration of this pulse sequence, a different G_y is applied to sample the next set of spins, and consequently fill in the next set of samples in k-space (see Figure 2.3). Once the entire k-space has been sampled, the RF pulse envelope is modified to contain a new range of frequencies that will excite a different cross section along z-axis.

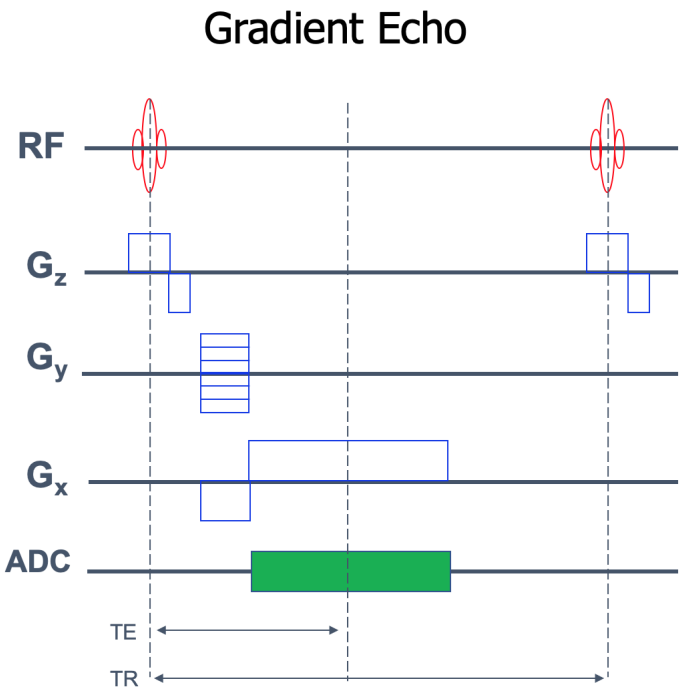

Figure 2.2: Sequence diagram for gradient echo: RF pulse (red), gradients (blue) and signal readout (green). Echo time (TE) and repetition time (TR) are marked. The phase encoding gradient, G_y, is applied repeatedly after each RF pulse at different strength, as marked by the stacked diagramatic representation. This conventional gradient echo readout is often used as a reference in studies that explore accelerated imaging and related artifacts that may arise henceforth.

2.4 Sequences for T2* mapping

2.4.1 Multi-echo GRE

It is possible to keep dephasing and rephasing the G_x gradient during readout, thereby forming multiple echoes following a single RF pulse over the decaying T2* signal. This forms the basis of the most basic approach to T2* relaxometry, termed multi-echo GRE (see Figure 2.4). This type of GRE acquisition has the least risk of artifacts as long as the patient stays still during the entire scan. However, such acquisition requires considerable time. This is because each RF pulse is applied with only one phase encoding (PE) gradient at one time, thereby allowing only one k-space line to be read. The time between successive RF pulses is referred to as repetition

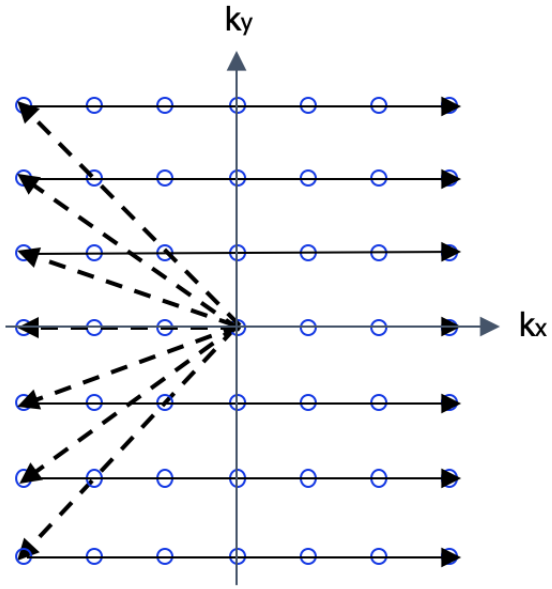

Figure 2.3: k-space acquisition for a conventional GRE readout with a single phase encoding step per individual RF pulse.

time (TR). GRE acquisition time is given by:

$$T_{acq} = NSA \cdot N_y \cdot TR \tag{2.8}$$

where NSA is the number of signal averages. As a result, the timing of this sequence scales linearly with the number of phase encoding gradient steps. In most scanning scenarios, a multi-echo GRE is used with spoiling at the end of each readout along the FE axis. Spoiling is a process by which any residual transverse magnetization is destroyed to avoid any crossover effects in the next readout. It can be achieved naturally by keeping TR very long. However, it is often impractical as the scan times become especially long, and therefore spoiling is performed by gradients or by RF pulses. Gradient spoiling is the preferred way of eliminating transverse magnetization in T2* relaxometry (Cohen Adad). RF spoiling may be chosen when strong eddy currents are generated (Epstein 1996), which we discuss in the follow up section 2.4.4.

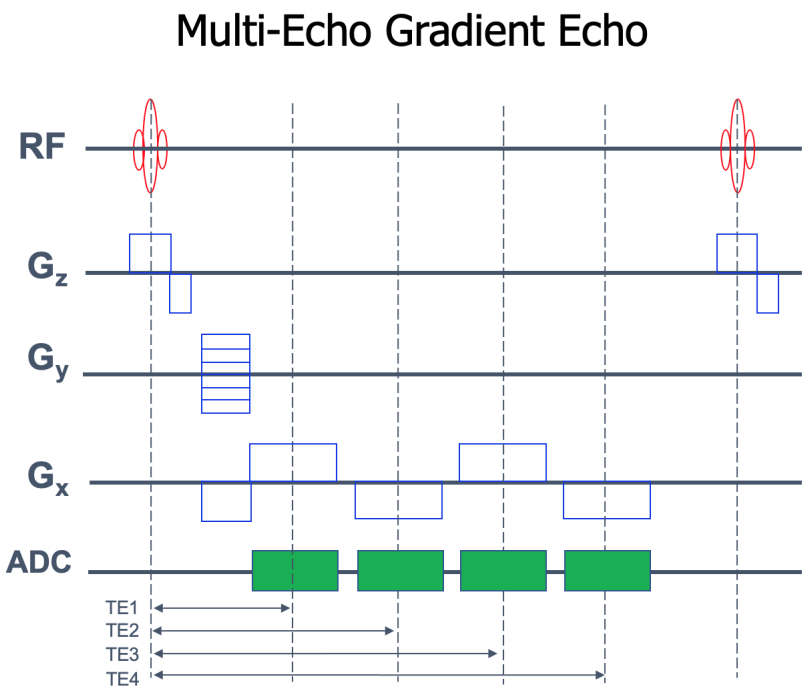

Figure 2.4: Sequence diagram for multi-echo gradient echo, which represents a 'golden reference' scan for T2* relaxometry. The signal is read by switching the G_x gradient backwards and forwards multiple times to acquire a single line of k-space at varying TE values, thus resulting in multiple signal readouts along the T2* decaying signal.

2.4.2 GRE EPI at varying echo times

GRE sequences can be combined with with echo planar imaging (EPI). EPI based GRE sequences can traverse multiple k-space lines after a single RF pulse by intermittently applying short "blip" gradients along the PE axis with strong switched gradients along the FE axis (see Figure 2.5). When an entire k-space of a single slice is covered after a single RF pulse, it is referred to as a single shot EPI scan (see Figure 2.6). As a result, a single slice can be acquired within a period that is less than 100ms, which eliminates most of the motion-related artifacts (Poustchi-Amin et al. 2001b). The entire adult brain image can typically be acquired within a time frame of 2-3s for an isotropic resolution of 2-3mm, thus requiring patient to be still for only a few seconds. Acquisition time for a GRE EPI sequence is given by:

$$T_{acq} = NSA \cdot TR \tag{2.9}$$

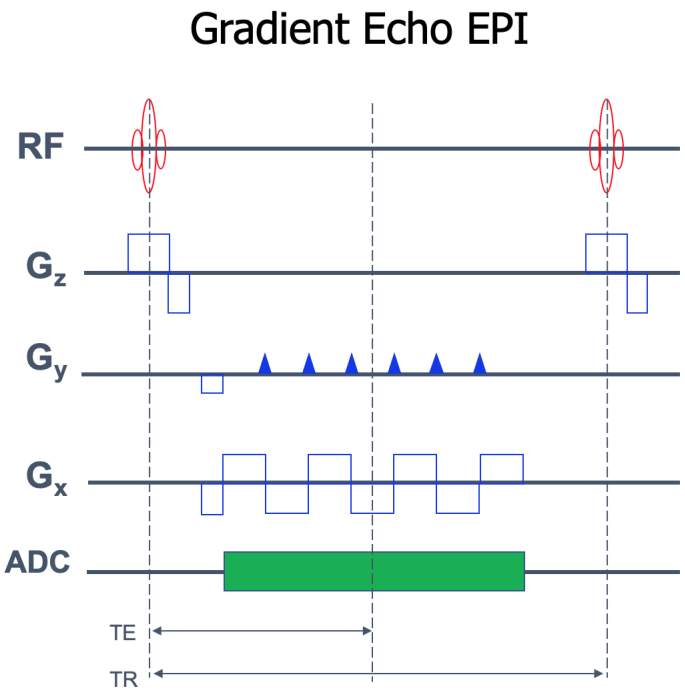

Figure 2.5: Sequence diagram for gradient echo EPI. The entire k-space matrix of 7x7 is filled following a single RF pulse by fast switching of G_x gradients alongside the signal readout. The phase encoding is incremented progressively by following each switched G_x gradient with a 'blip' in the G_y gradient.

We therefore note that this sequence length does not directly relate to N_y anymore.

It is further important to note that the TR in EPI must be long enough to incorporate the readout of the entire k-space. Therefore, generally, TR in GRE sequences is significantly shorter than TR in GRE EPI sequences.

To obtain multiple signal samples along the decaying T2* curve, GRE EPI scan can be repeated at different values of TE. As long as the patient stays still between the readouts of the entire slice stack, images can be concatenated temporally to form samples of the decaying T2* signal for each voxel. If small amount of motion occurred between the scans at different TEs, it can be corrected with rigid registration techniques by registering the 3D images via rigid registration (Ashburner et al. 2007). However, if motion occurred during one of the acquisitions, the scan would normally need to be repeated.

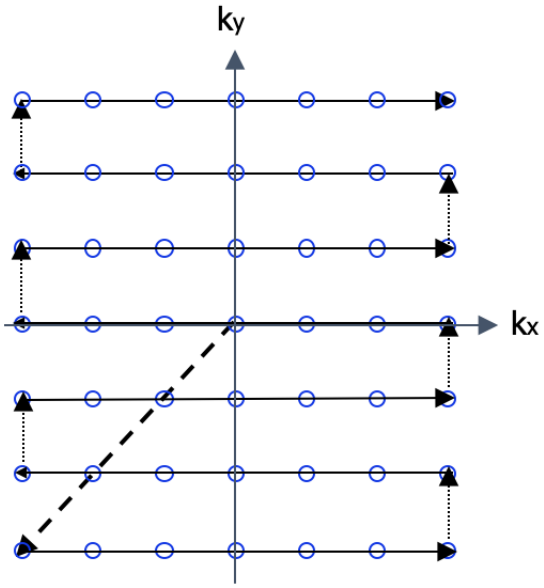

Figure 2.6: k-space acquisition for an accelerated readout with EPI. G_x and G_y rewinder gradients are turned on initially to navigate to the bottom of the matrix and then acquire the whole k-space by fast switching of G_x gradient. To move along k_y, the phase is incremented via 'blips' in G_y.

2.4.3 Multi-echo GRE EPI

It is possible to combine the principles of multi-echo GRE and GRE EPI to significantly shorten the overall scan time. In this instance, a single RF pulse is followed by a EPI readout of the entire k-space multiple times (see Figure 2.7). Rewinder gradients are added in both FE and PE axis after each whole k-space traversal is completed in order to return the scanning position to the initial starting point in k-space. Each RF pulse yields multiple images of a slice at different echo times, which constitute different samples along the decaying T2* signal.

The three methods that we described here are the key acquisition approaches for T2* relaxometry. Below we discuss the drawbacks and benefits of each in relation to fetal imaging.

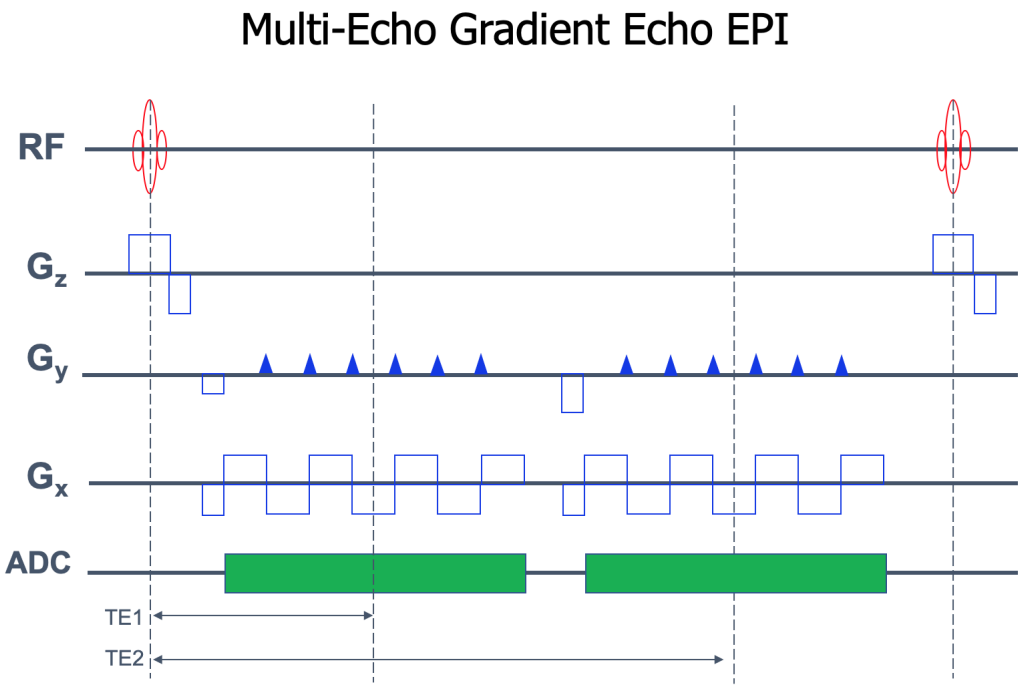

Figure 2.7: Sequence diagram for multi echo gradient echo EPI acquisition. In this example, k-space is traversed twice for each RF pulse, resulting in two fully sampled images of the same slice at two different TE values. This acquisition is highly resistant to motion.

2.4.4 Comparison of approaches

Multi-echo GRE

Multi-echo GRE is commonly used as a golden reference scan in T2* relaxometry as it can deliver very fine sampling of the T2* signal decay. The sampling rate of the T2* decay curve in this sequence is defined by the speed of rephasing and dephasing gradients along the FE axis, as well as the length of the acquisition matrix. On a modern Phillips Achieva 3T scanner, such as the one used for the work performed for this thesis, multi-echo GRE could yield a sampling rate of up to 3ms between successive signal samples. However, a typical multi-echo GRE sequence is very long and is therefore prone to artifacts if any motion occurs. As an example, an acquisition matrix of 256x192 with a TR of 100ms, would yield an acquisition time of approximately 20 seconds per slice. Any motion within this period would introduce artifacts into the imaged slice.

Single shot GRE EPI at varying echo times

In contrast, a single shot GRE EPI would require less than 100ms per slice. With additional accelerations, such as the ones that we describe in the next section 2.6, slice acquisition rate of 20-40ms can be achieved.

Nevertheless, if motion is encountered between successive slice readouts, the spatial position of slices will be inconsistent. In the case of fetal imaging, where maternal respiration is not constrained and the fetus is freely moving in space, the acquired slices would be scattered within a three dimensional space. In a conventional acquisition, such as a T2* weighted scan that requires acquisition at only a single TE value, a 3D image may be built retrospectively via slice to volume reconstruction (SVR) framework, which we describe in detail in section 3.7.

However, in T2* relaxometry each slice has be acquired multiple times at different TE values in order to sample the T2* decay curve. Slices that are acquired at long TE values, where the signal has dephased significantly, would yield images with very poor contrast and low SNR. If these slices are scattered in space, SVR could theoretically be used to align them to a 3D volume. However, SVR process relies on rigid image registration, which is highly dependent on image contrast and SNR. Therefore, SVR methods will yield to poor performance for slices that were acquired at late TE values. As a result, the use of single shot GRE EPI acquisition is limited in such case, unless only the slices at short TE values are used to estimate the T2* constants.

Multi-echo single shot GRE EPI

Multi-echo single shot GRE EPI is the fastest scanning technique that is available for sampling T2* decay curve that is highly robust to motion. The key advantage of this technique is formation of multiple echo images of a single slice at different TE values following a single RF pulse. As the timing between successive readouts of k-space is small, the reconstructed images of a single slice at different TE values are aligned with each other even if motion occurs.

Nevertheless, multi-echo single shot GRE EPI acquisition is highly limited by the rate of spin

dephasing due to the prolonged readout that is required for each k-space traversal. As a result, it is a suboptimal technique for performing T2* relaxometry in the adult brain, where the T2* values are short. However, the initial estimates of T2* values in the developing brain, conducted by the works in this thesis, confirmed that T2* values in the developing brain are considerably longer, thereby making it feasible to use this technique for T2* relaxometry in this specific patient group.

2.5 Controlling T2* signal loss

T2* signal decay is driven by the underlying biological tissue susceptibility properties as well as the imperfections in the imaging hardware. Spins that absorbed the energy from a RF pulse start to lose coherence in their phase, leading to concomitant attenuation in the recorded signal (L. Czervionke et al. 1988). When spins precess at a Larmor frequency that is different to that of the original RF pulse, the result is termed as off-resonance. Any inhomogeneities in the magnetic field cause off-resonance effects which in turn leads to faster decay of the T2* signal (T. B. Smith et al. 2010).

Accelerated imaging methods, such as the single shot EPI, are particularly prone to off-resonance effects. Therefore, it is vital to control the factors that may increase off resonance effects in gradient echo EPI scanning in order to safeguard the T2* dephasing for as long as possible. Prolonged decay of transverse T2* weighted signals allows for finer sampling of the decay curve and more robust T2* relaxation rate estimation.

This section gives a brief overview of the key factors that cause off-resonance and key MR parameters that may be used to control T2* signal decay. Specific attention is given to these factors in relation to single shot EPI, which is employed for data acquisition in this thesis.

2.5.1 Shimming

Three dimensional localization of signal in MRI scanners is estimated by the differences in frequency and phase of the spin precession. Estimation of signal is therefore highly dependent on the homogeneity of the underlying B_0 field.

Correction of baseline B_0 inhomogeneities that may arise from imperfections in the baseline hardware is performed via passive shimming (Wachowicz 2014). It consists of placement of ferromagnetic materials into the bore of the scanner. Overtime these shims themselves become magnetized and are therefore subject to regular adjustments (Jezzard 2006).

Placement of a patient inside the scanner distorts the magnetic field homogeneity due to the inherent and spatially variable susceptibility in the biological tissue. Therefore, a routine clinical examination is always preceded with active shimming to adjust the field. Active shimming techniques use additional gradient coils to superimpose spatially varying magnetic fields to adjust the B_0 inhomogeneities by formulating an ideal field as a set of spherical harmonic functions (Juchem 2013). While 1.5T scanners would typically shim with linear shim coils in the three spatial directions, 3T scanners are often equipped with second orders shims that can perform approximations within the skew lines of the 3D coordinate reference system as well as their curved surfaces.

Importantly, if the patient moves in the scanner during an examination, the field inhomogeneities will fluctuate with the new position. Therefore, re-shimming is advised in such scenario and is routinely performed.

2.5.2 Voxel size

The resolution and the extent of the reconstructed image can be formally defined in terms of k-space sampling, $\Delta k_{x,y}$:

$$FOV_{x,y} = \frac{1}{\Delta k_{x,y}} \tag{2.10}$$

And the number of lines acquired along each direction of k-space, N_x and N_y:

$$\Delta x = \frac{\mathrm{FOV}_x}{N_\mathrm{x}} \qquad \Delta y = \frac{\mathrm{FOV}_y}{N_\mathrm{y}} \tag{2.11}$$

In practice, these factors are controlled by the bandwidth of the RF envelope (BW_t), the strength of the gradients in all three dimensions as well as the sampling frequency of the receiver coil (BW_r).

While an ideal scan would possess the highest possible resolution, a set of trade offs must be made that will effectively limit the minimum voxel size. Voxel size directly affects a number of characteristics, such as signal to noise ratio (SNR), acquisition speed and partial volume effect. The SNR can be modelled by:

$$SNR = k\frac{\Delta x \Delta y \Delta z \sqrt{N_x N_y NSA}}{\sqrt{BW_r}} \tag{2.12}$$

where k is a scaling factor that depends on coil sensitivity, noise, B_0, TR, TE.

The robustness of T2* estimates is highly affected by the SNR as variation in noise can cause significant fluctuations in the T2* estimates. In sections 3.3 and 3.4 we discuss strategies of overcoming these via post-processing methods. However, reasonable effort should be taken to acquire sufficient amount of samples at high SNR levels.

For multi-echo GRE EPI, finer voxel sizes drive longer acquisition times, which directly limits the number of samples that are available to be acquired along the decaying T2* curve. Furthermore, smaller voxels lead to smaller SNR, which further limits the available samples that can be used for T2* estimation.

Conversely, larger voxel sizes bear higher risk of partial volume effects, leading to a signal that is weighted by contributions from different tissue types. This is particularly problematic for T2* relaxometry measurements where whole regions of interest may be taken for estimating average T2* relaxation values, as partial volume effects will skew the mean measurements. In these cases, image erosion may be applied along the boundaries between different tissue types

in order to limit the effect of skewed T2* estimates due to partial volume effects.

2.5.3 Receiver Bandwidth

BW_r is controlled by the dwell time, τ_{dwell}, which is the interval at which the samples are being read and digitized by the analogue to digital converter. As the signal readout takes place during the application of G_x, BW_r is controlled by:

$$BW_r = \frac{1}{N_x \tau_{dwell}} = \frac{\gamma}{2\pi} G_x FOV_x \tag{2.13}$$

Off-resonance effects, which cause faster signal decay, can be improved by shortening the EPI acquisition time as there is less time for signals to dephase between each k-space line. The simplest strategy may be to reduce the size of N_y, which is the time limiting step in EPI. However, this decreases image resolution. Alternatively, acquisition can be accelerated with higher BW_r. The downside of higher BW_r is decreased SNR (Jezzard et al. 1998). Furthermore, higher BW_r requires stronger gradients for large FOV. Performance of stronger gradients is limited by the slew rate.

2.5.4 Slew Rate

Gradient rise time, also known as the slew rate, plays a vital role in the speed of k-space sampling and quality of image reconstruction. Fast slew rates improve gradient coil waveforms as the signal reaches its desired level faster and can be sampled by the receiver coil. In some circumstances, particularly for ultra fast traversal of k-space, the signal may be sampled during the gradient rise time. This requires sophisticated reconstruction techniques as the received signal has to be adjusted for the temporally varying strength of the field while the magnetic field gradients are rising.

In EPI imaging, slew rate, alongside receiver BW_r, is the key factor that defines the speed of scanning. However, while fast slew rates are generally desirable, they can also be the key cause

of eddy currents. Eddy currents are time varying currents are induced within various conducting surfaces in the scanner and the human body. The strength of the eddy currents is proportional to the rate of change of the magnetic field and is therefore directly driven by the gradient switching during the EPI readout. Eddy currents produce time varying magnetic gradients within the locality of their surfaces, which distorts the B_0 homogeneity of the main magnetic field. This causes the Nyquist ghosting artifacts, which is described in the next section 2.7.3. Furthermore, the resulting off-resonance also induces faster dephasing of spins and therefore causes faster T2* signal decay.

Notably, the slew rate setting also plays another critical role in fetal and neonatal imaging. Fast slew rates, particularly for single shot EPI scanning, are associated with pronounced acoustic noise levels that can go above 120dB. Current clinical guidelines limits the level of acoustic noise that fetal and neonatal patients can be subjected to. Therefore, all new EPI protocols must adhere to the maximum allowable noise level. For example, in the study of Chapter 5, the slew rates of the acquisition had to be reduced in order to reduce the acoustic noise footprint in the scanner. In practice, quieter scans are also preferred for neonatal scanning as they bear less risk to wake up the newborn. Loud sequences can be disruptive to the clinical procedure and often lead to undesirable interruptions, and sometimes even abandonment of the scan. Therefore, a fine trade off needs to be achieved between the slew rate and the generated acoustic noise level.

2.5.5 Susceptibility induced signal loss

Susceptibility induced signal losses in GRE imaging are most visible near tissue-air interfaces of the nasal sinuses (see Figure 2.8). They arise due to the notable variation in magnetic susceptibility of air and tissue, which are paramagnetic and diamagnetic respectively (Cho et al. 1992). Soft tissues and bone tissue have similar susceptibilities of $\chi = -9.1 \times 10^{-6}$ and $\chi = -11.4 \times 10^{-6}$, while air has a susceptibility of $\chi = 0.4 \times 10^{-6}$ (Poynton et al. 2008). In 2D GRE sequences, this susceptibility manifests primarily in the through-slice direction and creates an additional background magnetic field gradient that acts on the MR signal (W. Deng et al. 2009). Typically, an assumption is made that this gradient is linear over the scale of small voxel

size (Fernandez-Seara et al. 2000).

To model this effect, consider the net magnetization that is detected through the receiver coil, and is modelled as a vector within two orthogonal planes. As a result, the additional effect of the background gradient, $G_{\Delta z}$, can be modelled a complex number via Euler's notation. Its effect will depend on the time that has passed since the RF pulse, which is recorded at TE, and the distance along the through-slice plane, z:

$$S(z, TE) = M(z, TE)e^{-i\gamma G_{\Delta z} TEz} \tag{2.14}$$

where $M(z, TE)$ is the term that includes all of the other factors that affect magnetization, such as spatially varying spin density, relaxation, off-resonance effects in the x-y plane and others.

To integrate the total signal across z, one must consider the slice acquisition profile. An ideal slice profile is a rectangular function (Rouphael 2009):

$$\Pi_{z_0}(t) = \begin{cases} S_0 & -\frac{z_0}{2} \leq t \leq \frac{z_0}{2} \\ 0 & \text{otherwise} \end{cases} \tag{2.15}$$

where z_0 is the slice thickness and S_0 is the magnetization immediately after the RF pulse. In Fourier terms, it is expressed as:

$$F\left\{\Pi_{z_0}(t)\right\} = S(k_z) = S_0 z_0 \operatorname{sinc}(\pi z_0 k_z) \tag{2.16}$$

where k_z is the frequency representation of the effect of the background gradient $G_{\Delta z}$. k_z is given by (Fernandez-Seara et al. 2000):

$$k_z = -\gamma G_{\Delta z} TE/2\pi \tag{2.17}$$

As the signal is not Fourier transformed along the z-dimension, we can convert back to image

domain. Substituting (2.17) into (2.16) and integrating with (2.14) we get:

$$S(TE) = M(TE)z_0 \,\mathrm{sinc}(\gamma G_{\Delta z} TE z_0/2) \tag{2.18}$$

Finally, we can account for the TE term in M by writing:

$$S(TE) = S_0 e^{-\frac{TE}{T2*}} z_0 \,\mathrm{sinc}(\gamma G_{\Delta z} TE z_0/2) \tag{2.19}$$

Signal losses are particularly evident in single shot EPI scans, where the fast switching of gradients, coupled with longer readout time, causes significant loss of coherence amongst the spins on the boundary between air and tissue interfaces over time. The effect is significantly enhanced as the echo time (TE) increases. Therefore, the final k-space samples within a multi-echo single shot GRE EPI readout are subject to the largest signal losses. Therefore, it is important to minimize the factors that may cause total loss of signal due to susceptibility induced difference near the air-tissue boundaries.

The general strategy to reducing susceptibility induced signal losses is to accelerate scanning as much as possible. Therefore, single shot EPI scanning is often combined with parallel imaging, partial fourier and other acceleration methods. Susceptibility differences may also be decreased with reduced voxel size (Ogawa et al. 1993). Smaller voxel size leads to less dephasing between the spins within different molecular environments that may reside within the same voxel.

Another strategy may be to alter the slice selection gradient (Frahm et al. 1988). The basic principle is to counterbalance the strong dephasing of spins by over compensating the slice selection gradient on the boundaries between tissue air interfaces. A conventional slice selection gradient is turned on simultaneously with the RF pulse in order to deliver the energy to the spins in a particular cross section of the body by matching the frequencies in the RF pulse with the Larmor frequencies of the selected spins. Typically, this slice selective gradient is then reversed in the opposite direction briefly after the RF pulse in order to reverse the dephasing effects that it has induced in the field and adjust the phase difference necessarily. The study by Frahm et al. (1988) that a small misadjustment to that reversal, such that the gradient

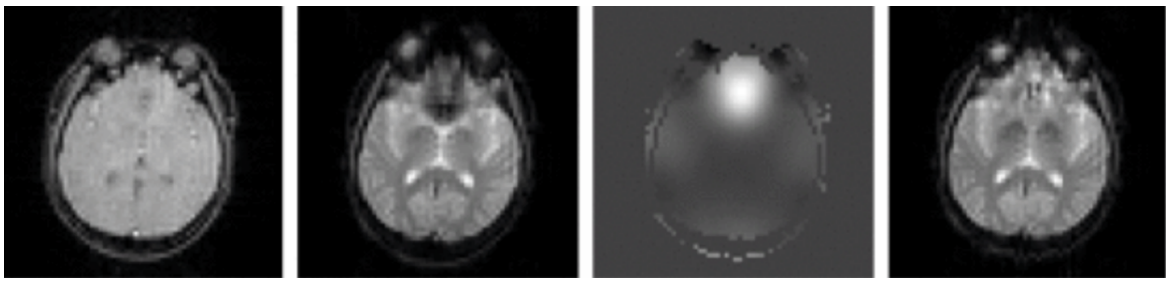

Figure 2.8: An example of signal susceptibility losses at air-tissue interfaces during EPI readout and their correction. a) conventional gradient echo with short TE of 10ms b) gradient echo EPI image at TE of 20ms with visible susceptibility signal losses c) a B_0 field map d) corrected gradient echo EPI image by the method presented in Liu et al. (2006). Image courtesy of Liu et al. (2006)

.

introduces small overcompensation in phase, would allow to recover the lost signal in magnitude temporarily. However, such approach would require the slice selection to be altered in both directions, thereby prolonging the scan.

It is possible to adjust the RF pulse such that its quadratic phase distribution along the slice selection direction is compensating for the dephasing of spins due to phase dispersion (Cho et al. 1992). This technique is effective but significantly reduces the SNR in the areas that are not subject to strong susceptibilities induced signal losses. Finally, lesser susceptibility signal losses may be corrected via field maps (Liu et al. 2006).

Two separate studies by Fernandez-Seara et al. (2000) and Dahnke et al. (2005) addressed the correction of susceptibility induced signal losses for correction of T2* relaxation time estimates. It is possible to use the multi-echo information to partially recover the signals, based on the formulation in (2.19). Fernandez-Seara et al. (2000) use an iterative post-processing method to calculate the optimal value of $G_{\Delta z}$. Dahnke et al. (2005) propose that the search for $G_{\Delta z}$ is initialized by considering the phase evolution across neighbouring slices over time. A detailed review of these methods is presented in Section 5.2.3 of Chapter 5.

Finally, it is important to note that the susceptibility induced signal loss near the sinuses is not present in fetal GRE imaging. This is because the fetal brain is surrounded by the amniotic fluid, which has smaller susceptibility differences with the biological tissue composition in the brain. Therefore, this correction would only apply to neonatal imaging in this thesis.

2.5.6 Phase Encode Direction

Susceptibility induced signal losses, geometric distortions, aliasing and other artifacts can often be improved for specific region of the FOV by positioning the phase encode direction to be along the axis of the shortest width of the brain. Typically, in neonatal subjects, this would be the left-right axis between the ears of the subject. In fetal scanning, this would be in the dorsa-ventral direction of the mother's anatomy. However, notably most fetal scans are acquired in the transverse orientation with respect to the fetal brain. As the brain can be positioning in any axis with respect to maternal anatomy, this phase encode 'trick' cannot always be used. The transverse plane is the most desired acquisition plane as it offers the best view for qualitative assessment of final images.

2.6 Motion Compensated Imaging

2.6.1 Parallel Imaging

Parallel imaging refers to methods that shorten the scan time by reducing the number of k-space lines that are read while keeping the extent of k-space the same. As a result, the Nyquist criterion is broken, which would normally yield aliased images (see Figure 2.9). Aliasing can be disentangled in image domain or k-space domain with SENSE or GRAPPA techniques respectively. GRAPPA reconstructs individual missing k-space samples by deriving the data from adjacent cells in k-space from multiple receiver coils. SENSE reconstructs the image by considering signals from multiple independent receiver coil with unique sensitivity profiles. For SENSE reconstruction, we can rewrite (2.20) as the signal recorded from multiple receiver coils:

$$S_j\left(k_x(t), k_y(t)\right) = \iint c_j\left(x, y, \Delta z\right) \rho\left(x, y, \Delta z\right) \mathrm{e}^{-i\gamma\left(k_x(t)x + k_y(t)y\right)} dxdy \tag{2.20}$$

The forward model of the reconstruction can be described as:

$$y_j = EFC_j x + \epsilon_j \tag{2.21}$$

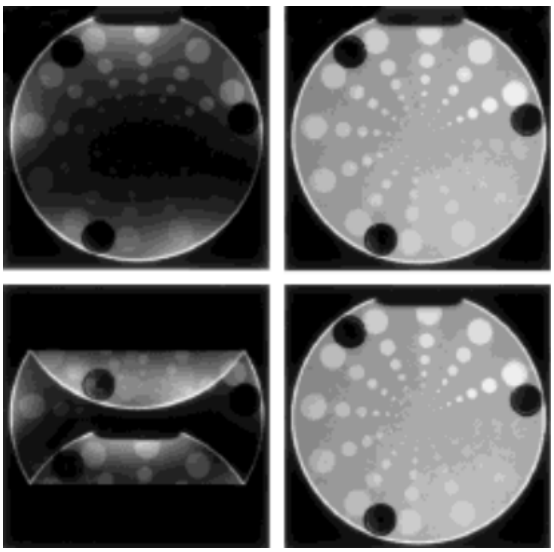

Figure 2.9: An example of parallel imaging reconstruction from the landmark paper that introduced the SENSE technique Pruessmann et al. (1999). Top left: conventional reconstruction from multiple coils by using sum of the squares of the signal of the full FOV acquisition. Bottom left: conventional reconstruction from multiple coils by using sum of the squares of the signal of the half of the FOV acquisition. Top right: SENSE reconstruction of the full FOV acquisition. Bottom right: SENSE reconstruction of the half of the FOV acquisition. The aliased data in the bottom left image is disentangled by the SENSE algorithm.

where y_j is the data sampled from the j-th coil, x is the image to be reconstructed, E is the encoding matrix, F is the Fourier transform, C_j is the sensitivity of coil j, and ϵ is the noise. If the undersampling factor in k-space is given as R, then a closed form solution to (2.21) can be modelled if the number of receiver coils is bigger or equal to R. However, typically, the method requires an iterative optimization method, such as gradient descent. A simple least squares solution is given as:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{EFCx}\|_2^2 \tag{2.22}$$

However, in practice this needs to be regularized (Majumdar et al. 2012), which yields:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{EFCx}\|_2^2 + \lambda\Phi(x) \tag{2.23}$$

where Φ is the regularization function and λ is the scaling factor. Parallel imaging is a necessary tool for EPI imaging as it allows alleviation of the strong susceptibility effects accumulated from the long EPI readouts (Yanasak et al. 2014). The reduced acquisition time leads to less distortion and signal losses. Therefore, it is particularly useful for multi-echo single shot EPI

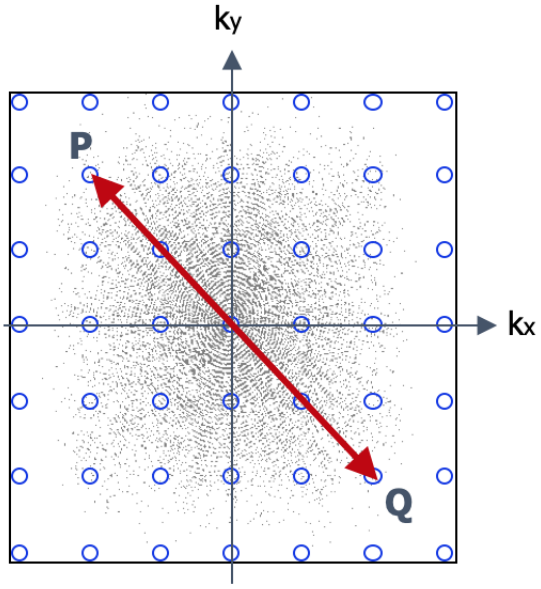

Figure 2.10: The conjugate symmetry of k-space makes P and Q equivalent. As a consequence, under the perfect field homogeneity conditions, it is possible to reconstruct the k-space from half of the data in either k_x or k_y axis.

sequences that we present in this thesis. Simultaneously, it is important to note that parallel imaging causes substantial losses to SNR and therefore should be adjusted carefully for very long TE values.

Generally, both SENSE and GRAPPA are widely used in clinical practice. The choices vary by manufacturer and neither method is considered to be superior in the general case. Overall, GRAPPA can be a slightly longer sequence for low resolution imaging as the calibration time of coils takes up a larger percentage of acquisition. At the same time, GRAPPA has a better tolerance to motion artifacts precisely for this reason, as autocalibration is performed at each scan. However, if a SENSE coil sensitivity calibration is performed just prior to the scan of interest, its performance to GRAPPA is comparable. In our case, we employed SENSE based parallel acceleration due to its commercial availability on the Phillips Achieva 1.5T and 3T MR systems that were used for the work in this thesis.

2.6.2 Partial Fourier

The definition of k-space yields it a conjugate symmetric property:

$$S(-k) = S^*(k) \tag{2.24}$$

As a consequence, it is symmetric with respect to the origin in both directions, FE and PE, with the reversed sign (see Figure 2.10). It is therefore possible to reconstruct the image with only half of the k-space in one of the chosen axis. This redundancy is used to speed up scanning by acquiring 60% or more of the k space. While 50% is the theoretical threshold, it would require ideal B_0 homogeneity and no noise. Therefore, in practice, a limited number of additional central k space lines are acquired to estimate the effects. When the symmetry is used to decrease the number of phase encoding steps, it is referred to as partial Fourier imaging. It is also possible to use an alternate strategy of acquiring a fraction of k-space in frequency encoding direction, however this alternative provides only marginal improvement in scanning time. The reduction of k-space matrix comes at the expense of SNR, which decreases by a square root of the fraction of the k space that is acquired.

2.6.3 Simultaneous Multi-Slice

An exciting new technique that has been developed and widely adopted across the MRI community relatively recently is simultaneous acquisition of multiple slices (SMS). The basis of the technique relies on composite RF pulses that carry phase shifted frequencies to the common slice gradient (Barth et al. 2016). The k-space that is generated from such images will carry information from multiple slices unless disentangled through coil sensitivity profiles, which is similar to SENSE (D. Larkman et al. 2001) or via a phase shift (Setsompop et al. 2012). The technique has been adopted widely in large scale studies such as the Human Connectome Project (Uğurbil et al. 2013) and Developing Human Connectome Project (Price et al. 2015, 2016a,b).

A basic RF pulse for can be modelled as a product of two functions:

$$RF(t) = A(t)e^{i\Delta\omega t+\phi} \tag{2.25}$$

where $A(t)$ represents the RF envelope with a select band of frequencies that will excite a range of spins if matched to their Larmor frequencies, and $e^{i\Delta\omega t+\phi}$ is the phase modulated profile of the pulse, which shifts the envelope to the necessary position along the slice direction. SMS modifies the second part of this formulation by packaging multiple phase shifted components into the RF pulse, which will excite multiple slices at the same time in conjunction with one slice selective gradient:

$$RF_{SMS}(t) = A(t)\sum_{N} e^{i\Delta\omega_n t+\phi_n} \tag{2.26}$$

SMS is particularly useful in EPI imaging as its SNR penalty is smaller than SENSE. One particular drawback of SMS imaging is the added deposition of energy through added RF pulses, which increase the relative specific absorption rate (SAR). Therefore, careful considerations and tests must be made when adopting it to fetal imaging domain. One common solution to reduce SAR in SMS imaging is to implement periodic excitation profiles with PINS imaging (Koopmans et al. 2013) or by varying the slice selective gradient over time to slow down the k-space traversal with VERSE RF pulses (Conolly et al. 1988).

2.6.4 Navigators

Additional RF pulses may be used to track motion during acquisition and are most commonly used to track diaphragmatic motion. A simple navigator pulse typically excites a small area perpendicular to the moving structure, which allows extraction of motion information in the post-processing step (Malamateniou et al. 2013). This design is very similar to respiratory gating that may be done with external devices. A more complex design allows for prospective correction of motion as the k-space acquisition is varied with the navigator readout during the 'dead' time of the scan, which is the time when no gradients or RF pulses are being applied within an active scan.

The use of navigators in fetal imaging has been limited (Kainz et al. 2015). Navigators have been used to trigger acquisition of single shot EPI when the brain anatomy is still and is along a certain axis (Bonel et al. 2008). However, this kind of acquisition would require a prolonged amount of time if the whole 3D brain was to be acquired in this manner. As a result, most fetal brain imaging studies choose to over sample the brain instead by acquiring individual 2D slices that are scattered in 3D space and perform a 3D reconstruction of the fetal brain via SVR algorithm (Studholme 2011).

2.6.5 Non-cartesian sequences

Most clinical fetal and neonatal imaging uses Cartesian grid to acquire the k-space data. Non-cartesian sampling schemes are another family of sequences that are considered motion resistant as they oversample the center of the k-space, from which the motion information can be inferred. However, most of these methods have not found wide adoption in fetal imaging yet (Malamateniou et al. 2013).

2.7 Artifacts

Fast imaging techniques, particularly those that use EPI, are subject to a variety of artifacts. In the next section 2.7 we describe these artifacts and their relevance to T2* relaxometry for motion compensated ultra fast imaging.

2.7.1 Chemical shift

The resonant frequency of hydrogen protons in fat molecules is different to that of water molecules due to the nuclear shielding effect described in section 1.4. The frequency shift is given by:

$$\delta f = 3.5 \text{ ppm } \omega_0 \tag{2.27}$$

where w_0 is the Larmor frequency at a given field strength and 3.5 is the ratio of resonant frequency shift between fat and water expressed in parts per million (ppm). Hydrogen protons in fat will shift by 220Hz at 1.5T (Brown et al. 2014). As a consequence, fat will be spatially shifted in the reconstructed image. In a conventional GRE scan, where each phase encode line is acquired with a dedicated RF pulse, the chemical shift will only be visible in the frequency encode direction:

$$\delta x = \frac{FOV_x \cdot \delta f}{BW_r} \tag{2.28}$$

where FOV_x is the extent of the image in frequency encoding direction, BW_r is the receiver bandwidth and δf is the chemical shift expressed in Hz (Hood et al. 1999). Notably, in EPI, the signal is sampled continuously in the frequency and phase encoding directions. Chemical shift will be significantly larger in the phase encoding direction as the signal is sampled at a much slower rate than in the frequency encode direction. The spatial shift in the phase encoding direction is given:

$$\delta y = FOV_y \cdot \delta f \cdot \tau_{pe} \tag{2.29}$$

where $\tau_{pe} = N_x \tau_{dwell}$ and is the timing between successive reads of k-space in the phase encoding direction.

A number of different strategies may be used to address chemical shift artifacts. Firstly, a higher sampling frequency in the frequency encoding direction will shorten the timing between the phase encode steps and thereby reduce the overall chemical shift in conventional GRE(Delfaut et al. 1999). Secondly, the number of phase encoding steps can be reduced at the expense of image resolution. More importantly, there exist a plethora of methods for cancelling the signal from fat via added RF pulses that come prior to the EPI excitation and readout. These are broadly divided into spectral and spatial saturation techniques. Spectral saturation, also known as Fat-Sat, or CHESS methods, uses narrow frequency range that is tuned to the Larmor frequency of hydrogen protons in fat to cancel its signal (Cameron 2013). The magnetization from fat is tipped into the transverse plane and is then spoiled with dephasing gradients in the frequency encode and phase encode gradients. Alternatively, a short TI inversion recovery (STIR) pulses may be used to null the effect of fat due to its short longitudonal relaxation values

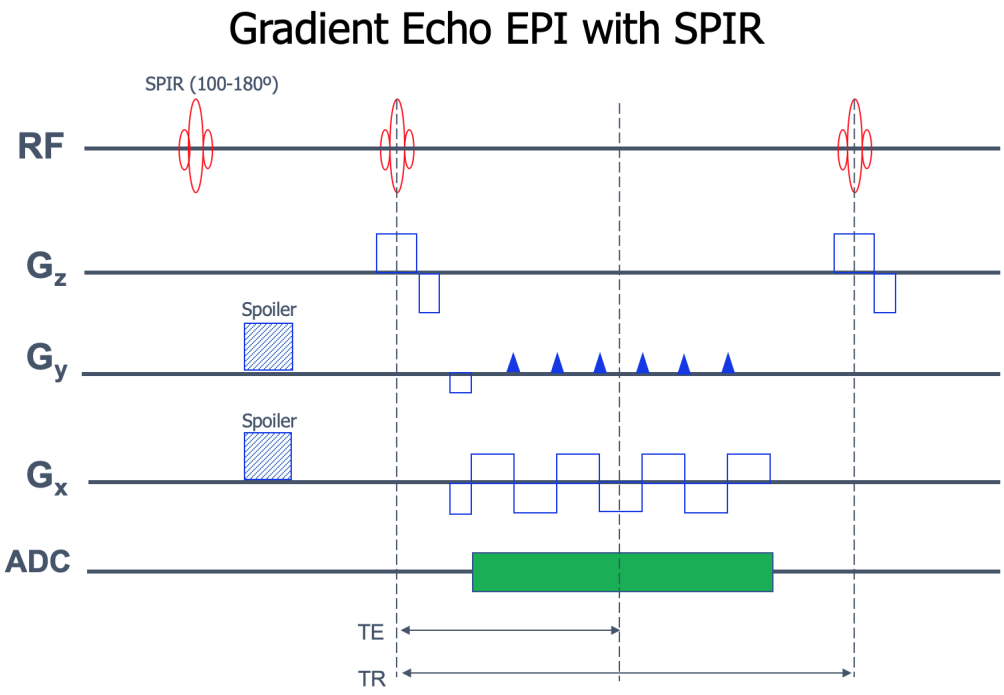

Figure 2.11: Sequence diagram for gradient echo EPI with spectral presaturation with inversion recovery pre-pulse that will suppress all fat signal in the tissue. The RF pulse can vary between 100-180°, which will invert the magnetization along B_0 and, at the same time, tip it partially into the transverse plane. The transverse signal is spoiled by G_y and G_x. The main RF pulse is timed such that the fat signal reaches zero mark along the longitudonal plane as it recovers its magnetization.

(Bydder et al. 1985). STIR uses 180° pulses to reverse the magnetization in the longitudinal plane. As longitudinal signal from fat recovers, an application of the main RF pulse for gradient echo signal can be timed such that it meets the fat signal as it crosses the zero component in the longitudinal axis.

However, both CHESS and STIR require ideal homogeneity in the B1 field, for a 90° or 180° flip angle respectively, which is hard to achieve in practice due to RF inhomogeneities. As a result, additional transverse magnetization may contaminate the signal during the main imaging sequence. Therefore, a preferred technique is a hybrid combination of both STIR and CHESS, called Spectral Presaturation with Inversion Recovery (SPIR) (Oh et al. 1988). SPIR applies a RF pulse within a range of 100-180° that inverts the signal and creates a transverse component at the same time (see Figure 2.11). The residual magnetization from fat can recover until it crosses the zero point along the longitudonal axis, like in STIR, which is when the RF pulse of

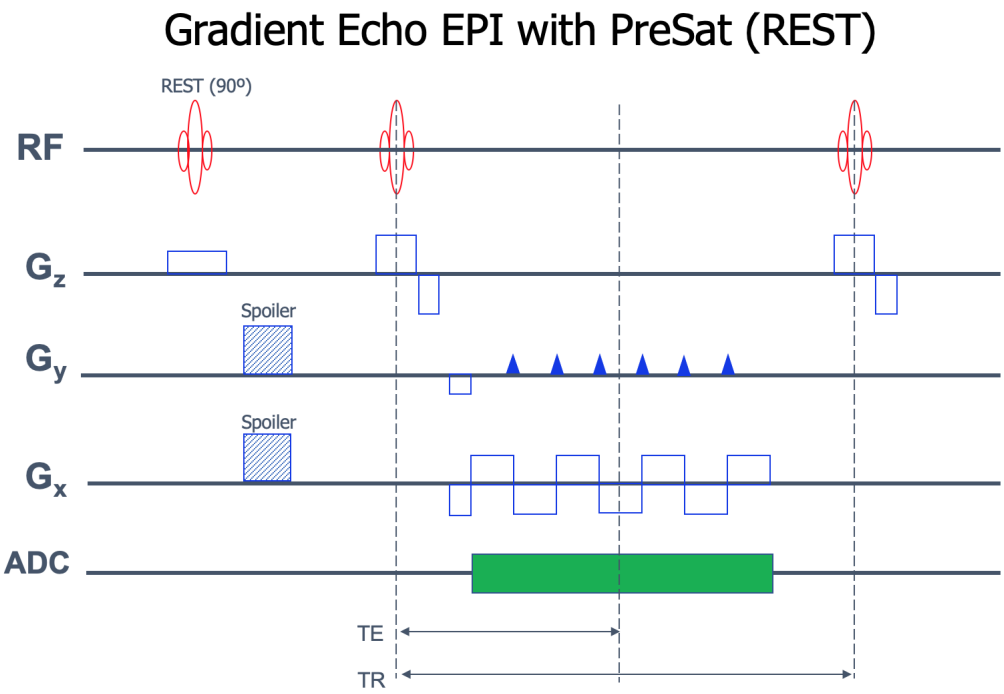

Figure 2.12: Sequence diagram for gradient echo EPI with spatial saturation preparatory pulse that will suppress the signal from a selected cross-section in the through-slice direction. Large spoiler gradients are turned on following the 90° RF pulse in order to destroy all transverse magnetization.

the main imaging sequence is applied. Just prior to that, any transverse magnetization in the RF field is spoiled via a set of strong gradients.

Finally, in some circumstances it may be beneficial to consider complete spatial suppression of signal in certain regions of the anatomy. For example, as pregnancy leads to substantial deposition of adipose tissue in the visceral compartments of the maternal tissue (Straughen et al. 2013), it creates extra demands in the effectiveness of fat suppression pulses, particularly for larger patients. Furthermore, the FOV that must be covered is significantly larger than a standard abdominal scan, which creates limitations around the resolution that is practically achievable. To address this, spatial saturation pulses may be used, also known as Regional Saturation Technique (REST), which allow suppression of signal from entire regions of tissue, rather than just fat (Felmlee et al. 1987). REST uses the same principle as Fat-Sat, except that the RF pulse is applied to the entire frequency spectrum that covers both fat and water based tissues (see Figure 2.12.

2.7.2 Distortion

Geometric distortions are caused by offsets in the phase of spins which shifts their reconstructed location away from true spatial coordinates in the reconstructed image (Jezzard et al. 1999). In EPI, distortions appear primarily in the phase encoding direction due to the slower sampling speed. The spatial shift in phase encoding direction follows the same calculation principle as in (2.29) except that the frequency offset component is dependent on a spatially varying magnetic field inhomogeneity $\Delta B0$ (Poynton et al. 2008):

$$\delta y_d = \frac{\gamma \Delta B_0(x,y,z)}{2\pi} FOV_y \tau_{pe} \tag{2.30}$$

While geometric distortions do not hinder the accuracy of T2* decay constant estimations, they affect image registration accuracy. In particular, geometric distortion correction is often necessary when registering a T2 weighted image to a T2* weighted image to perform group analysis or segmentation of brain structures. T2 weighted images are typically acquired with a single shot fast spin echo EPI sequence in fetal imaging, which is not prone to geometric distortions. As a result, T2 weighted images contain accurate structural information and are a preferred image contrast in fetal brain segmentation methods (Gholipour et al. 2017). Geometric distortions can be minimized with improved shimming techniques. Many 3T MR systems today are tooled with second order shim coils that improve the general homogeneity of the scanner and offer advanced performance. However, single shot EPI scans are still rife with distortions. Distortions can be corrected by estimating ΔB_0. The most common method for doing this is to acquire a field map, which is a separate low resolution GRE acquisition sequence that acquires a dual-echo image with two closely matching TE values (Jezzard et al. 1995) (see Figure 2.13). The change in phase, $\angle$, between the two generated images is proportional to the field inhomogeneity value:

$$\gamma \Delta B_0 = \frac{\angle I_2 - \angle I_1}{TE_2 - TE_1} \tag{2.31}$$

While it may be possible to calculate phase displacements directly from multi echo single shot GRE EPI sequence, and thereby remove the need for an additional field map acquisition that

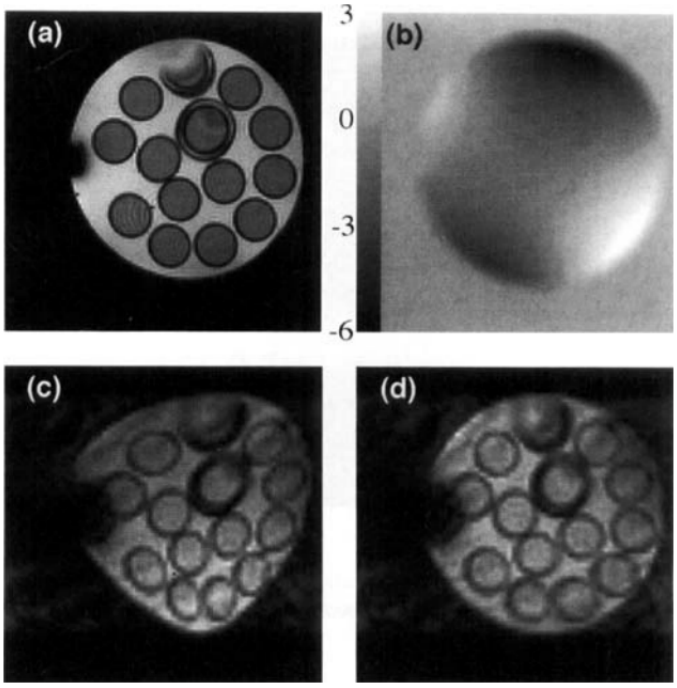

Figure 2.13: Example of EPI distortion correction by a way of a field map as published in the classic paper by Jezzard et al. (1995). a) shows the image acquired with a conventional gradient echo b) a map of field inhomogeneity, as described in (2.31). The numbers alongside a scale indicate the pixel shift c) gradient echo EPI acquisition, which resulted in significant geometric distortion and susceptibility signal loss d) geometric distortion is corrected via a field map

may prolong the overall scan time, there are typically a few challenges with this approach. Firstly, the difference in TE must be of a small order, which is often not feasible if the FOV is large. Secondly, in many MR systems, a field map acquisition will often perform automatic re-shimming of the B_0 field, thereby leading to better homogeneity to start with prior to the main scan.

2.7.3 Nyquist ghosts

Nyquist ghost artifacts appear as a uniform copy of the image which is shifted by half the field of view in the phase encoding direction (see Figure 2.14). They arise exclusively during EPI acquisition due to the irregularities of the k-space sampling that cause a phase offset in the reconstructed data. A common EPI sequence traverses the k-space sequentially in the backwards and forwards directions along the frequency encoding axis, while phase of the spins is incremented via gradient blips in the phase encoding direction. During image reconstruction,

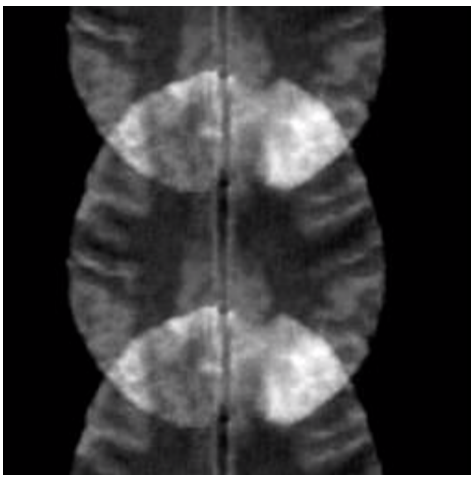

Figure 2.14: An example of an EPI nyquist ghost that's manifested as a repeating pattern shifted by half the FOV. These artifacts are most commonly caused by eddy currents generated due to fast switching of the gradients in EPI sequences, which in turn causes inconsistencies between the odd and even k-space readouts. Courtesy of Chung (2019)

the backwards passes through the k space are reversed. If there were any irregularities in the gradient waveforms or their timing, the echo formation in the reversed backwards pass is not a symmetrical version of itself. As a result, a Nyquist ghost is formed in the reconstructed Fourier transformed image as its shifted with $FOV_y/2$ pixels relative to the center of the image in the phase encoding direction. Nyquists ghosts are most commonly induced by off-resonance effects. The most frequent stimuli for the local variations are eddy currents that are generated through fast application of the gradients during the EPI readout and add phase shifts to the spins in the local magnetic field.

The best strategy to reduce Nyquist ghosts is to calibrate the field homogeneity through a shimming sequence prior to the scan of interest and keep the acquisition matrix size in the phase encoding direction as small as possible.

2.7.4 Aliasing

Anatomy that lies outside of the FOV produces a signal that is detected by the receiver coils. Due to the circular nature of phase, this signal cannot be distinguished from its conjugate that

lies within the FOV. Aliasing will occur if:

$$A_y \geq \frac{2\pi}{\Delta k_y} \tag{2.32}$$

where A_y is the width of the anatomy along the phase encoding axis. If aliasing cannot be avoided due to limitations in the FOV, the phase encoding direction may be switched to avoid the wrap around on the anatomy of choice. This can occur with larger patients in the scanner during fetal imaging. Alternatively, in order to reduce aliasing, REST slabs can be used to suppress the signal completely that is outside of the field of view. As described in the previous section 2.7.1, Regional Saturation Technique (REST) prolongs the scanning time but allows alias free imaging. Finally, many scanners also have the option of 'fold over suppression', which scales FOV_y and N_y linearly in the phase encoding direction at the expense of the prolonged scanning time.

2.7.5 Bias Field

Bias field artifacts are low frequency drifts in signal intensities over the spatial domain that typically arise due to non-homogeneous B1 fields (see Figure 2.15). Bias field artifacts are not always immediately evident from visual inspection alone but can cause a significant difficulty when performing intensity based post-processing techniques such as registration, classification and statistical parametric mapping. The accuracy of intensity based post-processing domains such as image registration and segmentation suffer significantly from bias field and require its correction. Bias field correction can be performed retrospectively via low frequency filtering, histogram matching or surface fitting based methods (Song et al. 2017). One of the most popular algorithms is the intensity based N4 algorithm that is available under the Insight Toolkit software (Tustison et al. 2010). A few notable fetal MR studies that investigated slice to volume reconstruction methods proposed comprehensive frameworks for estimation and correction of bias fields in each slice to improve iterative registration schemes (Kim et al. 2011; Kuklisova-Murgasova et al. 2012). We review these in the context of slice to volume reconstruction in section 3.7.

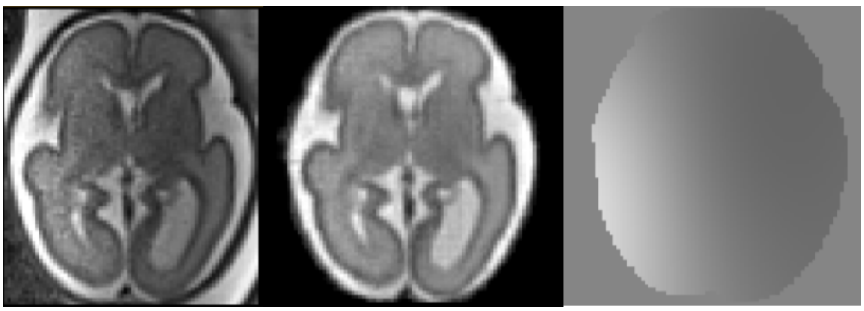

Figure 2.15: Example of a single slice from a fetal brain slice stack acquired in the transverse orientation - a) with bias field uncorrected b) bias field corrected and 3D reconstructed c) smoothly changing bias field calculated for the given slice a). Courtesy of Tourbier et al. (2018)

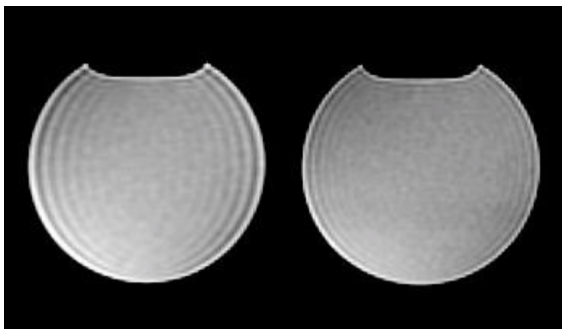

Figure 2.16: A phantom scanned at two different k-space matrix sizes resulting in varying Gibbs' artifact. Left: Substantial artifacts appear as a result of the coarse sampling of the signal (small matrix size). Right: improved version of the same scan due to a larger matrix size. Courtesy of Chung (2019)

2.7.6 Gibbs' artifact

Gibb's artefacts appear as fine parallel lines near the edges of high contrast interfaces and fade quickly from the boundaries (see Figure 2.16). Their cause is given as the finite rise of gradients, leading to under sampling in high spatial frequencies at the edges of k-space, which are responsible for fine details of the image (L. F. Czervionke et al. 1988). These artifacts are normally considered to be harmless in qualitative evaluation of fetal images as long as they are recognized as artifacts (Levine et al. n.d.). If Gibbs' artifact is persistent, one of the most common ways to treat it is by reducing the FOV, while keeping the number of Ny the same, or increasing the phase encoding lines. Gibbs' artefacts can also be reduced in the post-processing stage by applying a Hanning window filter.

2.8 Patient comfort

Ensuring the maximum possible comfort to the patient during the examination can be one of the most effective methods for motion artifact prevention (Malamateniou et al. 2013). For neonatal imaging, newborns are scanned with a team that consists of a pediatrician, a radiologist and pediatric nurse staff (Pennock 1992). Newborns are normally prepared for the examination in a separate room with minimum disruptions and noise. Sedated or unsedated newborns are equipped with dedicated earmuffs to reduce noise exposure (Malamateniou et al. 2013). The newborn is dressed and wrapped in plenty of comfortable linen and clothing to ensure soft padding around them and provide the necessary warmth while in the scanner room. To ensure the best possible SNR, a dedicated neonatal coil should be used during such an examination. It is not uncommon for babies to wake up during an examination, particularly during exposure to GRE EPI sequences. The team of specialists may then attempt to put the newborn back to sleep, which may require considerable amount of time. Therefore, it is paramount that the acoustic noise is reduced as much as possible for GRE EPI sequences in particular by reducing the slew rates of the gradient coils (Malamateniou et al. 2013).

For fetal imaging, the mother is typically advised not eat anything for 4 hours before the scan to minimize the motion of ingested food within the gastrointestinal tract, which may cause signal susceptibility artifacts due to air boluses that are trapped in the gut (Saleem 2014). Patients are normally laying in the scanner in the supine position, or sideways, depending on their most comfortable position(Saleem 2014). The scanning is done with a phased-array torso surface coil. Frequent communication with the patient is established throughout the scan to make sure that they have enough comfort. If necessary, certain patients are offered a break during the scan to ensure that they restore blood circulation in the limbs and unwind. This often helps to prevent motion due to discomfort. If the patient does need to move during the procedure itself, recalibration is performed to ensure homogeneity of the B_0 field.

Finally, examinations for neonatal and fetal imaging are normally limited to one hour scanning time. The majority of the time is dedicated to the clinical portion of the scan, which leaves a window of 10-15 minutes for an acquisition of a research based investigation. As a result, all

investigations bear significant constraints on time that is available for scans, which has to be taken into account when designing protocols for research studies.

Chapter 3

Reconstruction

This chapter presents the key aspects of post processing methods that are used for reconstruction of T2* relaxation maps. A discussion of optimal sampling rate and range for voxel based T2* estimates is first presented. It is followed by presentation of common techniques for estimation and correction of relaxation values in the presence of noise. Finally, we introduce the reader to the methods of slice to volume reconstruction (SVR) in the presence of motion in fetal imaging.

3.1 Optimal sampling rate and range

To obtain an ideal estimate of T2* relaxation rate one should aim to measure the T2* weighted signal along the entire decay curve (Cohen-Adad 2014). Typical sampling in adult brain T2* relaxometry in literature ranges between 4-8 echoes depending on the protocol (Deistung et al. 2013; Gati et al. 1997; W. Lee et al. 2012; T.-Q. Li et al. 2006; Peran et al. 2007; Rivkin et al. 2004; Rodrigue et al. 2011; Yao et al. 2009), with some studies measuring up to 64 echoes (Peters et al. 2007). An acquisition protocol that is highly accelerated will contain smaller SNR, and may suffer from other irregularities such as off-resonance effects or sampling errors. As a result, a finer sampling of the signal decay is advised. An earlier work by Peran et al. (2007) performed a quantitative evaluation of the optimal sampling rate for accurate estimation of T2* with a multishot GRE EPI sequence in adult subjects and found that a sample of 6 echoes gives

sufficient sampling accuracy.

When considering the upper and lower bound of TE values that need to be captured, it is important to note variation of values across different regions of interest. For example, typical T2* values for frontal white matter are statistically higher than those of deep gray matter structures. Therefore, a suitable range of values must be captured to accurately measure both. In general, as a rule of thumb, the lower and upper bounds of T2* relaxation must first be identified throughout the regions that are to be measured. Then, one would sample at least at half the lower bound and at double of the highest bound to cover the sufficient sampling range (Cohen-Adad 2014). More importantly, the estimates would be highly effected by noise levels. If the SNR in the later sampled echoes is small, substantial errors may be made in the estimates. To avoid this, denser sampling is advised. Furthermore, considerations must be taken for clinical subjects. For example, certain pathologies are associated with increased iron concentrations, which may lead to fast signal dephasing (Hagemeier et al. 2013; Mitsumori et al. 2009). Therefore, the sampled T2* signal should have a lower effective bound in such cases.

Further difficulties may arise when measuring T2* in areas of the brain that are subject to strong susceptibility induced signal losses. A standard formulation for estimation of through slice susceptibility induced gradients uses a cardinal sine approximation and is estimated iteratively (Fernandez-Seara et al. 2000). Sparse sampling of the decay curve may substantially hinder the accuracy of the results in this situation due to the cyclic nature of the cardinal sine. Therefore, it is generally desirable to sample as densely as possible if strong susceptibility induced losses are expected.

In summary, estimation of exponential decay constants leads to a "chicken and egg problem", such that an optimal sampling rate and range parameters require one to know the T2* decay value itself. Therefore, typically, a pilot experiment is necessary to obtain the optimal sampling strategy. Furthermore, we note that while adult T2* values exhibit very narrow range of normative values, the first reports of newborn T2* values by W. Lee et al. (2012) and Rivkin et al. (2004) had shown significant variation of relaxation rates between subjects. The variance was concluded to be attributed to intense structural and functional changes of the brain during

the early maturation period. Therefore, a robust estimate of fetal and neonatal T2* should include as dense and as wide sampling as possible, in order to safeguard against the large variance in values from subject to subject.

3.2 T2* estimation algorithms

Once the suscepibility induced effects, which we described in section 2.5.5, are accounted for and corrected, the signal should follow a smooth monoexponential decay model. A number of algorithms can be used to estimate the decay constant T2*. The simplest way to formulate the solution is by transforming the decay into a log-linear plot and fitting a straight line through the measured signal points.

This can be achieved via a simple ordinary least squares model (OLS). Nevertheless, many research studies would typically use the OLS estimate as the initial value for T2* relaxation rate through which the next step is achieved (Cohen-Adad 2014). This is because OLS does not account for changing SNR levels through the sampled points. Signals sampled at a later TE have a significantly smaller SNR due to dephasing effects and should therefore be downweighted in their contribution to the estimates. The latter can be achieved via the generalized least squares (GLS) model, which treats the residuals with a varying degree of correlation to the model. The best formulation of GLS for exponential data would be the weighted least squares, which assumes non equal variance in the observed entries.

The most common preferred method for fitting residuals to an exponential model is the Levenberg Marquadt algorithm (Levenberg 1944), which belongs to non linear gradient based methods. For $F(x)$ which is a function of independent variables of x, the minimum is found at $\Delta F(x) = 0$. The solution can be found by following the descent of the gradient $-\Delta F(x)$. Through iterative search with varying step sizes, the optimal value for T2* and S0 parameters can be found. A more detailed implementation of this is shown in section 4.2.4.

3.3 Truncation, baseline and offset models

In multi-echo MR acquisitions, small voxel size, which leads to decreased SNR, has a particularly detrimental effect when considering the signals measured at long echo times. The presence of low SNR signal samples in a set of multi-echo measurements manifests in consistent bias and increased variance in the estimates of the exponential decay rate. To increase the SNR, it is possible to reduce the bandwidth per pixel or increase the number of signal averages. Nevertheless, the former is limited by the desirable resolution and the latter increases scan time, both of which are not ideal. Due to the inherent noise component, a monoexponential signal fit does not have a zero floor in MRI data. A number of approaches had been previously proposed to address T2* estimation in the conditions of low SNR (Feng et al. 2013, 2014; T. He et al. 2008, 2013). Truncation, offset and baseline are simple post-processing techniques that could improve the estimates in low SNR (T. He et al. 2008).

The offset model considers noise and other effects as an additional constant in the exponential fit equation, which is normally estimated by calculating the background noise of the scanner:

$$S(TE_i) = S_0 e^{-TE_i/T_2^*} + C \tag{3.1}$$

Baseline model removes the noise component from the acquired signal, instead of adding it to the estimate:

$$S = \mu(S_m) - \mu(S_n) \tag{3.2}$$

where S_m is the measured signal, S_n is the noise and μ is the average operator. To obtain the measure of the S_n, signal is typically measured in the background of the generated data where no tissue is present in the image.

Finally, truncation technique relies on removal of data points at the end of the TE train that do not meet the necessary threshold for SNR. As such, at every iteration of the exponential model estimate, a pearson correlation coefficient (PCC) is calculated for the model. If the PCC is lower than a certain amount, e.g. 0.995, than the last data point is removed and the fit is recalculated. This is repeated until the difference between recalculated T2* decay points is less

than a certain empirically chosen threshold (T. He et al. 2008). Generally, truncation model produces better accuracy and repeatability for estimation of T2* in the presence of noise and low SNR than the baseline and offset models (Feng et al. 2014; T. He et al. 2008).

Noise corrected models can incorporate the fitting of the squared signal to the second moment, for rician non-centrally distributed noise (Feng et al. 2013), which yields better performance than the truncation method at very low T2* values:

$$S_m^2 - 2C\sigma^2 = S_0^2 e^{-2TE/T_2^*} \tag{3.3}$$

where C is the number of receiver coils and σ^2 is the measured variance of the noise.

3.4 Spatial smoothing models

Another type of model that can yield superior accuracy for T2* estimation in very low SNR conditions are spatial smoothing models. The basic principle is to fit the data with a smoothing parameter, which convolves the signal with a spatial filter, thereby removing the effect of noise by using knowledge of the pixel neighourhood (Feng et al. 2014). For a pixel x_i, spatially filtered output is given as:

$$\bar{S}(x_i) = \sum_{x_j \in W_i} w(x_i, x_j) S(x_j) \tag{3.4}$$

where W_i is the spatial window over the current pixel x_i and $w(x_i, x_j)$ is the weight function over the intensity $S(x_j)$ of the neighbourhood pixel x_j. The most commonly used filters are Gaussian low-pass (LPG) and non-local means (NLM) filter. While the former introduces unwanted blurring into the data, which then has to be removed, NLM has been shown to constistently outperform LPG (Feng et al. 2014). In particular, its effects work well in presence of Rician non centrally distributed noise in the MRI signals.

Finally, pixel-wise curve fitting with adaptive neighborhood regularization (PCANR) is a more advanced algorithm that uses adaptive regularization according to inter pixel signal similarity (C. Wang et al. 2018). One can assume that pixels with similar intensity values at different TEs

would have similar decay values. As a consequence, during the fitting procedure, a regularization term is added that considers the signal decay in neighbouring pixels. The regularization term is tuned according to noise and proximity to the given voxel. Such approach produces T2* maps with less root mean squared error, particularly in noisy conditions. The downside of PCANR are long processing times which can hinder computation.

Overall, the models that we described in this section and section 3.3 are most often used for T2* estimation under very low SNR conditions. Most commonly, these are the T2* estimates obtained in studies of iron-overload for diagnosis of thalassemia in cardiac (T. He et al. 2008) and liver tissue (C. Wang et al. 2018). While these models have not found wide adoption in the study of relaxometry of the brain, we note that under certain conditions it may be possible to achieve highly accelerated T2* mapping of the brain that will yield very low SNR images. In such circumstances, these kinds of methods may allow for accurate fitting strategies.

3.5 Region based relaxometry

The most common way to perform estimation of relaxation rates across a region of interest relies on calculation of exponential decay curve for each voxel individually followed by a calculation of the mean across all voxels for that region, which is termed 'voxel based relaxometry'. Marro et al. (2011) introduced an alternative approach which proposed to average the signal intensities from all voxels for each echo first, prior to T2* estimation. The study demonstrated and tested this approach to show the benefit of such calculation for liver T2* estimation and argued that it yields more accurate and stable measurements for homogeneous regions such as the liver.

3.6 Multi-exponential models

Although most transverse relaxometry studies consider mono-exponential models of signal decay, the signal contributions can be split into multiple components. Protons in water molecules that are trapped in the myelin sheaths undergo spin dephasing at different rate to that of protons of

the intra- and extra- cellular water (Deoni et al. 2012). As a result, white matter regions can modelled with a multi-exponential transverse relaxation model (Hwang et al. 2010b):

$$S(\text{TE}) = S_{0,m}e^{-TE/T^*_{2,m}} + S_{0,a}e^{-TE/T^*_{2,a}} + S_{0,w}e^{-TE/T^*_{2,w}} + C \tag{3.5}$$

where m, a, w subscripts represent myelin, axon and mixed water compartments respectively, and C represents the residual baseline signal. Nevertheless, measurement of the multi-exponential model parameters is challenging even in adult subjects as very fine sampling of the relaxation curve is required to estimate the signals from all three compartments. Typically, 18-64 echoes are needed to fit such models (Hwang et al. 2010a) (Van Gelderen et al. 2012).Furthermore, due to the extra short decay values of the first component, myelin water fraction, TE values in the order of single digits are necessary. Moreover, a very high SNR is also needed with noise levels of 1% (Graham et al. 1996). As a consequence, we note that such a model, at present, would present a significant challenge within the fetal brain imaging domain, which would not be achievable with current accelerated imaging methods in the presence of motion.

3.7 Slice to volume reconstruction

To address inter-slice motion in fetal MRI, a group of methods have been developed to prospectively reconstruct a 3D volume from individual 2D slices that are scattered within the acquired image space (see Figure 3.1). These techniques are referred to as slice to volume reconstruction (SVR) methods. The approach requires dense oversampling of the target area of interest to ensure sufficient coverage of the brain (Ferrante et al. 2017). Typically, 6 to 8 sets of repeated acquisitions in at least 2 orthogonal views of the brain are acquired for this purpose. Each acquisition results in a set of slices that are termed as a slice stack.

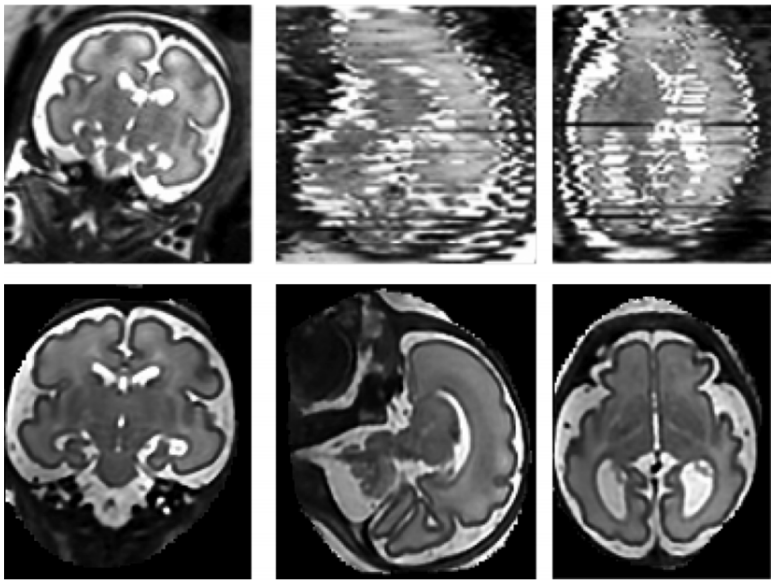

Figure 3.1: Top row: an example of a single slice stack acquired in the coronal plane. Significant motion is visible in the transverse and sagittal planes. Bottom row: High resolution slice to volume reconstruction of the fetal brain in 3D. Reconstruction is made possible by combining multiple slice stacks in three orthogonal planes, which yields a dense oversampling of the image space. Courtesy of Kainz et al. (2015).

3.7.1 Slice alignment

Formally, the slice alignment problem is defined as a set of K slices $s_{1,n}...s_{K,n}$ acquired N number of times. A given slice k belongs to a slice stack $\Sigma_{k=1}^{K} s_{k,n}$, where n corresponds to the number of the repeated acquisition. The position of $s_{k,n}$ is known relative to the scanner coordinate system Ω_s but unknown with respect to the anatomy of the 3D brain structure Ω_a, which is moving relative to Ω_s. The goal of the alignment is to find a rigid transformation matrix, T, which consists of 3 translation and 3 rotation elements, such that it allows one to map $s_{k,n}$ from Ω_s to Ω_a. There are three distinct approaches to tackle this first initial step of slice alignment.

The first approach utilizes a target reference volume, V_{ref}. The first step requires construction of the initial target reference volume, $V_{ref,0}$, either from a stack of slices with the least motion, $\sum_{k=1}^{K} s_{k,min}$, or from an average of all the slice stacks combined, $\sum_{n=1}^{N}\sum_{k=1}^{K} s_{k,n}/N$. $V_{ref,0}$ is constructed at the lowest resolution to ensure the widest possible range of effective rotation and translation estimates during the first iteration. Next, each individual slice stack $\sum_{k=1}^{K} s_{k,n}$

is registered to $V_{ref,0}$. After this, an updated reference, $V_{ref,1}$, is constructed by using the information from the newly aligned positions of the slice stacks that produced a mapping, $T_{k,n}(s_{k,n})$, to $V_{ref,0}$. $V_{ref,1}$ is constructed at a higher resolution, to allow for finer registration accuracy for the next iteration. Simultaneously, all of the slice stacks are split in half, $\sum_{k=1}^{K/2} s_{k,n}$. The target update and stack splitting process is repeated until the final registration is performed at a single slice level. In other words, the registration is performed as individual slices that are registered to $V_{ref,final}$.

The second approach uses the intersection of individual slices acquired in orthogonal planes (Kim et al. 2009). Orthogonal pairs of slices can be matched by minimizing an energy function to find the optimal relative position between them. This is performed for all slices such that a global minimum is reached.

The third approach uses convolutional neural networks (CNNs) to learn a canonical representation of a spatially consistent 3D brain space - an atlas. The rigid transformation matrix can be written as $\hat{T}_{k,n} = \Psi(s_{k,n}, \Theta)$, where Θ are the learned parameters of the CNN. Note that each $\hat{T}_{k,n,}$ is only an approximation as it maps the slices to the canonical atlas space Ω_{atlas}. The next step requires a mapping to be found from Ω_{atlas} to Ω_a, so to align the slices to the individual 3D brain anatomy of a given subject. Therefore, the CNN based approach is used an initializer and is typically followed by iterative registration scheme, which was described in the first approach above.

3.7.2 Volume reconstruction

Once the correct spatial alignment of $s_{k,n}$ in Ω_a has been determined by $T_{k,n}(s_{k,n})$, the data must be interpolated to form a single 3D volume. At this point, individual voxels from $T_{k,n}(s_{k,n})$ are irregularly scattered in Ω_a. To determine the intensities of voxels that are placed on a regular grid in Ω_a, a super resolution reconstruction is performed. Gholipour et al. (2010) presents the forward model with the added knowledge of the slice acquisition:

$$y_{k,n} = DBST_{k,n}x + \epsilon_n \tag{3.6}$$

where $y_{k,n}$ is the vector of voxels formed from $s_{k,n}$, x is the vector of ground truth voxels representing the high resolution 3D brain image, D is the downsampling matrix, B is the point-spread function matrix representing the MR signal acquisition model in x-y plane, S is the matrix that represents the slice acquisition model along z-axis as described in (2.15), and ϵ_n is the noise. Given the knowledge of $D, B, S, T_{k,n}$, the high resolution 3D brain image can be computed via maximum likelihood estimation, which is equivalent to minimization of the sum of the $l2$-norm:

$$\hat{x}_{MLE} = \arg\min_{x} \sum_{k=1}^{n} \|DBST_{k,n}\hat{x} - y_{k,n}\|_2^2 \tag{3.7}$$

where $\hat{x}$ is an estimate of the reconstructed volume.

While the above represents a basic approach to super resolution reconstruction, there are a number of other considerations that must be taken into account when performing this step. Therefore, below, we present the reader with a brief overview of the evolution of SVR methods to give a better understanding of how these components and considerations evolved and developed over time.

3.7.3 Evolution of SVR methods

The first SVR methods were presented by Rousseau et al. (2006) and (Jiang et al. 2007). Rousseau et al. (2006) acquisition was based on thick slices, which resulted in substantial blurring in the final reconstructed high resolution volume. The point spread function (PSF) was based on a weighted sum of Gaussian kernels. This kind of PSF model is typically used in light based imaging techniques, such as light microscopy or photography. Jiang et al. (2007) proposed thin slice reconstruction scheme with 1mm thickness. Interpolation for each target reference update was performed with multi-level B splines for Jiang et al. (2007), giving superior performance to that of Rousseau et al. (2006). At a later date, Gholipour et al. (2010) introduced a model that incorporated the knowledge of acquisition model. Importantly, the slice profile was modelled by a cardinal sine windowing function rather than a Gaussian kernel. The cardinal sine represents the true representation of MR sampling in the through-slice direction by the slice

select gradients and is therefore a better approximation for the PSF. Furthermore, Gholipour et al. (2010) used super resolution reconstruction for interpolating between the scattered data points, which allowed one to resolve at sub sampled scale. The reconstruction was performed via maximum likelihood with a Tikhonov regularization term to further constrain the registration to minimize noise and registration errors. In 2011, Rousseau et al. (2010) presented an edge preserving technique that to further improve the super resolution reconstruction.

Kuklisova-Murgasova et al. (2012) combined the previous advancements presented in Gholipour et al. (2010) and Rousseau et al. (2010) into a robust framework. Notably, the earlier models required manual segmentation of the brain from the rest of the in-utero tissue. Furthermore, it is common that a very small subset of slices from a single shot EPI acquisition would be corrupted due to severe fetal motion. Kuklisova-Murgasova et al. (2012) incorporated an automatic brain segmentation method and an outlier removal method to tackle these issues. The framework also included bias correction to remove the slow frequency drift of intensities throughout the image that arises due to B1 field inhomogeneities Kim et al. (2011).

Kainz et al. (2015) presented a GPU accelerated procedure of SVR, that accelerated the reconstruction up to ten times in comparison to a conventional CPU based reconstruction. Other notable approaches, such as Alansary et al. (2017), introduced SVR methods for non rigidly moving structures to allow reconstruction beyond the fetal brain. Alansary et al. (2017) segmented the whole in utero image into multi-scale patches and reconstructed parts of the uterus of the fetal and mother's anatomy.

In 2017, Hou et al. (2017) addressed the challenge of the local minima during the slice to volume reconstruction. While prior methods, such as the work from Gholipour et al. (2010) and Kuklisova-Murgasova et al. (2012) introduced outlier rejection strategies, Hou et al. (2017) proposed a framework that would provide the initial alignment of slices in a common canonical 3D space that represents a fetal brain atlas with a single step mapping. The follow up work (Salehi et al. 2018), introduced a more accurate measure of misalignment between the slice oriented slice and its ground truth by using Special Euclidean group SE(3) representation of the coordinate system.

Chapter 4

Fast motion robust fetal T2* relaxometry at 1.5T

4.1 Introduction

The goal of this study was to directly quantify T2* in key brain regions in the fetus in utero and explore possible variations with gestational age, and compare these with previously reported values in preterm and term neonates at the same field strength. At the time of submission of this work, to the best of our knowledge, there have been no fetal brain T2* values reported to date. Knowledge of T2* may be useful for assessing brain development, brain abnormalities and for optimizing functional imaging studies.

Maternal respiration and unpredictable fetal motion mean that conventional multi-shot acquisition techniques used in adult T2* relaxometry studies were not practical. Single shot multi-echo echo planar imaging (EPI) was used as a rapid method for measuring fetal T2* by effectively freezing intra-slice motion.

T2* determined from a sample of 24 subjects correlated negatively with gestational age with mean values of 220ms (+/-45) for frontal white matter, 159ms (+/-32) for thalamic grey matter and 236ms (+/-45) for occipital white matter.

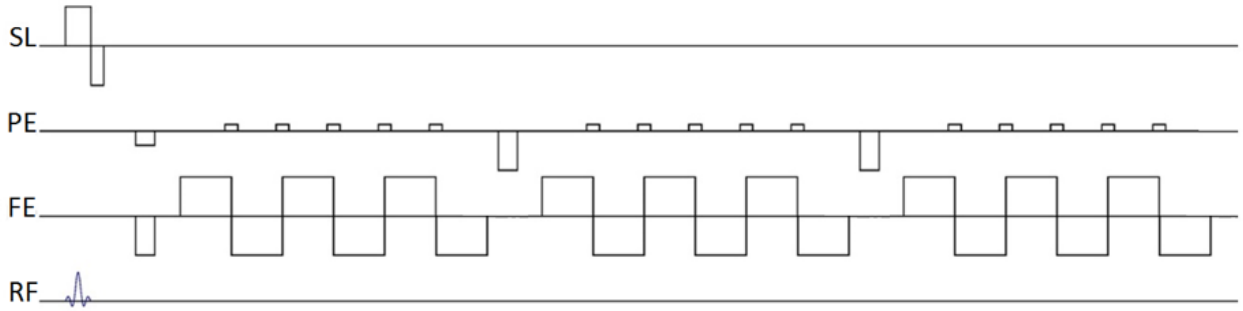

Figure 4.1: A schematic representation of a single-shot fast field echo multi-echo EPI sequence used in the study, which results in a set of spatially aligned T2*-weighted images at different TEs. SL - slice selective gradient, PE - phase encoding gradient, FE - frequency encoding gradient, RF - radio frequency pulse.

Our study concluded that fetal T2* values are higher than those previously reported for pre-term neonates and decline with a consistent trend across gestational age. The data suggests that longer than usual echo times or direct T2* measurement should be considered when performing fetal fMRI in order to reach optimal BOLD sensitivity.

4.2 Method

4.2.1 Acquisition

A simple calculation of brain T2* requires acquisition of two T2*-weighted images with different echo times (TE) that are aligned to each other and combined as a log ratio. Inclusion of additional TEs significantly improves the robustness of estimation and allows measurement errors to be ascertained. Collection of 6 separate T2* weighted images is thought sufficient for a reliable quantification of T2* decay rate in adults (Peran et al. 2007), but prolongs acquisition time considerably. To reduce scanning time, a multi-echo approach is often used to acquire data at all required echo times following a single radio frequency (RF) pulse excitation.

Sensitivity encoding (SENSE) (Pruessmann et al. 1999) and reduced flip angles with SPGR

sequences (Haase et al. 1986) allow faster acquisition. Further speed-up can be achieved by increasing the coverage of k-space per echo formed, using spirals (B. Wu et al. 2012) or phase encoding multiple lines (combining multiple echoes) with echo planar imaging (EPI) allowing complete images to be obtained in a single shot (Speck et al. 1998). Such methods trade acquisition speed of a slice with temporal resolution within the decay curve that needs to be sampled. For adult studies where the subject can remain still and brain T2* is typically 48ms for grey matter and 67ms for white matter at 1.5T (Siemonsen et al. 2008), multi-shot, multi echo methods have proved effective. For fetal relaxometry, motion makes multi-shot methods liable to gross artefacts and the available data from neonates suggest that T2* values are extended so that temporal resolution of sampling during the decay process is likely to be less critical. In view of these very different conditions, we have used single shot multi-echo gradient echo EPI to achieve multiple readouts of complete fetal brain slices following single RF pulse excitations (see Figure 4.1). This allows full T2* decay curves to be sampled for each imaged voxel from each shot. Figure 4.1 illustrates the basic diagram for this sequence. Acquiring multiple slices allows whole brain coverage.

Fetal imaging was performed on a 1.5T Philips Achieva scanner with a 32 channel cardiac/torso phased array coil for signal reception. Pilot experiments were performed to select scan parameters, to balance spatial and temporal resolution, signal to noise ratio and spatial distortion. The field of view (FOV) was set to 290x290 mm2 with fold-over suppression applied to provide full maternal abdomen coverage. Using an in-plane resolution of 3mm2 combined parallel imaging using a SENSE factor of 2 resulted in an echo train length of 51 and allowed full slice images to be encoded in 43.2ms. The sequence finally employed provides whole brain coverage, producing a set of 5 images per anatomical slice with an echo time range between 27-199.8ms, inter-echo spacing of 43.2ms and a TR of 10.7s. Fat suppression was achieved using the spectral pre-saturation with inversion recovery (SPIR) method (Oh et al. 1988). All scans were planned in the transverse orientation relative to the fetal brain. It may be helpful to note to the reader that a fetus can be oriented in any arbitrary plane with respect to the mother's abdomen, and therefore the transverse orientation of the fetal brain is often in a plane that is not orthogonal to the scanner coordinate system. For each subject the sensitivity profiles of

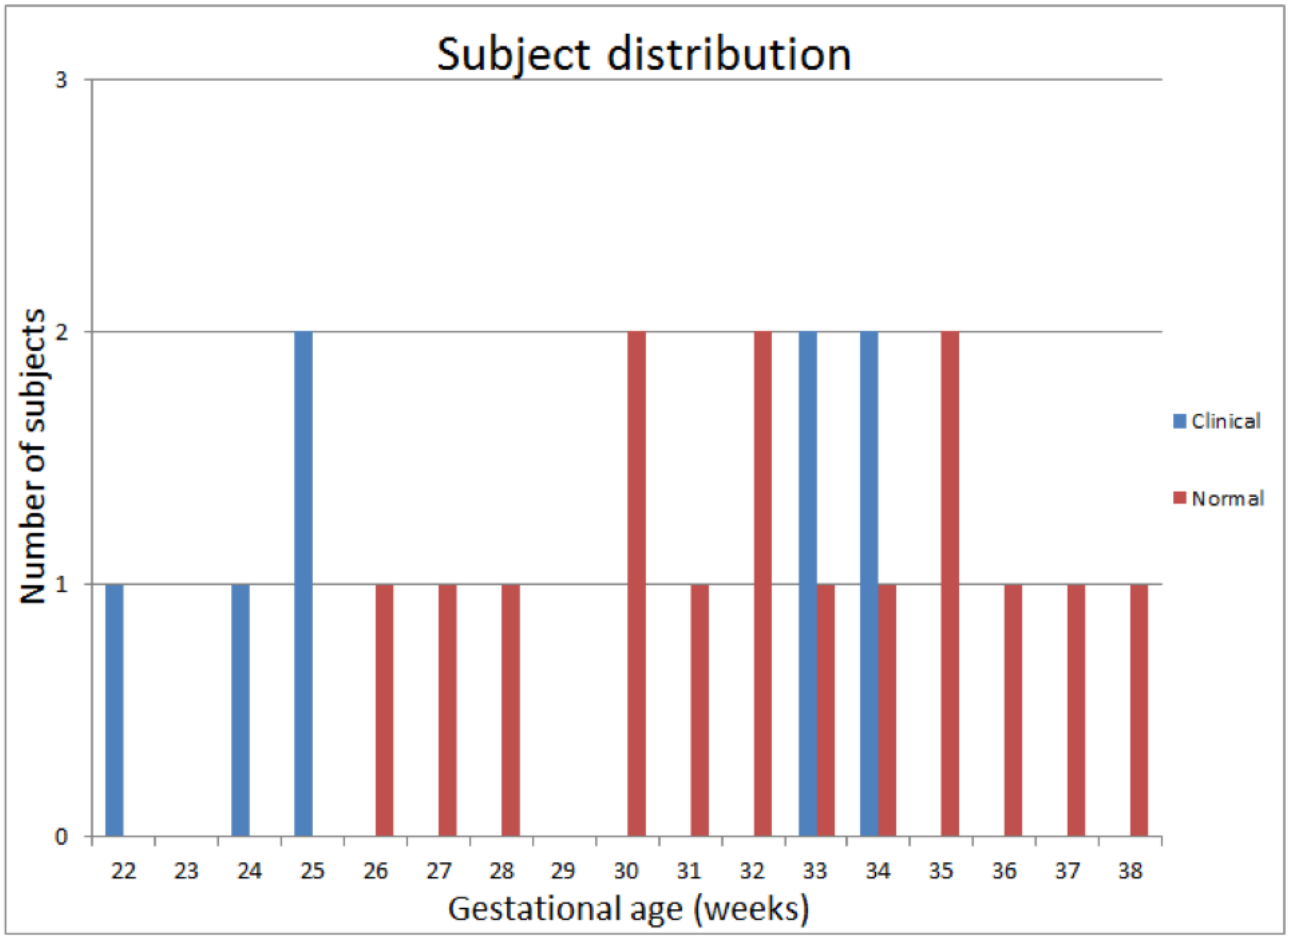

Figure 4.2: Distribution of scanned subjects across gestational age in whole weeks.

the phased array coils were measured immediately prior to the multi-echo EPI acquisition to ensure that optimal sensitivities were used for SENSE reconstruction and minimize the risk of large motion of the mother between the reference scan and main acquisition. In each scan 43 slices of 3mm thickness were acquired with no overlap. This provided sufficient coverage for the full brain anatomy of the fetuses from early to late gestational age, with a margin to allow for the chance of a change in fetal position between piloting the scan geometry and acquiring the data. Each whole brain acquisition took 10.7 seconds and was repeated multiple times for each subject. All subjects received a full structural clinical imaging examination consisting of single shot fast spin echo T2 weighted images (2) and snapshot inversion recovery (SNAPIR) T1 weighted images (Malamateniou et al. 2011).

Gestational Age (weeks+days)	Condition
22+1	Monochorionic diamniotic surviving twin (normal structural brain appearance at time of scan)
24+2	Absent septum cavum pellucidum
24+4	Transposition of great arteries Balanced chromosomal translocation 16/19 (normal structural brain appearance at time of scan)
24+4	Bilateral mild ventriculomegaly Down's syndrome
33+1	Maternal toxoplasmosis infection (normal structural brain appearance at time of scan)
33+2	Isolated mild ventriculomegaly
33+4	Isolated mild ventriculomegaly
34+3	Isolated mild ventriculomegaly
34+3	Bilateral mild ventriculomegaly Rounded dilatation of the anterior third ventricle

Table 4.1: Diagnoses of the clinical group of subjects

4.2.2 Subjects

23 fetuses (see Figure 4.2) with gestational ages ranging from 22 weeks to 38 weeks (mean 31.1, SD 4.3) were scanned. Eight of these subjects were diagnosed with a clinical condition, as verified by the clinical MRI component of the examination (see Table 4.1). This included two isolated mild ventriculomegaly cases, one bilateral mild ventriculomegaly with Down's syndrome and another bilateral mild ventriculomegaly with rounded dilation of the anterior third ventricle. Three subjects had normal structural appearance of the brain at the time of scan, but had other clinical conditions associated - transposition of great arteries with balanced chromosomal translocation, maternal toxoplasmosis infection and monochorionic diamnotic surviving twin pregnancy respectively. Finally, one of the scanned subjects had an absent septum cavum pellucidum. The remaining 15 subjects were confirmed to have no detectable abnormalities from the images obtained during their respective examinations.

The mothers were carefully positioned, with emphasis on patient safety, comfort and image quality. They were placed on the examination table with a mild lateral tilt, supported by additional padding, to avoid inferior vena cava syndrome. Maximizing comfort reduced the risk of gross maternal motion during the examinations. The anterior part of the coil was placed as close as possible to the known position of the fetal head to maximize signal to noise ratio (SNR).

The study protocol was approved by the local research ethics committee and informed consent was obtained prior to each examination. The data for this study was typically taken towards the end of the examination, the length of which should not exceed 60mins under the local scanning guidelines. To increase the data available for analysis, up to 7 repeats of the multi-echo EPI scans were taken per individual fetus, although because of the limited time allowed for a patient to be in the scanner, fewer repeats were acquired in some scans. No pharmacologic sedation was administered during scanning to any subject.

4.2.3 Regions of Interest

Three regions of interest (ROI) were manually selected in each individual scan for each subject with help of an expert fetal neuroradiologist - frontal white matter (FWM), thalamic grey matter (THA) and occipital white matter (OWM) (see Figure 4.3). These were chosen as they have robust definitions and are clearly identifiable throughout the age range, as well as for their clinical and developmental significance. At 22 weeks of age, for example, more subtle structures, commonly reported in adult studies, such as caudate and lentiform nuclei, would be difficult to select with sufficient accuracy. ROIs were selected on later echo images (at TE of 172.8ms or 199.8ms) as these provided better contrast between tissues (see Figure 4.3). Data that contained image artefacts within the fetal brain was excluded from the study. Fetal movement means that each stack of slices acquired may not provide continuous volume coverage. Regions were drawn individually on each slice, since neighbouring slices can be misaligned due to fetal motion between shots. The definition of the ROIs was based on the work of I. Gousias et al. (2012) to ensure robust definitions for each. By using upsampled images it was possible to accurately select patches within identifiable regions of interest with sub-pixel precision. The minimum size of any ROI was 9 interpolated voxels ensuring that at least one full native resolution voxel was sampled.

T2* measurements from multiple anatomical slices over multiple repeats (intra-session) were then averaged to provide a mean T2* value for each ROI.

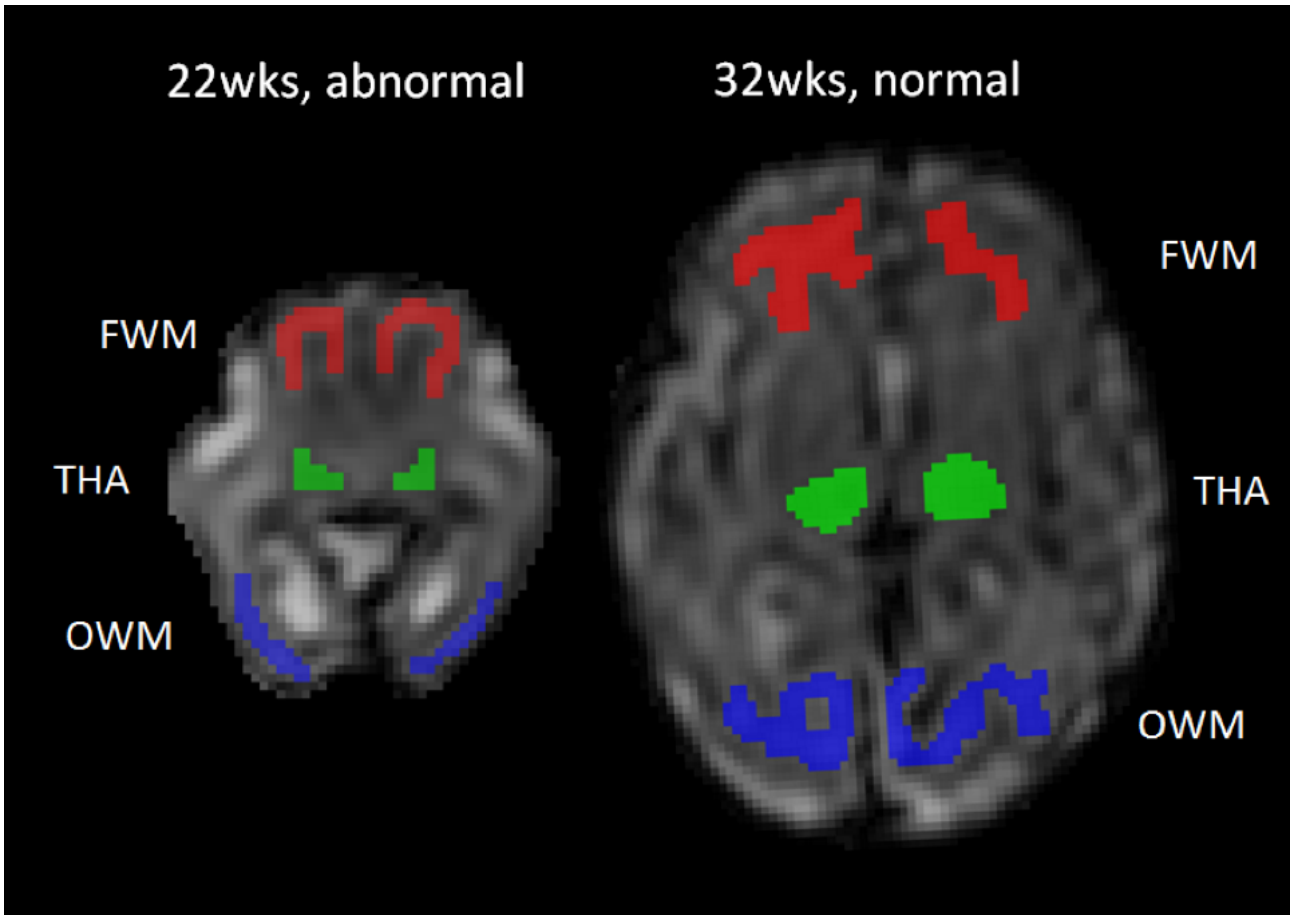

Figure 4.3: Measured regions of interest include: FWM - frontal white matter, THA - thalamic grey matter, OWM - occipital white matter. Cerebral appearance at gestational ages of a) 22 weeks and b) 32 weeks. Echo 4 (199.8ms) is shown in cases.

4.2.4 T2* Estimation

Single shot multi-echo EPI acquisition resulted in a set of 5 images each weighted at a different TE for each slice. For subjects with particularly large motion between the slices, small misalignments were occasionally observed between individual echo images of single slices. This was visually observed in approximately 5% of all slices. To ensure that corresponding brain voxels were aligned between individual echoes rigid body registration was performed in plane between the echoes, separately for each slice in all datasets. To align individual echoes, the image for echo 1 was first processed by manually extracting the brain and removing all maternal tissue. Echo images 2 to 5 were then rigidly registered to the brain extracted first echo with IRTK software (Studholme et al. 1999). B-spline interpolation was used to place the echoes on the same voxel grid. Since the time difference between the first and last echo for each slice was 173msec the changes in position were small in all cases, with maximum rotation of 2 degrees, and maximum in plane translations of 0.6mm and 0.4mm.

Once the echoes were registered to each other, voxel wise T2* constants were calculated by fitting a mono-exponential decay model to the measured intensities (S_i) and their respective echo times (TE_i):

$$S_i\left(TE_i\right) = S_0 e^{-\frac{TE_i}{T2^*}} \tag{4.1}$$

where S_0 is the signal at zero TE. To calculate initial estimates of T2* and S_0, ordinary least squares regression was used to fit a linear trend to the log of signal intensities versus TE for each voxel:

$$E = \sum_i \left(\ln\left(S_i - \frac{TE_i}{T2^*} + \ln\left(S_0\right) \right) \right)^2 \tag{4.2}$$

These estimates were then used to initialize a nonlinear iterative regression based on a standard Levenberg-Marquardt method (Levenberg 1944), with gradient (G), Hessian (H) and scaling factor λ:

Parameters	**Scan 1**	**Scan 2**
Echo train length	1	51
SENSE factor	1	2
Interecho spacing (ms)	8	43.2
Sampled echoes	14	3
Sampled range (ms)	5-109	27-113.4

Table 4.2: A conventional gradient echo sequence was used to measure T2* values in an adult volunteer and compared against an accelerated single shot EPI readout on the same subject.

$$\begin{pmatrix} S_0 \\ T2^* \end{pmatrix}_{i+1} = \begin{pmatrix} S_0 \\ T2^* \end{pmatrix}_{i} - (H + \lambda I)^{-1} G \tag{4.3}$$

A robust fit was implemented by applying Tukey's bi-square (32) as a weighting function. At each iteration weights are set based on values of the previous one. The process repeats until the weights converge.

A normalized root mean square error (nRMSE) was calculated to measure goodness of the exponential fit to raw data points. The normalization constant for each subject was calculated as the signal average over all echoes and repeats in the extracted brain only.

4.2.5 Protocol Verification

To verify the accuracy of the proposed method for acquisition and estimation of T2* values tests were performed on a healthy adult volunteer. The sequence described above was compared to a slower, conventional approach that is typically used to sample T2* decay times in adult subjects (Table 4.2). Acceleration of k-space readout was reduced by switching off EPI mode to reduce the echo train length from 51 (single shot readout) to 1. Spatial resolution, slice thickness and TR were kept the same as those used on the fetal acquisitions. FOV was reduced to 250x202 to cover only the anatomy of adult head. The T2* decay constant is much shorter in adults than in neonates at 1.5T (26), there is a more rapid loss of signal and so only 3 echoes (covering a range from 43.2ms to 113.4ms) were used to calculate T2* from the single shot acquisitions.

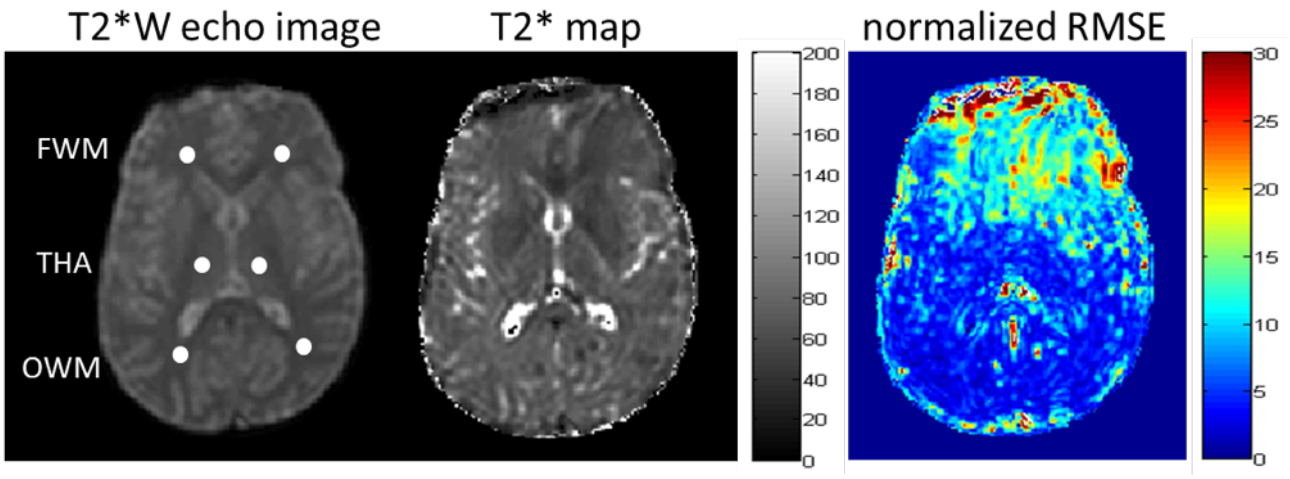

Figure 4.4: a) Definition of 6 ROIs measured in a healthy adult brain, marked on a single shot multi-echo GRE EPI scan b) Estimated T2* map (in ms) c) Estimated nRMSE (as percentage)

4.3 Results

4.3.1 Protocol Verification

The adult test data was sampled in regions of interest that matched those of the fetal study – thalamic grey matter, occipital white matter and frontal white matter (see Figure 4.4). Data was captured in both right and left hemisphere so that a total of 6 ROIs were sampled. Each ROI consisted of a region of 49 voxels covering an in-plane area of 441mm^2.

Figure 4.4 shows the first echo from a T2* weighted multi-echo scan (a), together with corresponding T2* map (b) and spatially varying nRMSE for the EPI data (c). The nRMSE is increased in frontal parts of the brain, which is as expected due to image distortion and signal loss in EPI associated with air in frontal sinuses. Figure 4.5 shows examples of T2* fits for a single voxel in a non-EPI and single-shot EPI scans in occipital white matter (nRMSE of 0.02), and compares these to a fit in a frontal white matter voxel (nRMSE of 0.09). Measurements of T2* in adult frontal white matter from EPI data were highly variable and not in agreement with the conventional scan. Measurements in occipital white matter and thalamus had nRMSE below 8% and resulted in consistent T2* values across both scans, with no significant differences found, providing evidence that in regions remote from air-tissue interfaces, the EPI method achieves reliable T2* values.

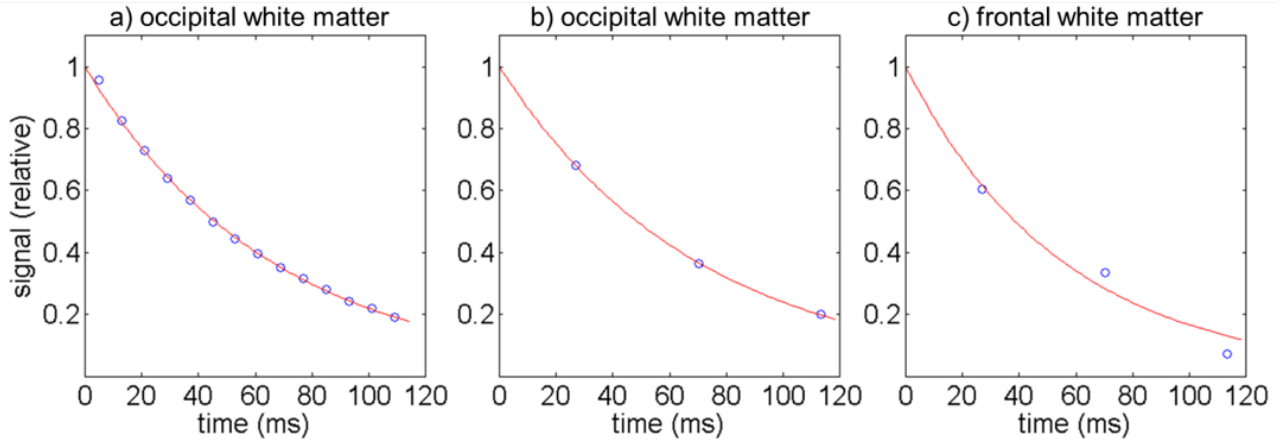

Figure 4.5: a) single voxel T2* fit in a multi-echo GRE scan. T2* = 66ms, nRMSE = 0.02, region = OWM b) single voxel T2* fit in a multi-echo single-shot GRE EPI scan. T2* = 70ms, nRMSE = 0.02, region = OWM c) T2* fit in a multi-echo single-shot GRE EPI scan with large fitting error. T2* = 56ms, nRMSE = 0.09, region = FWM.

Based on the adult results, voxels for which the T2* fitting resulted in nRMSE larger than 0.08 were excluded from further analysis in the fetal EPI data.

4.3.2 T2* in Fetal Subjects

There was substantial movement in 17 out of the 23 fetuses such that at least one slice in a full stack of images was misaligned through-plane by more than the slice thickness (3mm). The extent of inter-slice motion decreased for fetuses of greater gestational age, which is consistent with previous observations (Hayat et al. 2011). Figure 4.6 shows examples of single brain slices after inter-echo registration and the associated T2* maps for three different gestational ages. Table 4.3 presents a summary of results for normal and clinical fetal T2* measurements in the specified ROIs, together with previously reported values for pre-term neonates by Rivkin (Rivkin et al. 2004) and Lee (W. Lee et al. 2012). There is no significant difference between average T2* values of the clinical and normal group. For FWM, the mean T2* value across all subjects (pooled across both groups) is 236ms (SD 33). The average T2* of OWM is a slightly higher value of 257ms (SD 35). Thalamic T2* (THA) was measured at 157ms (SD 24ms). In all regions the fetal values are larger than the published data for neonates of similar age.

To assess reproducibility, mean T2* estimates were considered separately for each repeat for

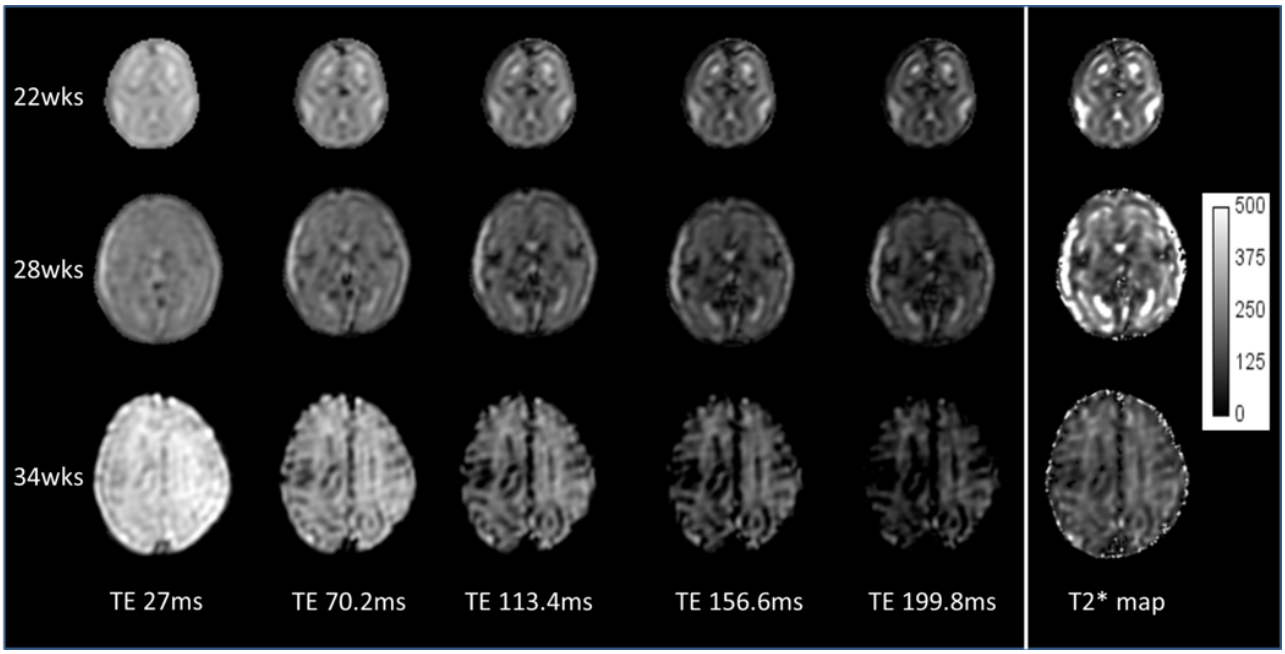

Figure 4.6: Examples of fetal scans at 22, 28 and 34 weeks of gestation. T2* weighted images for TE range 27ms-199.8ms and corresponding reconstructions of T2* maps.

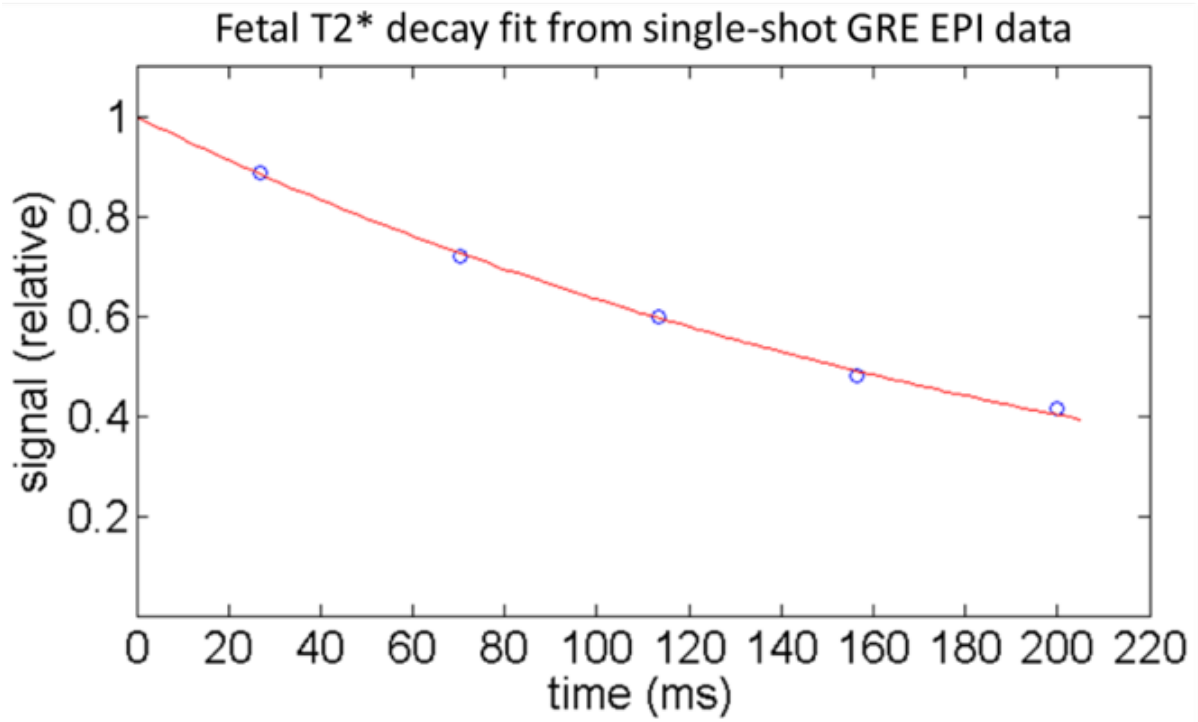

Figure 4.7: An example of a T2* decay fit for a voxel in the fetal frontal white matter. T2* = 220ms, nRMSE = 0.015

	Our study			Rivkin 04		Lee 12	
	Clinical	Normal	Total	Pre-Term	Term	Pre-Term	Term
subjects	9	15	24	4	10	8	8
age (wks)	29.4	32	31.1	33	42	30.8	43.5
stdev	5.3	3.3	4.3	0.6	0.9	2.2	1
	Frontal White Matter						
T2* (ms)	242	234	236	180	152	-	-
stdev	24	38	33	38	51	-	-
	Thalamus						
T2* (ms)	163	154	157	137	127	-	-
stdev	24	24	24	13	23	-	-
	Occipital White Matter						
T2* (ms)	269	249	257	148	142	193	122
stdev	23	40	35	39	37	49	14

Table 4.3: Fetal brain T2* values (mean and standard deviation) for normal and clinical subjects. Values from previous investigations into neonatal T2* are also listed. Both neonatal studies reported values for medial and lateral occipital lobes separately. For comparison these are displayed here as an average aggregate for occipital white matter. There is no significant difference between lateral occipital lobe (LOL) and medial occipital lobe (MOL) at 95% confidence interval for all cases, given the variance and size of these samples.

each subject. The standard deviation of the mean T2* estimates across repeats was divided by the mean to calculate relative variation for each subject. Results show that T2* estimates vary by 0.15 (SD 0.02) across repeats for frontal white matter, 0.17 (SD 0.02) for occipital white matter and 0.12 (SD 0.02) for thalamus. There was no significant correlation between this variation and gestational age.

Figures 4.8,4.9,4.10 show scatter plots of T2* values against gestational age with a linear regression fit to normal and clinical groups of data for FWM, OWM and THA. Mean T2* data from pre-term and term neonates, based on previous studies (19)(20), has also been added for comparison.

Table 4.4 summarizes linear regression analysis for fetal data, by reporting on the rate of change of T2* (slope), its standard deviation, coefficient of correlation and corresponding significance levels. There is a significant correlation between gestational age and reported T2* value for all regions tested in the normal group and for THA in the clinical group.

For slope comparison between clinical and normal groups there was significant difference found in FWM ($p=0.026$) and OWM ($p=0.001$), but not in THA ($p=0.274$).

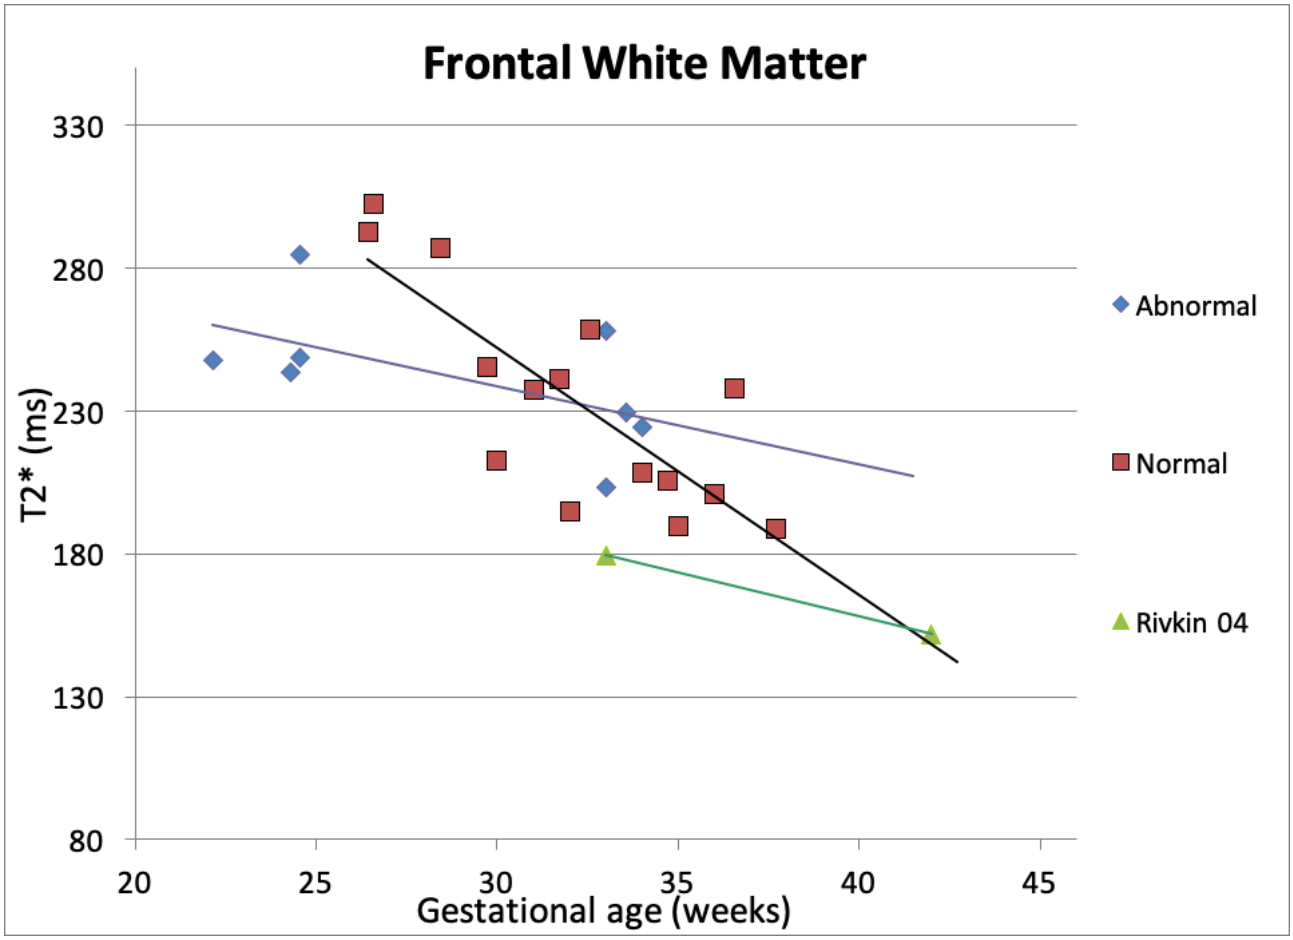

Figure 4.8: Scatter plot for Frontal White Matter displaying individual subject T2* values against gestational age. Trend lines are drawn for normal and clinical subjects. Previous measurements reported by Rivkin. Linear trend lines for clinical subjects in frontal and occipital white matter were removed as the correlation between gestational age and T2* was not significant, as reported in Table 4.4.

	Clinical		**Normal**	
	Slope (Std Error)	PCC (p-value)	Slope (Std Error)	PCC (p-value)
FWM	-2.7 (1.6)	-0.58 (0.14)	-8.6 (1.8)	-0.75 (0.00)
THA	-4.5 (0.5)	-0.97 (0.00)	-5.8 (1.0)	-0.83 (0.00)
OWM	-2.8 (1.4)	-0.64 (0.09)	-10.1 (1.4)	-0.90 (0.00)

Table 4.4: Parameter estimates for the linear correlation of T2* with gestational age for normal and clinical fetal subject groups. Pearson correlation coefficient (PCC) and corresponding p-values indicate congruency and significance of these estimates.

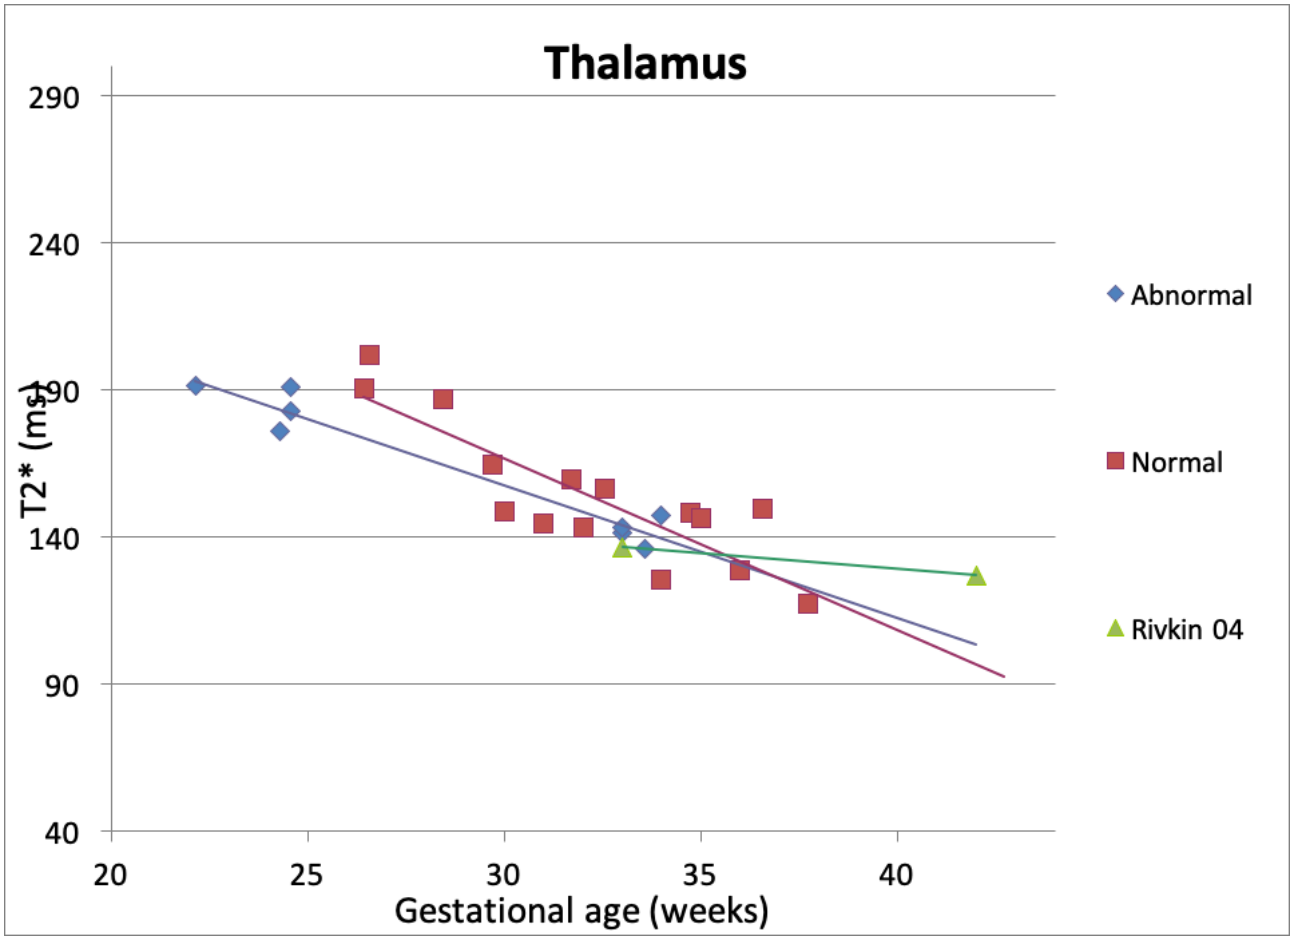

Figure 4.9: Scatter plot for Thalamus displaying individual subject T2* values against gestational age. Trend lines are drawn for normal and clinical subjects. Previous measurements reported by Rivkin.

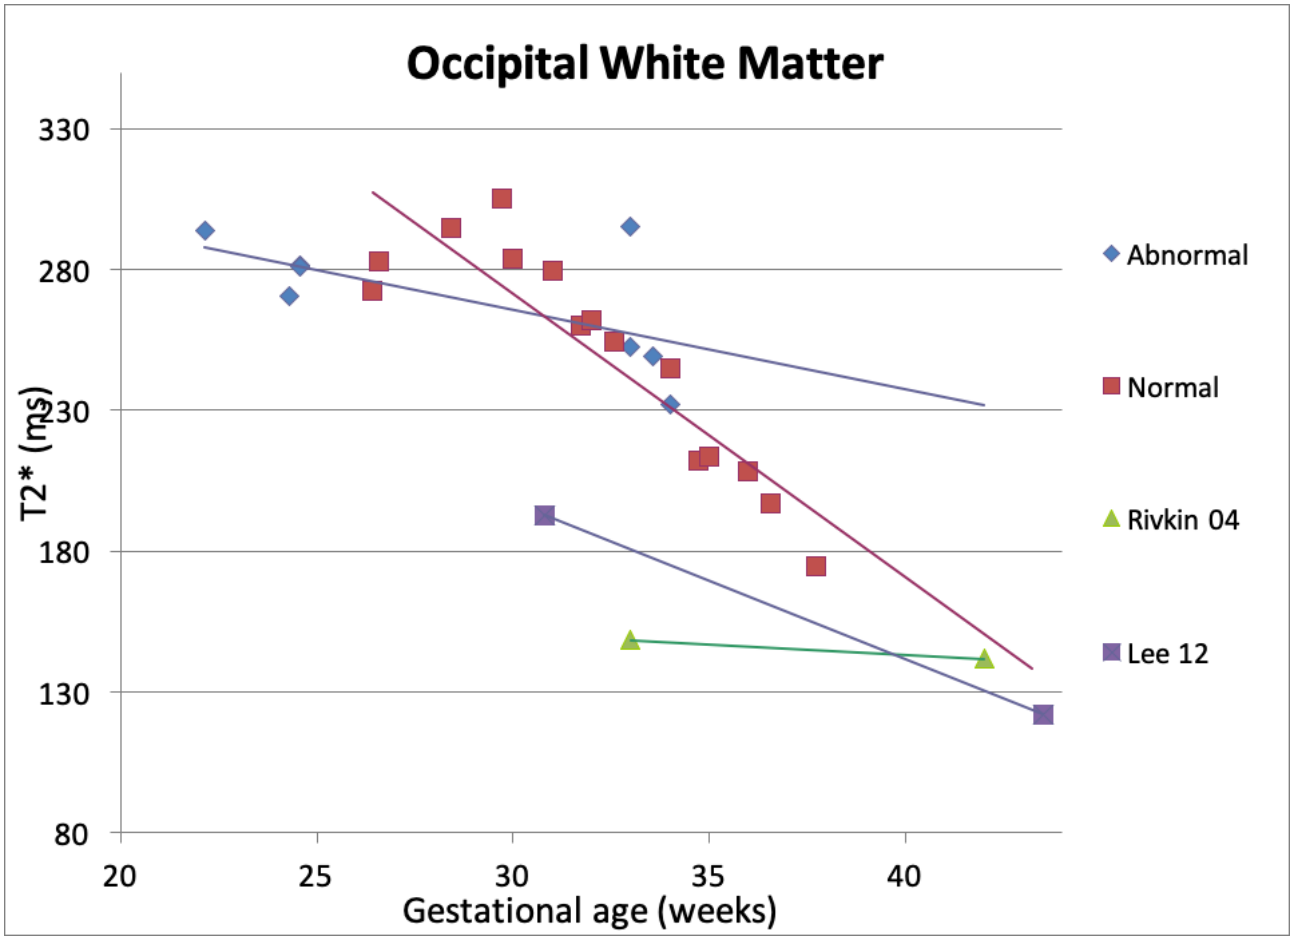

Figure 4.10: Scatter plot for Occipital White Matter displaying individual subject T2* values against gestational age. Trend lines are drawn for normal and clinical subjects. Previous measurements reported by Rivkin and Lee are also displayed.

For both FWM and OWM in the normal group, extrapolation of the linear trends for fetal T2* variation with gestational age to a gestational age of 43 weeks of age shows that it is consistent with the previously published (19)(20) term data from neonates at calculated uncertainty levels (Table 4.3). This was not the case for the thalamus however, for which extrapolated trends predict a lower T2* at 43 weeks of gestation than is found in neonates at equivalent age, for both normal and clinical groups. Extrapolating results from the clinical group for OWM and FWM regions were higher at 43 weeks of gestation than in previously reported neonatal values.

4.4 Discussions and Conclusions

This study presents measurements of T2* relaxation time for key anatomical regions in the fetal brain. Results suggest that mean T2* values in the fetal brain across the studied gestational age range are larger than relaxation times in the published literature on neonates and show a pronounced decrease with gestational age. Decline in T2* relaxation times may be related to increase in synaptic density, myelination and cerebral blood flow (19) as well as rapidly declining water content (Dobbing et al. 1973). Data was obtained both from normal subjects and from fetuses with diagnosed conditions (see Table 4.1).

The results were obtained using a robust method for performing T2* relaxometry based on a single shot multi-echo gradient echo EPI protocol. It provides whole head coverage and is highly tolerant to motion which is a common and often unavoidable problem in fetal scanning. The method was tested using data obtained from an adult and found to be consistent with a slower standard multi-shot measurement method in regions of the brain away from air-tissue interfaces, such as close to the frontal sinuses. As the minimum temporal spacing between points (inter-echo spacing) used to sample the decaying signal is limited by the requirement to spatially encode full images, the method is most suitable when T2* is long, which turns out to be the case for the fetus. The data obtained confirm that fetal T2* values are as long or longer than comparable data from neonates and substantially larger than values for adult brain. A contributing factor both to lengthening T2* and to achieving reliable measurements may be the

lack of air in close proximity to the fetal brain, since the amniotic fluid and surrounding tissues provide a much more closely matched magnetic environment. Gas bubbles in the maternal alimentary canal can cause local distortion and signal loss in some cases, but in our examples these effects never penetrated very far into the fetal head.

Average fetal T2* estimates obtained in this study are higher than those of pre-term infants of similar age in previous studies by Rivkin (19) and Lee (20) for frontal and occipital white matter. Previous data presented by Rivkin and Lee on OWM is heteregenous. In our case, OWM also differs to these studies and is measured higher than the average value for FWM. Linearly extrapolating fetal T2* measurements of the normal group to 44 weeks of gestation shows convergence of results to the reported term neonatal measurements for frontal and occipital white matter, but not for the thalamus. This is of particular interest because the thalamus is known to be a highly vulnerable structure that is reduced in size in premature infants (Srinivasan et al. 2007). The evolution of the T2* with gestational age may be more complex than the simple linear trend we could infer from the present data, and remains to be explored as larger datasets become available. Obtaining longitudinal data on individual subjects would be particularly interesting in this context. In the absence of overt lesions, preterm infants show numerous differences in brain development to that of term neonates (Holland et al. 1977). A comparison between the measured fetal T2* values and previously reported T2* values in preterm infants is therefore only used here as a reference. As more fetal relaxometry data becomes available, a further assessment and comparison of fetal and neonatal T2* values would be more appropriate.

The findings of this study may have significant implications for the optimal design of sequences for performing fetal fMRI, since the lengthened T2* values would imply longer echo times would be required for achieving optimal signal contrast. In addition, a direct measurement of T2* would also be possible for fMRI measurements. The use of multi-echo fMRI was previously shown to increase sensitivity of BOLD signal detection (Posse et al. 1999). Given the particular complexity of fetal scanning and associated artifacts, measuring BOLD changes based on quantitative T2* relaxometry maps, instead of T2*-weighted images, could improve the overall quality of fMRI results (Seshamani et al. 2013a). These factors will need to be balanced against the decreased data rates implied by lengthening TE or sampling T2* decay curves directly as recent work on

fMRI in adults suggests that faster sampling has advantages (S. Smith et al. 2013)(D. Feinberg et al. 2010a)(Ugurbil et al. 2013).

Further research may consider the benefits of T2* measurements outside of the fetal brain. A recent study Goitein (Goitein et al. 2013) adopted conventional multi-shot methods to assess the normal range of T2* in fetal liver. Decreases in fetal heart iron are associated with anemic subjects and is another area of valuable interest (Rao et al. 2007). In these applications the much shorter T2* presents a challenge for the motion tolerant, and hence robust, methods adopted here.

In summary, this study presents for the first time systematic measurements of T2* in major anatomical regions in the fetal brain. The decline of T2* with gestational age is consistent with trends shown in previously published neonatal data and could provide insight into brain development processes. The measurements made on the normal cohort can provide normative values by which to assess changes seen both in clinical fetal examinations and also preterm infants. Further work is needed to enlarge the study population, particularly to allow exploration of T2* changes in more homogeneous clinical groupings. However, since the method appears to work robustly, larger scale studies should be eminently feasible.

Chapter 5

Fetal and neonatal whole brain T2* mapping at 3T

5.1 Introduction

The goal of this study was to measure and compare fetal and neonatal T2* decay constants at 3T in multiple regions of the brain. To the best of our knowledge there are no other studies that reported normative values in either fetal or neonatal subjects at this field strength. This study extends on the previous methods introduced in chapter 4 by adding a fully automated segmentation of the brain for 14 distinct regions of interest. As part of the reconstruction step, it also presents a method for iterative slice to volume registration (SVR) between a stack of motion corrupted T2* weighted slices, acquired in transverse orientation, and a 3D reconstructed T2 volume.

5.1.1 Prior literature

Following the publication of our research in Vasylechko et al. (2013) and Vasylechko et al. (2015), which is described in chapter 4, a group of studies have emerged that confirmed the results for fetal T2* values in Frontal White Matter (FWM), Occipital White Matter (OWM) and

Thalamus (THA) regions. These new investigations were led by Anna Blazejewska (Blazejewska et al. 2015) and Shramistaa Seshamani (Seshamani et al. 2015a,b). Their publications made two novel contributions to the study of T2* relaxometry in the fetal brain. The first contribution made a new clinical observation and the second introduced a novel technical framework. In this section we describe the basic principle of the methods introduced by these studies, and discuss their relevance to the extended novel technical and clinical contributions presented in this chapter.

The key clinical contribution by Blazejewska et al. (2017) is the measurement of T2* values of fetal Cortical Gray Matter (CGM) and Cortical Plate (CP). CP T2* values were measured in fetuses with gestational age of 25 weeks and younger, while CGM was reported for subjects over 25 weeks of age. The knowledge of fetal brain CGM/CP T2* decay constants is valuable for optimization of BOLD fMRI in fetal resting-state analysis. However, CGM and CP appear as thin elongated grooves in the fetal brain and cover the area that is significantly smaller than the achievable resolution for single shot inplane T2* weighted acquisition of fetal brain scans. To address this, Blazejewska et al. (2017) devised a new method to obtain motion corrected $1mm^3$ volumetric T2* weighted images of the fetal brain at two different echo times, which were combined to form an estimate of quantitative T2* maps.

The reconstruction used the principle of cascaded slice to volume reconstruction, presented in Seshamani et al. (2013b). The data consisted of 100+ dual-echo single shot GRE acquisitions in the transverse orientation, which would provide sufficient coverage for all the datapoints in a single 3D volume. An explicit formulation for slice acquisition can be used to reconstruct the full 3D mapping with:

$$\hat{S}(x_k) = FBTS(x) + n \tag{5.1}$$

where $\hat{S}(x_k)$ is the estimated signal value for voxel k, F is the subsampling matrix, B is the blur matrix representing the point spread function, T is the motion estimate, $S(x)$ is the measured signal, and n is the observed noise. The initial process starts with selection of one of the acquired motion-corrupted T2* stacks of slices as the reference, to which all the remaining stacks are aligned. The stacks are progressively split through iterative stack to volume registration and

target update steps.

To estimate T2* decay constant maps, Seshamani et al. (2013b) combined two separately reconstructed 3D T2* weighted images at different TEs. A linear combination of signal intensities from both images was combined via a log linear relationship derived from simple monoexponential decay formulation.

While these studies were of significant clinical and technical value, we note a number of weaknesses in the presented approach.

The first remark relates to the robustness of the T2* quantification from dual echo acquisition. A standard formulation of exponential signal decay $S_i = S_0 e^{-TE_i/T2^*}$ has two degrees of freedom - the original signal intensity S_0 and the decay time $T2^*$. An analytical solution is derived by taking a logarithm of both sides to linearize the equation, which makes it possible to estimate the values with only two datapoints, such as two signal measurements at two different echo times. However, in the presence of variable noise $S_i = S_0 e^{-TE_i/T2^*} + \epsilon$, which is inherent to every MRI image reconstruction, a reliable measurement requires numerical estimation from at least three sample points along the decay curve. The accuracy of two-point fetal brain T2* constant estimation is addressed in the conference abstract of Blazejewska et al. (2015), which concluded that dual-echo T2* estimation is comparable to five echo T2* estimation when measured as an average across the whole brain. However, no remarks were provided to the effect of measurements across much smaller regions of interest where variation of T2* decay constants is expected to be much narrower, with a significantly smaller number of total available samples, such as the deep gray matter and cortical gray matter. Furthermore, an earlier work by Peran et al. (2007), performed an extended study of optimal number of echoes, echo time values and number of measurements required to estimate normative T2* decay constants in the adult brain, which is void of motion corrupted data and presents a simpler challenge. The study concluded that the optimal accuracy and precision for T2* estimation is achievable with a minimum of six echo samples that cover the time period up to the value of the T2* decay constant itself. For example, a reliable estimate of a voxel with a true T2* decay constant of 80ms would require a minimum of 6 sample points that will cover a 80ms time period.

The second remark relates to estimation of T2* constants on two independently reconstructed 3D volumes with different TE weighting. We note that any errors incurred in each of the reconstruction steps prior to T2* estimation, particularly during interpolation steps, carry a multiplicative effect onto the final estimate of the T2* constants. Any deviations in true signal intensities will significantly affect the T2* decay estimate in the small regions of interest. An alternate approach to this problem would be to perform T2* estimation on the acquired data first, and then follow with the 3D reconstruction steps. There are two possible approaches to such formulation. The first may utilize the reconstructed quantitative T2* maps as the source for 3D reconstructions, while the second may use original T2* weighted signal intensities from one of the echo times, to reconstruct a 3D volume, and then use the resulting transformation matrices on the estimated quantitative T2* maps to reposition the voxels to the correct destination in the final 3D volume.

Lastly, an altogether different approach would be to estimate both the T2* constants and the labels to which each data voxel belongs in the original coordinate frame in which the data had been originally acquired. As a consequence, the estimated T2* constants would not be subject to errors from interpolation of data samples during 3D reconstruction.

5.1.2 Contributions

Based on the above discussion, we present an alternate method for measurement of whole brain T2* constants, which performs estimations of average T2* values of each region of interest entirely in the acquired image coordinate space. T2* decay constants are estimated with multiple echo time points to ensure robust measurements. The difference between this new method and the work presented in chapter 4 is measurement of T2* values for all of the acquired data voxels, rather than just a few seed voxels.

Our second technical contribution is a fully automatic segmentation framework that propagates region of interest labels into the acquired T2* space for each voxel without the need for manual intervention. This presents a significant decrease in the amount of time that it takes to measure T2* values across multiple regions of interest for a single subject. Consequently, our study is

able to report T2* measurements for 14 unique regions of interest, in contrast to the previously reported 3 regions of interest. This extends the current clinical knowledge of T2* decay constants in the developing brain into a finer set of brain regions.

Our third contribution is an extension of T2* measurements to the neonatal domain. Previous studies were only able to compare fetal T2* decay constants to neonatal T2* decay constants by searching for relevant values reported in other works of literature. Our study had been designed specifically to make the direct comparison between fetal and neonatal cohorts possible. It uses the same scanner, the same acquisition sequence and the same postprocessing framework. As a result, our work is able to maximize comparability and make a number of extended evaluations between each cohort, including a direct comparison between neonatal and fetal subjects of similar gestational age. Knowledge of these values has valuable clinical relevance to the study of developmental trends between fetal and newborns.

Finally, we apply this framework to study T2* decay constants at 3T field strength. To the best of our knowledge, this is the first study of fetal T2* decay constants as well as the first study of neonatal T2* decay constants at this field strength.

5.2 Method

5.2.1 Acquisition

Setup

All scans were performed on a Philips Achieva 3T scanner. Parameters were optimized individually for fetal and neonatal protocols to achieve a good balance of spatial and temporal resolution, signal to noise ratio and clinically safe noise level that would yield comparable images between these two groups of scans. Both scan protocols used a combination of single shot multi-echo gradient echo EPI with SENSE reconstruction and partial Fourier acceleration. Three repeats were acquired for each multi-echo T2* weighted whole brain acquisition. All T2* weighted scans

	Fetal	Neonatal
FOV (mm^2)	320x320	200x170
Resolution (mm^3)	2.5x2.5x2.5	2x2x2.5
TE range (ms)	36-396	35-378.8
ΔTE (ms)	60	57.3
SENSE factor	2.5	2
Partial Fourier	6/8	6/8
Fat Suppression	SPIR	-

Table 5.1: A summary of key acquisition parameters used for fetal and neonatal sequence variants that were optimized for single shot multi-echo GRE EPI on a 3T system.

were acquired in transverse orientation.

Resulting parameters for fetal acquisition were (see Table 5.1): TE/ΔTE 36/50ms, 8 echoes per each shot, 3mm isotropic resolution, 320x320 FOV, SENSE 2, echo train length of 51, TR 10s. A 32 channel cardiac/torso phased array coil was used for signal reception. Fetal scanning also used spectral presaturation with inversion recovery (SPIR) for fat suppression. Each scan consisted of 43 slices of 3mm thickness with no overlap, resulting in 10.7s per each whole brain acquisition.

Neonatal acquisition parameters were (see Table 5.1): TE/ ΔTE 50/80ms, 8 echoes per each shot, 2.5mm isotropic acquired resolution, 270x180 FOV, SENSE 2.5, echo train length of 58, TR 13s. A special neonatal head array coil was used for signal reception. Depending on the fetal head size, each scan produced 30 slices of brain data, resulting in 13s acquisition.

Both acquisition variants were optimized to clinically acceptable noise levels, within the range of 101-103 dBA. Measurements were performed with a microphone placed directly in the bore of the scanner. Adjustments were made to gradient slew rates in order to reduce the noise levels.

All scans were preceded by a full clinical examination, which included a set of structural T2 weighted spin echo acquisitions in all 3 orthogonal planes with a 1.17x1.17x1.25 mm^3 resolution with full brain coverage. These structural T2 scans were used to 3D reconstruct the perinatal brain via the Slice to Volume Reconstruction algorithm and then delineated with the automatic segmentation framework. The resulting 3D reference segmentation and the 3D structural T2 image of the whole brain was then used to map individual low resolution T2* slices to the

correct position. Furthermore, a SENSE re-calibration scan was also performed prior to each T2* weighted acquisition in order to subsume the best possible calibration of coil sensitivities. This helped to reset coil sensitivity maps and therefore aided to improve the quality of the single shot multi-echo GRE EPI acquisitions.

Verification

T2* relaxation constants decrease non-linearly at higher field strengths. As a result, it was necessary to verify the robustness of the fast acquisition method for T2* estimation at 3T field strength. Single shot multi-echo GRE EPI acquisition was compared to a conventional non accelerated multi-echo gradient echo acquisition that is commonly used to accurately measure T2* decay constants in adult subjects. The conventional method included no SENSE and no EPI acceleration as well as a much denser sampling of the T2* decay constant. The first TE was set 12ms, the interecho spacing at 12ms and 15 echoes were acquired for each image in total. The FOV was set to match that of the adult brain volunteer at 250x200. All other parameters of the scan were kept identical, including resolution, TR and half fourier. Total scan time for a single slice with the conventional non accelerated method was 3mins, and 250ms for the accelerated acquisition. As the T2* decay in adults is much shorter than neonates and fetal subjects, only the first echoes were used from the accelerated scan to estimate voxel T2* values.

5.2.2 Subjects

The final study group consisted of 10 fetal control and 16 neonatal subjects. All subjects were assessed by a trained neuro-developmental specialist and were confirmed to have no clinical conditions at the time of scanning. In the neonatal cohort, 8 subjects were born and scanned at preterm age, while 8 subjects were born preterm but scanned at term equivalent age (see Figure 5.2). The study received approval from the local ethics research committee and all scans had received informed parental consent. No sedation was administered to any subject.

Group	Age at Scan	Number of Subjects
Fetal Control	31.9±1.6	10
Neonatal Preterm	33.9±2.1	8
Neonatal Preterm at Term	41.5±0.8	8

Table 5.2: Distribution of the acquired subject data. Due to limitations imposed by acquired resolution, all patients with gestational age of less than 30 weeks of age were excluded from the study.

5.2.3 T2* estimation

Fetal imaging largely avoids signal losses that arise from macroscopic field inhomogeneities in gradient echo scans. This is due to the absence of tissue-air interfaces, except for incidental occasions where the fetal head is positioned close to a bolus of air near in the mother's gut. As a result, fetal gradient echo imaging can be modelled with a simple monoexponential decay in all parts of the fetal brain.

Conversely, estimation of T2* decay in neonatal scans suffers large signal losses at air tissue interfaces. In order to account for these signal losses, we used an iterative algorithm presented by Fernandez-Seara et al. (2000) to estimate and remove the effect of macroscopic field inhomogeneities prior to computing T2*. In the presence of linear field variation with the largest effect coming through slice direction, T2* decay can be modelled as:

$$S_i\left(TE_i\right) = S_0 e^{-\frac{TE_i}{T2^*}} z_0 \operatorname{sinc}\left(\frac{\gamma}{2} z_0 G_b TE_i\right) \tag{5.2}$$

where γ is the gyromagnetic ratio, z_0 is the slice thickness and G_b is the unknown field inhomogeneity gradient that has a linear effect in the through slice direction.

The value of G_b is not known and must be estimated iteratively by finding a local minima with the chosen loss function between the data and the model, which in our case is done via weighted least squares approximation. Estimating all three parameters in parallel is challenging. Therefore, the first step is to estimate initial values of T2* and S_0 using an ordinary least squares regression for a simplified monoexponential model, as described in Vasylechko et al. (2015). The optimal values are then fed into an expanded model with an additional *sinc* term.

Following this, the process of nonlinear iterative regression is repeated using the initial values for T2* and S_0 for every chosen G_b, increased in steps of from 0 to an empirically selected G_{bfmax}. The non linearity of G_b has a number of local minima, and therefore convergence to global minima is not always guaranteed. Thus, if the estimated optimal G_b at a given step is greater than the value of the first zeroth crossing of the sinc function, which is defined by:

$$G_{bfmax} = \frac{2\pi}{\gamma z_0 TE} \tag{5.3}$$

then an echo is removed from the end of the echo train. Removal of the last echo extends the G_{bfmax} and expands the search space for a global minimum. This process is repeated until G_b optimal is smaller then G_{bfmax}, which signifies the global minimum. The alternative approach, presented by Dahnke et al. (2005), presents a more robust selection of initial values for G_b, estimated from the evolution of signal phase in the through slice direction by considering the gradient that lies across neighbouring slices. However, due to the risk of inter-slice motion in neonatal scanning, this initialization may not be optimal as estimation of through slice gradients would not represent the true gradient maps between shifted slices.

In order to account for through slice susceptibilities, we applied the aforementioned correction method introduced by Fernandez-Seara et al. (2000) while computing T2* estimation in the neonatal images.

5.2.4 Automatic segmentation of T2* weighted slices

Acquired T2 weighted and T2* weighted data was processed in a series of steps in order to obtain the final segmentation in the image coordinate space of the acquired T2* weighted images into 14 distinct regions of interest. An overview of the entire framework is given in Figure 5.1. Figure 5.2 shows an example of the output images. The full pipeline is described in the following subsections below.

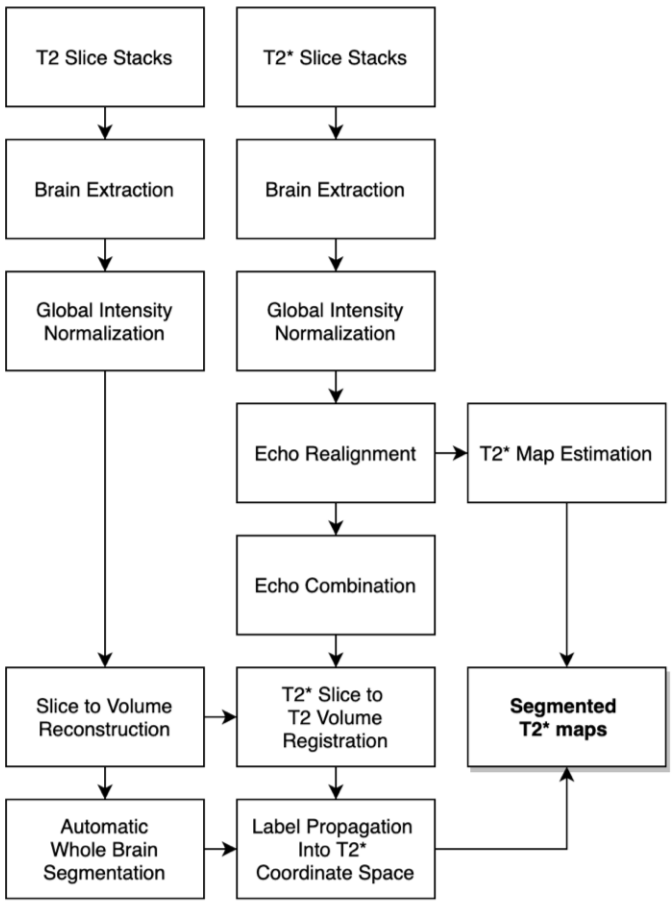

Figure 5.1: An overview of the proposed processing framework used for obtaining segmentations on the individual T2* slices. The steps are described in detail in each subsection of Section 5.2.4

Brain Region Extraction

A spherical 3D brain mask centered on the fetal brain was created in snapITK to delianeate majority of maternal tissue in T2 and T2* images of the fetal cohort. Such bounding box aided in registration of final T2 and T2* images, as well as slice to volume reconstruction step for T2 images. The masks had to be large enough to contain full range of fetal head motion. Unlike fetal brain images, which are surrounded by maternal tissue, neonatal images did not necessitate manual brain region extraction. Instead, Brain Extraction Tool (BET) (S. M. Smith 2002) from FSL library was used to extract intra-cranial tissue of the brain and exclude the skull.

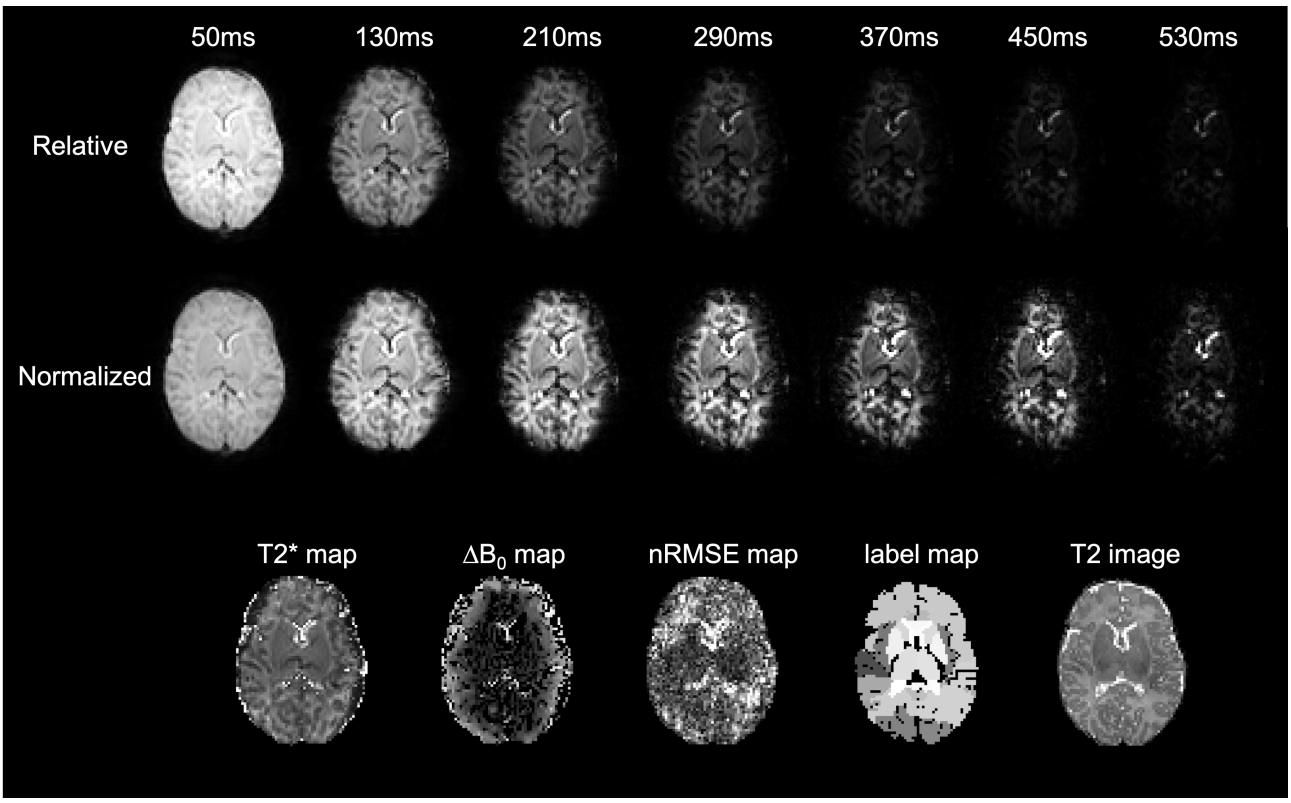

Figure 5.2: Top and middle rows: Consecutively acquired T2* weighted scans of the same brain slice, with increased TE value. Top row - the signal intensity in all images is shown relative to the first image (TE1). Middle row - the signal intensity in each image is normalized with respect to its min-max range. Bottom row: For the same slice - estimated T2* map, through slice susceptibility map, nRMSE map, brain segmentation, corresponding T2 slice to which the T2* weighted data was fitted.

Global Intensity Normalization

Global intensity normalization was performed with ITK's N4 algorithm (Tustison et al. 2010) to remove intraslice and interslice signal intensity drifts. Such drifts are caused by inhomogeneities of the receiver coil, causing slowly changing 'shading' effects that impact registration algorithms with intensity based methods. Global intensity normalization was performed on both T2 and T2* images. This is particularly true for subjects with a lot of motion. In order to reduce the effect, a field map was acquired before and after each T2* acquisition to reset the coil sensitivities even before N4 algorithm was used.

T2* echo re-alignment

For images with particulaly large motion small misalignments between individual echoes of T2*W data were observed. This had only occured in fetal scans where free motion of the fetus in the amniotic fluid is a primary concern. Changes in position between echo images were small and similar similar to those observed in the prior study at 1.5 T. IRTK rigid body registration (Studholme et al. 1999) between echoes was used to align the echoes and improve the accuracy of T2* estimates. Registration was cascaded by registering a later echo onto its predcessing neighbour, and then propagating the resulting transformation again onto the earlier echo. This was repeated until the first echo was reached for all registrations.

Slice to Volume Reconstruction

Fetal T2 images were reconstructed using slice to volume reconstruction algorithm from Kainz et al. (2015) to obtain a 3 dimensional view of the brain (see Figure 5.3). It was only possible to perform such reconstruction on T2 images as they contained at least 7 stacks covering the whole brain of the subjects with at least 2 orthogonal views. It was not possible to perform the same procedure on T2* weighted images as only 3 repeats in a single orthogonal view were acquired as part of the protocol. The aim of slice to volume reconstruction (SVR) is to build a complete 3 dimensional representation of the brain by aligning multiple stacks of slices of a moving fetal or neonatal subject. This process is construed by a set of iterative estimations of motion from temporally contiguous stacks of the acquired stacks, starting from large blocks of stack to stack alignment, to progressive splitting of each stack into smaller temporally contiguous blocks and individual realignment of each to a reference volume, until individual slices are aligned into the 3D representation of the brain. At the end of each iteration, a reference image is updated by estimating an average of realigned stacks of slices. Final image is a self consistent high resolution, high SNR 3 dimensional representation of the whole brain. The full description of the method can be found in Kainz et al. (2015).

Occasional slices were corrupted by very large motion to an extent that anatomical details were

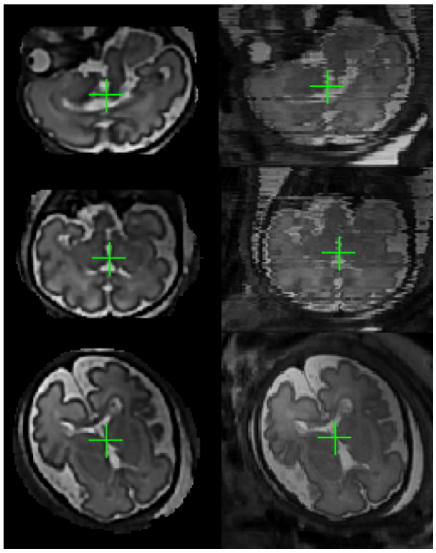

Figure 5.3: An example of a single stack of slices acquired in the transverse orientation (right column). When these slices are viewed in other planes, such as saggital or coronal planes (rows 1 and 2 in the right column), the anatomy of the fetal brain is not coherent. Slice to Volume Reconstruction (SVR) is used to align all of the acquired slices from multiple stacks into a coherent 3D representation of the fetal brain (left column).

not distinguishable and were removed from the stacks of slices prior to reconstruction. For some subjects, manual landmarks were necessary to be defined to aid the initial registration process. This consisted of 4 key pairs, which usually were placed on both eyes and dorsal/ventral midlines. The algorithm developed by Kainz et al. (2015) includes a parallelization of motion estimations that runs on a set of GPUs. As a result, each reconstruction took on average 30mins per fetal / neonatal scan, as opposed to 3 hours on a standard CPU architecture.

T2* to T2 registration

Individual T2* weighted slices were mapped to the reconstructed 3D T2 fetal brain volume images. An example of the final output is shown in Figure 5.5). An iterative scheme was designed to reach a fine alignment between T2* slices and T2 volume, as described in the flowchart shown in Figure 5.4.

At first, motion corrupted T2* weighted stacks of slices were rigidly registered wholly to a 3D T2 volume, as a volume to volume registration. Next, by splitting the T2* stack into smaller and smaller batches of time-continuous slices, the registration was refined. This approach had

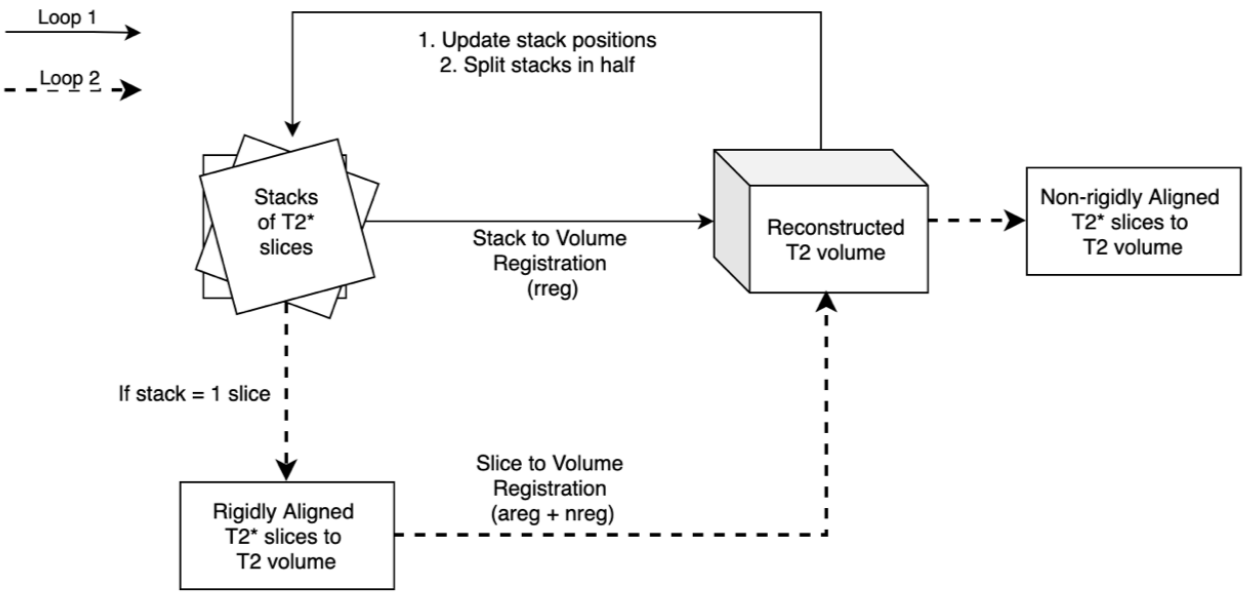

Figure 5.4: An overview of the processing scheme for mapping individual T2* slices to 3D reconstructed T2 volumes.

been found to yield better overall registration accuracy than direct registration of individual slices into the T2 volume. Direct registration of T2* weighted slices into T2 volume had on average performed well, however it failed in few instances with individual slices that had a relatively large amount of motion. The T2* weighted slices that were used for registration were constructed by combining the first three readouts of the T2* weighted echo train at the first three TE values. A set of auxiliary experiments had shown this combination to achieve better registration accuracy than other variants. The additional variants that were tested were - individual registration of the first, second or third T2* weighted images; registration of the combined first two echoes; registration of more than 3 echo images; registration of the estimated quantitative T2* map slices. The target update step was not implemented for two reasons. First, there was insufficient data to construct a 3D target from the acquired T2* weighted images. Secondly, a 3D reconstructed T2 brain images were already available as a whole 3D representation of the brain and could be used to roughly align the slices with rigid registration in an iterative manner. Affine registration was added to account for distortions in T2* weighted space that vary with the acquired direction. The affine registration parameters were optimized to provide only small adjustments to the final estimates. The output consisted of transformation

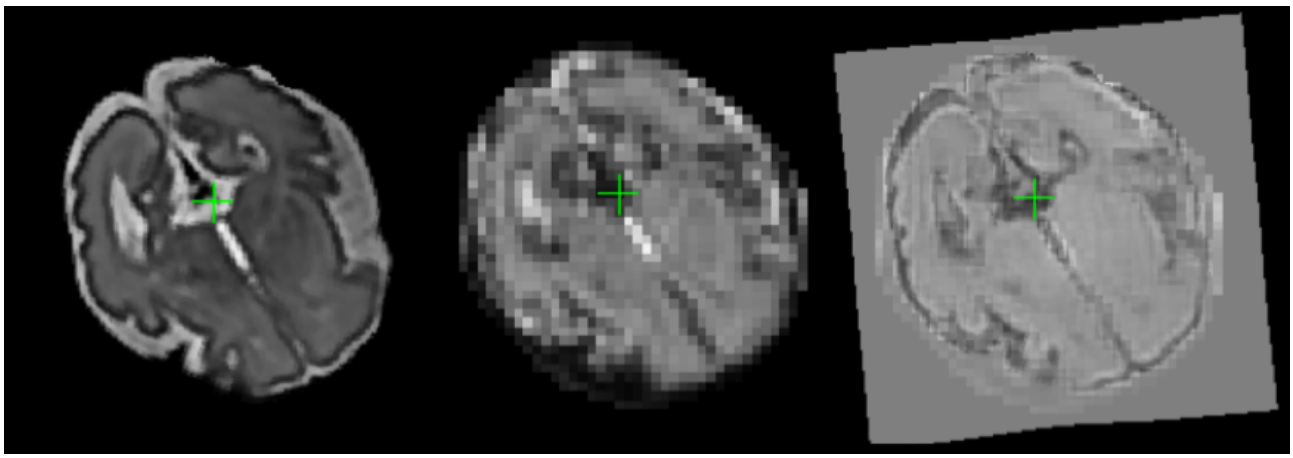

Figure 5.5: Example of registration between a 3D reconstructed T2 image (left) and a stack of T2* brain slices (center). A visual error comparison is shown on the right.

matrices between the reconstructed 3D T2 volume and individual motion corrupted T2* slices. As a result, the inverse transform allowed for each voxel in the 3D T2 volume to be known in the original image coordinate space of the T2* weighted data. As each 3D T2 volume had a corresponding segmentation map, it was therefore possible to assign a class label for each voxel in the acquired T2* image coordinate space and match it to the estimated T2* decay constant. Measurements from the same class label were aggregated to report an average value for each region of interest.

Segmentation

An automatic multi-atlas segmentation framework (Makropoulos et al. 2014) was applied to 3D reconstructed T2 volumes to segment the brain into 83 distinct regions of interest. The framework utilizes age specific atlases of the neonatal brain known as ALBERTs (I. S. Gousias et al. 2013) that cover gestational age range of 28 to 46 weeks. Of the 83 labels provided by the segmentation, 32 labels are assigned to cortical tissues for both gray and white matter segments, 18 deep brain structures and 1 label for the CSF. The framework was used for both fetal and neonatal images. Initial processing required the images to be void of non brain tissue and to be bias corrected. The approach to each subject group for these two steps is described in the sections above. Next, a multi step registration was performed by cascading rigid, affine and nonlinear registrations to align the selected ALBERT atlas to the T2 volume. In this

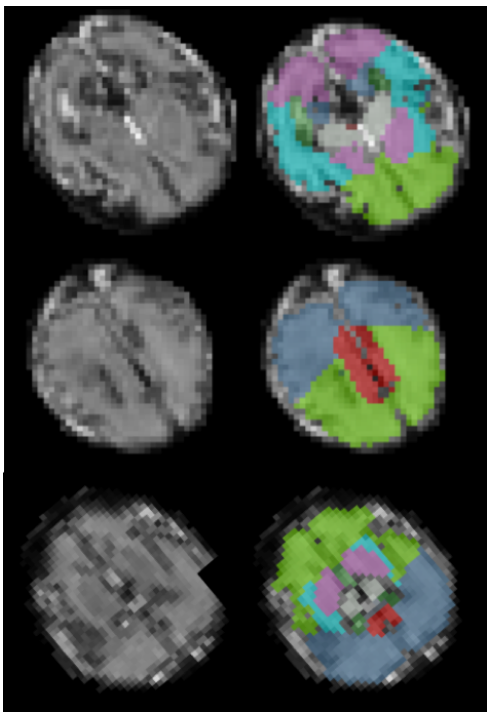

Figure 5.6: Example segmentation on T2* weighted slices (echo 2), propagated from a 3D reconstructed T2 image. The labels correspond to the following regions: green - frontal lobe, blue - parietal lobe, light blue - temporal lobe, purple - basal ganglia, red - limbic system.

study, rigid and affine registration parameters had to be optimized for lower resolution of the reconstructed fetal data. Additionally, in 3 fetal cases, manual landmarks had to be assigned to initialize the registration between a given T2 image and the corresponding ALBERT atlas. The final segmentations are obtained in two steps. First, the T2 image is segmented into separate tissue classes via a 4 class k-means clustering algorithm - WM, GM, CSF and non-brain tissues. Then the tissue priors are combined with ALBERT atlas priors and fed through an Expectation-Maximization framework to obtain the segmentation within the T2 image space.

Based on the above approach, two considerations were made with respect to the segmentation framework to optimize it to this study. Firstly, it was noted that the original ALBERT atlases, for which this segmentation was developed, were based on 0.85x0.85x1 mm^3 resolution, whereas the 3D reconstructed T2 images, on which it was performed, had an isotropic resolution of 1.25 mm^3. It was further noted that the final resolution on which these labels would be applied on, within the acquired T2* image space, was 3 mm^3. As a result, the 83 output labels were combined into a superset of 15 labels to account for the 3.5 times reduction in the effective resolution at which they would operate on. See an example of slice segmentations in Figure

5.6. The second consideration was made with respect to the number of voxels that would be available for reliable measurement of aggregate T2* decay constants in the smallest regions of interest. Based on empirical observation, it was found that for reliable estimates of cortical GM, within the acquired T2* resolution, the minimal gestational age at which such automatic segmentation would be reliable is 30 weeks of age. This is based on the observation for the thickness of cortical gray matter in the sulci of the brain at 30 weeks of age. Therefore, in order to allow for robust measurements of T2* decay constants using the proposed automatic segmentation and reconstruction method, all fetal data with gestational age below 30 weeks of age was excluded.

5.3 Results

Each slice of each individual T2* weighted scan was visually inspected at the time of acquisition and any signal dephasing due to air boluses in mother's gut, signal losses due to significant motion or other artifacts were noted. All slices that contained irregularities were excluded from the study. A final clinical examination report was reviewed to confirm diagnosis of each patient. A qualitative graded evaluation was performed for each patient for each step of the post-processing pipeline. This included overview of quality of slice to volume reconstruction of T2 images, segmentation of T2 images, alignment T2* weighted echo images and slice to volume registration between T2* and T2 images. The evaluations were carefully considered and any instances of notable reductions in quality were discarded. Finally, all fetal subjects below 30 weeks of age were excluded from the study due to the limits imposed by the automatic segmentation algorithm. As a result, the fetal data had been reduced from a total of 16 to 10 subjects, and neonatal data was reduced from 21 to 16 subjects.

Estimated T2* maps, nRMSE maps and through-slice susceptibility gradient maps were compared between the subject groups and their respective ROIs. Any deviant observations were further assessed by plotting individual decay curves for selected voxels to ensure robustness of the fitting method. The measures of mean, standard deviation, skewness, kurtosis, median and

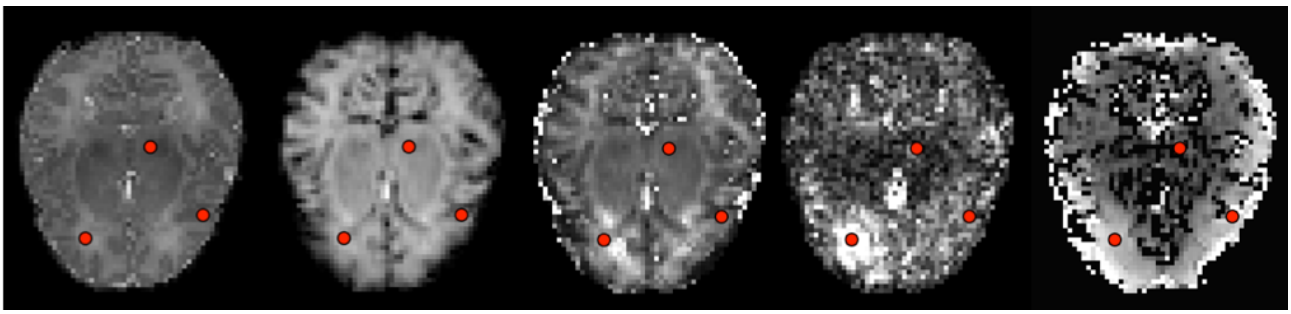

Figure 5.7: a) T2 weighted image b) T2* weighted image (echo 2) c) T2* map d) normalized root mean squared error map for the T2* relaxation fit e) through slice susceptibility component estimate This example depicts three regions of interest in the neonatal brain – one in the thalamus and two in frontal lobe. One of the frontal lobe areas exhibits high field inhomogeneity effects and signal loss due to proximity of air-tissue interface near the sinuses. This region has a high nRMSE with respect to the other areas that is not recoverable via through-slice susceptibility correction method.

interquartile ranges were taken to assess data distribution properties. All of the processing was performed in the original T2* acquisition coordinate space. No resampling of the T2* data had taken place.

5.3.1 nRMSE threshold

Normalized root mean squared error (nRMSE) threshold was re-estimated by empirical assessment of individual voxel fits in multiple regions of the brain across in multiple subjects. The nRMSE threshold from the prior study at 1.5T could not be used due to the difference in T2* model that was used for the neonatal data, which included through-slice susceptibility component. Therefore, it had to be re-estimated.

Figure 5.7 shows example of three regions of interest in a neonatal brain scan image. The first region depicts thalamic gray matter which is unaffected by through slice susceptibility effects, and therefore the difference between monoexponential and through-slice susceptibility corrected model fits is minimal. The second and third regions show the frontal cortex values, taken from the same slice. The first measurement shows high signal dephasing effect which cannot be recovered with our model. The nRMSE is measured at 0.10 and the model does not provide good fit to the data. In contrast, another voxel in the same slice of the frontal cortex has high through slice susceptibility effect that is recoverable with the model. Its corresponding nRMSE

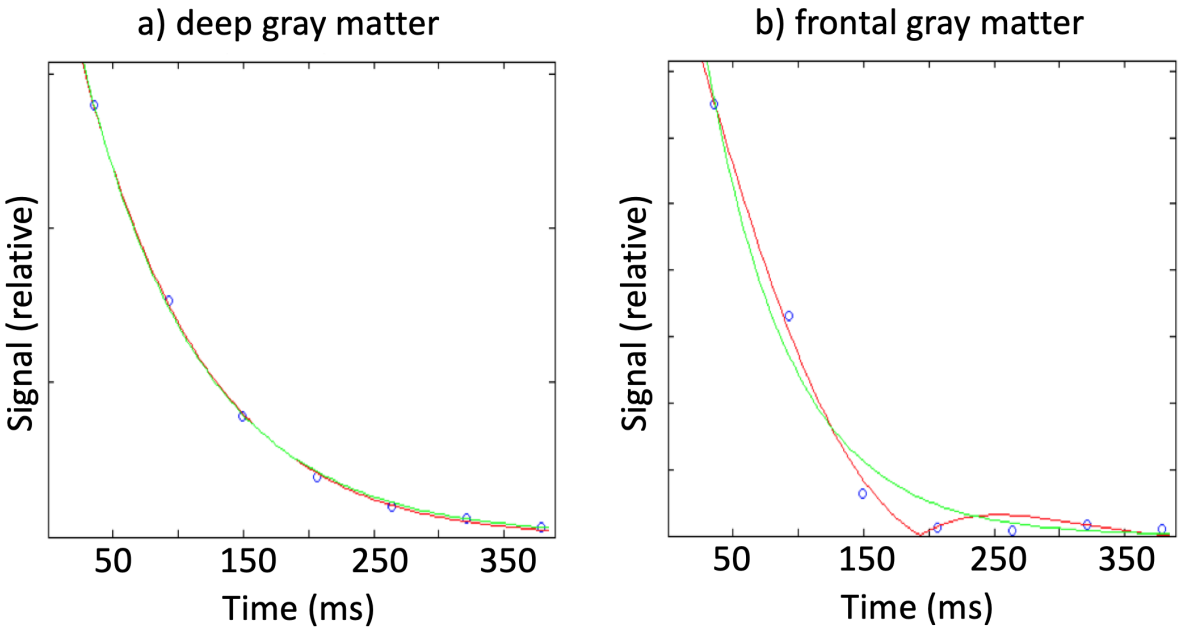

Figure 5.8: Individual voxel T2* signal decay fit with a mono-exponential model (green) and a through slice susceptibility corrected model (red) in the neonatal brain. a) signal in the deep gray matter in the thalamus. Both fitting methods achieve alignment as the signal is unaffected by through slice susceptibility gradient with T2* estimates of 97ms and 95ms, with nRMSE estimates of 0.019 and 0.013, respectively b) signal in the frontal gray matter with high through slice susceptibility signal loss. Through slice susceptibility correction adjusts the T2* estimate from 146ms to 64ms, and lowers nRMSE from 0.118 to 0.054.

measure is 0.05. An extensive empirical study across 200 individual voxel plots across a variety of regions and subjects had concluded that a threshold of 0.07 nRMSE would yield good quality of fit (see an example in Figure 5.8). Therefore, any T2* estimate that above 0.07 nRMSE had been automatically discarded.

5.3.2 General Observations

Mean T2* decay constants in fetal group were higher than those of preterm group in every region of interest, while the mean T2* decay constants in preterm group were in turn higher than those of preterm at term group (see Figure 5.3). These observations are consistent with the observations from previous studies at 1.5T for the three regions of interest that were measured - in frontal WM, occipital WM and thalamus. We further note that fetal occipital WM has the highest estimated T2* decay constant at 230ms, which is significantly lower than frontal WM,

	Fetal (μ, σ)		Preterm (μ, σ)		Preterm At Term (μ, σ)	
Temporal GM	190	24	126	17	99	7
Temporal WM	208	16	166	13	148	22
Frontal GM	171	16	143	19	120	16
Frontal WM	202	23	191	10	172	14
Parietal GM	175	15	139	18	130	16
Parietal WM	206	25	198	21	186	10
Occipital GM	200	23	105	18	98	14
Occipital WM	230	27	127	30	124	25
Limbic system	167	12	151	8	144	10
Basal ganglia	148	15	149	13	129	12
Thalamus	152	17	137	12	122	10
Corpus Callosum	186	23	159	24	149	19
Cerebellum	156	12	92	16	85	22
Brainstem	145	14	134	9	127	12

Table 5.3: T2* estimates (in ms) in 14 regions of interest. All measurements were made in the acquired image space via an automatic brain reconstruction and segmentation framework.

at 202ms, and thalamus, at 152ms. The order of these observations coincides with our previous measurements in Vasylechko et al. (2015). The lowest T2* decay constants are measured in the brainstem at 145ms and basal ganglia at 148ms. The T2* measurements of the cerebral cortex, both GM and WM, were higher in all subdivided regions of interest, than the structures in the diencephalon and the brainstem.

5.3.3 Fetal vs Preterm

Pairwise comparisons using ANOVA were performed between the three subject groups (see Table 5.4). Significant statistical difference at 95% confidence level was found between fetal and preterm subjects for the following regions - temporal GM, temporal WM, occipital GM, occipital WM, parietal GM, limbic system and cerebellum. The remaining 7 regions had no significant difference - parietal WM, frontal WM, frontal GM, thalamus, corpus callosum, brainstem, basal ganglia. By deducting p-value estimates from sample mean and standard deviation measurements in the previous fetal studies (Blazejewska et al. 2017; Vasylechko et al. 2015) with sample mean and standard deviation measurements in the previous preterm studies (W. Lee et al. 2012; Rivkin et al. 2004), a comparative analysis to previous findings could be made. It was found

Group 1 vs Group 2	**Fetal vs Preterm**	**Fetal vs Preterm At Term**	**Preterm vs Preterm At Term**
Temporal GM	0.00	0.00	0.40
Temporal WM	0.00	0.00	0.58
Frontal GM	0.07	0.00	0.24
Frontal WM	0.70	0.01	0.25
Parietal GM	0.00	0.00	0.90
Parietal WM	0.93	0.26	0.76
Occipital GM	0.00	0.00	0.99
Occipital WM	0.00	0.00	1.00
Limbic system	0.04	0.00	0.75
Basal ganglia	1.00	0.07	0.06
Thalamus	0.31	0.00	0.32
Corpus Callosum	0.15	0.02	0.91
Cerebellum	0.00	0.00	0.99
Brainstem	0.36	0.03	0.81

Table 5.4: Pairwise comparisons of T2* mean estimates between subject groups using two tailed t-test with p-values reported at 95% confidence interval. Significant differences are found in cortical regions, cerebellum and limbic system between fetal and preterm subjects, indicating non linear change in T2* values. No significant differences are detected in any areas between preterm and preterm at term subjects.

that our ANOVA tests between fetal and preterm groups were consistent with previous findings for occipital WM regions. However, the significance tests' outcomes had differed for frontal WM and Thalamus regions between fetal studies of Vasylechko et al. (2015) and Blazejewska et al. (2017), when compared to the preterm group from Rivkin et al. (2004). The results from current study are consistent with significance tests derived from Vasylechko et al. (2015)/Rivkin et al. (2004), but not from Blazejewska et al. (2017)/Rivkin et al. (2004). The samples means and standard deviations from Blazejewska et al. (2017)/Rivkin et al. (2004) indicate a significant difference in both regions between fetal and preterm measurements. However, we note that the mean gestational age for Blazejewska et al. (2017) study was significantly lower, at 29 weeks, in comparison to the mean age of Vasylechko et al. (2015) study, at 32 weeks. Furthermore, both Vasylechko et al. (2015) and Blazejewska et al. (2017) had indicated significant correlation of T2* change with age in fetal subjects. As a result, the difference in outcomes for the significance test which uses Blazejewska et al. (2017) results is most likely to be attributed to lower gestational age of its subject group. Finally, it is also noted that the gestational age of fetal and preterm

groups of the current study are highly similar to that of Vasylechko et al. (2015) / Rivkin et al. (2004), which further supports the argument for similarity of the outcomes in ANOVA.

5.3.4 Preterm vs Preterm At Term

ANOVA analysis between preterm and preterm at term subjects had shown no significant difference in any of the 14 regions of interest. These results were consistent with group comparison reported for thalamus, frontal WM and occipital WM in the previous findings by Rivkin et al. (2004). As noted in the previous chapter, occipital WM comparisons from Rivkin et al. (2004) and W. Lee et al. (2012) were heterogeneous. Consequently, the significance tests in this study are also heterogenous with those reported by W. Lee et al. (2012). We note that the average gestational age of the preterm group from W. Lee et al. (2012) was 30.8 weeks, while it was 33.0 weeks for Rivkin et al. (2004) and 33.9 for the current study, which is most likely reason for the disparity.

5.3.5 Coefficient of Variation

Assessment of coefficient of variation (CoV) for fetal, preterm and preterm at term groups was lower or equal to CoV in previous studies at 1.5T. The changes in CoV between the subject groups were heterogeneous (see Figure 5.5). CoV decreased, increased or stayed relatively constant with gestational age depending on the region of interest. Mean and standard deviation of CoV across all regions of interest was measured at 10±1.9%, 12±5% and 12±5% across fetal, preterm and preterm at term groups respectively. The largest CoV, with values above 15%, were detected in occipital WM, occipital GM and cerebellum. This observation is notable for the discussion of disparity between the previous findings of W. Lee et al. (2012) and Rivkin et al. (2004) for occipital white matter. W. Lee et al. (2012) reported that variation of T2* decay constant estimates was much larger in preterms than preterm at term subjects. Meanwhile, Rivkin et al. (2004) study had shown no change in variation, which is also confirmed in our study. Both Rivkin et al. (2004) and W. Lee et al. (2012) recorded CoV for preterms above 20%.

	Coefficient of Variability			% change between groups	
	Fetal	**Preterm**	**Preterm At Term**	**Fetal vs Preterm**	**Preterm vs Preterm At Term**
Temporal Lobe GM	13%	13%	8%	-33%	-22%
Temporal Lobe WM	8%	8%	15%	-20%	-11%
Frontal lobe GM	9%	13%	13%	-17%	-16%
Frontal lobe WM	11%	5%	8%	-5%	-10%
Parietal lobe GM	8%	13%	12%	-20%	-7%
Parietal lobe WM	12%	10%	5%	-4%	-6%
Occipital lobe GM	11%	17%	14%	-47%	-7%
Occipital lobe WM	12%	24%	20%	-45%	-3%
Limbic system	7%	5%	7%	-10%	-5%
Basal ganglia	10%	9%	9%	1%	-14%
Thalamus	11%	9%	8%	-10%	-11%
Corpus Callosum	13%	15%	13%	-14%	-7%
Cerebellum	8%	17%	25%	-41%	-7%
Brainstem	10%	6%	10%	-8%	-5%

Table 5.5: Coefficient of Variability has no apparent trend between regions or groups. The highest percentage change in mean T2* estimates is detected between fetal and preterm groups in the occipital lobe, temporal lobe and cerebellum, which is consistent with prior post-mortem studies of fetal development.

High CoV in occipital lobe in all of the previous studies, as well as the current study, suggests that the measure of T2* decay constants in this region can vary significantly from sample to sample. Therefore, a larger sample size in this region specifically would be beneficial for a more accurate estimate of average T2* decay constants in the future.

5.3.6 Correlation with age

Correlation of T2* decay constants with gestational age was evaluated via linear regression for estimates taken from all subjects combined (see Table 5.6). Non-zero correlation between T2* decay constants and gestational age was found in all regions of interest with p-values not exceeding 0.01 in any region. However, the R-squared values had not exceeded 0.25 in any regions except for basal ganglia ($R^2 = 0.51$) and thalamus ($R^2 = 0.46$). This suggests that, while the correlation between T2* decay constants and gestational age are significant, the measure of variance in T2* decay estimates is not well modelled by linear regression by gestational age

	ALL groups				Neonatal only			
	Slope	Std Error	R^2	p value	Slope	Std Error	R^2	p value
Temporal GM	-7.1	1.1	0.25	0.00	-2.3	1.4	0.25	0.03
Temporal WM	-4.6	0.9	0.03	0.00	-1.4	1.2	0.03	0.24
Frontal GM	-4.1	0.8	0.10	0.00	-1.9	1.1	0.10	0.13
Frontal WM	-2.7	0.7	0.13	0.00	-1.5	0.9	0.13	0.10
Parietal GM	-3.3	0.9	-0.05	0.00	-0.5	1.1	-0.05	0.62
Parietal WM	-2.4	0.8	0.24	0.01	-2.1	1.0	0.24	0.03
Occipital GM	-7.4	1.4	-0.04	0.00	-0.6	1.8	-0.04	0.51
Occipital WM	-7.9	1.7	-0.06	0.00	-0.6	2.2	-0.06	0.71
Limbic system	-2.1	0.5	0.09	0.00	-0.8	0.6	0.09	0.13
Basal ganglia	-2.5	0.5	0.51	0.00	-2.7	0.7	0.51	0.00
Thalamus	-3.1	0.5	0.46	0.00	-2.2	0.7	0.46	0.00
Corpus Callosum	-3.5	1.0	0.08	0.00	-1.8	1.3	0.08	0.16
Cerebellum	-5.2	1.1	-0.02	0.00	-0.9	1.4	-0.02	0.44
Brainstem	-1.8	0.5	0.08	0.00	-0.9	0.7	0.08	0.16

Table 5.6: Evaluation of correlation of T2* with age. R^2 and its p-value estimate indicate the strength and significance of the estimated correlation. The comparisons are made for all subject groups combined (ALL groups) and for neonatal subjects separately (Neonatal only). Only basal ganglia and thalamus show strong linear correlation of T2* estimates with gestational age in both comparisons.

alone. A separate analysis was performed with fetal subject excluded. This was premised by the observation that multiple regions of interest had shown significant difference in T2* estimates between fetal and preterm subjects while also sharing a similar gestational age. New results, which measured correlation for neonatal subjects only, showed significant correlation of T2* decay constants with age for only Temporal GM, Parietal WM, Basal Ganglia and Thalamus. Furthermore, only Basal Ganglia and Thalamus had sufficiently high $R2^2$ estimates.

5.4 Conclusions and Discussion

This work presents the first systematic study of brain T2* decay constants between fetal and newborn subjects using a comparable motion robust acquisition protocol and reconstruction technique. To the best of our knowledge, it is also the first report of T2* decay constants for both of these subject groups at 3T. Our results are consistent with previously observed T2* decay constant trends reported in individual studies of fetal and of preterm subjects at

1.5T in four regions of interest. We extend the previous observations into 11 more regions of interest in the brain. Finally, we introduce a fully automatic segmentation technique that allows measurement of T2* decay constants across whole regions of interest rather than on manually selected seed regions.

T2* weighted imaging is sensitive to size, density and orientation of myelinated fibers (Hwang et al. 2010b; Pitt et al. 2010). As a result, changes to T2* decay constants are an indicator for changes to myeloarchitecture (Cohen-Adad 2014). Our findings show that fetal T2* decay constants differ significantly from preterm born neonatal T2* decay constants at an equivalent mean gestational age in the cortical areas and cerebellum. This suggests that the transition from fetal to neonatal stage of life produces intense changes in the cortical and cerebellar organization in a very short period of time and has a non linear trend. This observation is particularly valuable as prior studies of fetal T2* decay constants showed strong linear correlation of T2* decay constants with gestational age in frontal white matter, occipital white matter, thalamus amd cortical gray matter, but no prior direct comparison had been made for the transition phase to the postnatal period.

Relative changes in T2* decay constants between fetal and postnatal transition phases are strongest in the occipital lobe, temporal lobe and cerebellum, which are responsible for sensory and motor functions of the brain. This observation is in line with prior conclusions from premortem studies by Kinney et al. (1988), which had outlined that brain maturation is highly asynchronous and projects from sensory regions to associative regions. Kinney et al. (1988) further outline that occipital lobe undergoes strongest changes, which is also confirmed in our findings.

Our study observes relatively high coefficient of variation (CoV) in T2* estimates across all subject groups, within the range of 10-20% for most regions of interest. This variation is similar to previously reported values at 1.5T in fetal and preterm subjects. The patterns of CoV do not have a homogeneous trend between the groups or regions which suggests no systematic effect on variation that may be caused by post-processing bias or by gestational age.

There was no significant difference in any regions of interest between preterm subjects and

preterm subjects scanned at term. These findings are consistent with prior study reported by Rivkin et al. (2004) in occipital white matter, frontal white matter and thalamus. Both Rivkin et al. (2004) and our results in this study differed from those of W. Lee et al. (2012) for occipital white matter, who reported a significant difference in T2* between subjects with a mean gestational age of 30 weeks of age and 43.5 weeks of age. We note that infants in W. Lee et al. (2012) study were significantly younger on average than the group in our study (34 GA) and in Rivkin et al. (2004) study (33 GA), which may explain the difference.

Furthermore, our analysis showed significant and strong correlation between T2* values and gestational age only in basal ganglia and thalamus. Neither Rivkin et al. (2004) nor W. Lee et al. (2012) had explicitly reported correlation of T2* with gestational age at 1.5T despite that each study had a similar number of subjects and a similar coverage of gestational age range if compared to our study. A careful observation of the reported values and graphs in those studies suggests that it is likely no significant correlation would have been detected due to high apparent variation in the T2* estimates between individual subjects. If those values would be reported as such, they would match our results as well.

5.4.1 Sensitivity to iron concentration

Interestingly, prior studies of T2 relaxation showed strong correlation of T2 estimates with age for preterm born infants. Counsell et al. (2003) measured 22 subjects across GA of 25-41 weeks and found significant correlation of T2 estimates with age in frontal white matter, occipital white matter and central white matter. However, we note that, unlike T2 weighted imaging, T2* weighted imaging is also sensitive to iron levels (Fukunaga et al. 2010) and macroscopic B0 field inhomogeneities (Hernando et al. 2012). There is strong indication from prior literature that the disparity between general trends with gestational age in T2 decay constants in Counsell et al. (2003) study, and those of our study for T2* decay constants, may be attributed to iron deficiencies in the preterm brain.

Notably, studies of preterm born neonates show that between 25-85% of preterm infants develop iron deficiency (Halliday et al. 1984; Rao et al. 2009). The large variance in occurence

is attributed to geographical and socio-economic factors. Results also show that the more premature infants are at higher risk of iron deficiency (Lundström et al. 1977). Finally, it is also observed that the deficiency in iron tends to level up and eventually normalize between preterm born and term born infants starting from 9 weeks post-natally (C. Zhu et al. 2019). The deficiency is thought to arise in preterms due to the fact that iron deposits in the body are mostly accumulated in the last weeks of normal fetal gestation (Seip et al. 1956). As a result, a premature birth rapidly depletes the normal distribution of iron stores, thereby making a newborn premature baby at risk of iron deficiency in the longer term. Iron metabolism has shown to have particularly significant effect on brain development in later stages of infancy, particularly in the regions of high myelination in hippocampus, visual system and auditory system (Georgieff et al. 2005; Thompson et al. 2001). Iron directly affects myelin production and oxidative metabolism (Stephenson et al. 2014; Todorich et al. 2009). Infants with iron deficiency test lower in tests for cognitive, motor and general neurophysiologic developments (Lozoff et al. 2006).

The aforementioned evidence suggests that the link between high incidence of iron deficiency in preterms, particularly as the iron deficiency tends to stabilize with respect to term born infants towards the later periods of post-natal development, may have a complex impact on the changes of T2* decay constants in the early stages of life. This may explain the disparity between the trends observed in T2 decay constants to those of T2* decay constants in preterm subjects. Furthermore, we suggest that due to high incidence of iron deficiency in preterm subjects it may also be the cause of high variation in T2* estimates that our study had observed between individual subjects, which in turn may have led to non significant differences recorded in ANOVA analysis between preterm and preterm at term subjects.

Nevertheless, two distinct regions of interest in our study showed low CoV and high r-squared values for linear model fits, thereby indicating strong correlation between T2* estimates and gestational age between preterm and preterm at term subjects. These regions are basal ganglia and thalamus. We note that basal ganglia is known to have the highest concentrations of iron in the adult brain (Griffiths et al. 1993). We argue that the iron stores in the basal ganglia may be sufficiently stable already in the very early gestational stages of development, which may explain

the lower variability in T2* estimates in that region. Similarly, for thalamus, multiple studies had shown significant correlation between changes in iron concentrations in the thalamus and various neurological pathologies (Hagemeier et al. 2013; Mitsumori et al. 2009), which suggests that there may be a strong interplay between regulation of thalamic iron and neurological development in the early stages of life as well.

However, localizations of iron in the developing brain is an area that is still poorly understood. To the best of our knowledge, there are no studies of brain iron concentrations for early preterm infants, although the correlation between preterm iron deficiencies from blood samples and its impact on cognitive development in the later stages of infancy are extensively studied. Furthermore, we also recognize that this study was not specifically designed to study iron levels in the brain and therefore the conclusion made from our results are only indicative. A future extension of the methods presented in this work to study the brain iron concentrations in the developing brain could be of valuable contribution. We further note that quantitative T2* imaging is routinely used as the optimal non-invasive method for measurement of iron concentration in vivo in adults, particularly in the brain, cardiac and liver tissue (Carneiro et al. 2006; Chavhan et al. 2009). However, as the brain undergoes intense changes in myeloarchitecture and water compartmentalization during late fetal and early neonatal development, quantitative T2* relaxometry alone would not serve as a good biomarker for iron concentration in the brain. It should be paired with another quantitative method that would help disentangle the gradient echo's sensitivity to changes in myeloarchitecture from senstivity to changes in iron.

5.4.2 Sensitivity to BOLD

Significant difference in T2* decay constants reported in this study at 3T between fetal and preterm subjects confirm that longer TE values should be considered for BOLD fMRI experiments for these subject groups. However, while it is generally considered that an optimal TE value should be set equal to the T2* decay constant for optimal contrast in the tissue of interest (Posse et al. 1999), this practice cannot be directly transferred to fetal and neonatal domains. Longer TE settings may lead to more pronounced motion artifacts in fetal scans and significant

irrecoverable signal losses from susceptibility artifacts in neonatal scans. Longer TEs will also prolong the scanning time, which may hamper the temporal resolution of fMRI studies, thereby decreasing its accuracy. Moreover, longer TEs will impact the signal to noise ratio, particularly for scans with higher resolution. As an example, the study of preterm fMRI by W. Lee et al. (2012) had to discard 20% of its data due to significant losses in signal quality. The TE in that study was set to 130ms, despite that the average T2* decay constant in the occipital lobe, for which a task based fMRI experiment was conducted, was measured at 180ms. Increasing the value of TE from 130ms towards 180ms would have resulted in further irrecoverable losses of signal from susceptibility effects rendering a larger portion of data unusable. Therefore, in conclusion, we recommend that each specific fMRI experiment should test against these factors while keeping a higher T2* decay constants in mind.

5.4.3 Variation

Importantly, our results showed that the variation between T2* decay constants in the cortical regions is much larger in fetal and preterm subjects than in adults. For example, we observed 37% difference between T2* decay constants in occipital and frontal gray matter in preterm subjects. Thus, careful considerations should be made when setting TE values with respect to a specific region of interest that is being imaged in task-based fMRI studies. For resting state fMRI, a multi-echo fMRI or a quantitative fMRI approaches may be of considerable advantage for improved accuracy.

Furthermore, our study also showed that T2* decay constants vary significantly more on an individual basis from subject to subject, than in adults. As a result, BOLD fMRI activation patterns may vary significantly from subject to subject from a confounding by TE alone. In an ideal setting, each subject's T2* decay constants should be evaluated during each individual scan and be taken into consideration during fMRI analysis.

5.4.4 Simultaneous multi slice imaging

A valuable direction for future work may be an investigation into simultaneous multi-slice (SMS) acquisition (D. Feinberg et al. 2010b; Moeller et al. 2010). This technique has gained a considerable attention in recent years thanks to its ability to significantly shorten the acquisition time on gradient echo scans. The effect of multi-echo GRE EPI scanning with SMS acquisition has not yet been studied. If successful, such optimization may allow for quantitative T2* mapping of the brain to be obtained with a much shorter timeframe, which could simplify acquisition of quantitative T2* maps in routine clinical examinations.

5.4.5 Segmentation of subjects below 30 GA

Another valuable extension to the current study would be a framework that performs automatic whole brain segmentation of T2* maps for subjects younger than 30 weeks of age. The decision to limit the minimum gestational age to 30 weeks in this study was made after a careful observation of thicknesses of the selected regions of interest, particularly those in the cortical GM, with respect to the chosen acquisition parameters, and the proposed reconstruction and segmentation frameworks. The simplest way to solve this challenge may be to consider acquisition at finer resolution levels. However, increased resolution would hamper the signal to noise ratios, thereby making signal losses from susceptibility artifacts even stronger in neonatal subjects in later echoes. Another solution may be to acquire multiple multi-echo acquisitions with shifted TE values in each scan, thereby obtaining a denser sampling of the signal decay during the early TEs. Any subject motion in this case would introduce irregular spatial coverage of the brain. Therefore, reconstructions would need to be made in 3D space by aligning data points from multiple acquisitions. Ultimately, such a method would combine the approach used in this study together with that of Blazejewska et al. (2017). However, as mentioned at the beginning of this chapter, an additional investigation into propagation of errors from data interpolation, and its effect on estimation of exponentially decaying signal, would be required. More importantly, we note that Blazejewska et al. (2017) acquired 100+ images for its 3D reconstructions. Obtaining a similar number of repeats for a multi-echo approach would incur a significant increase in

scanning time, which may be prohibitive within most clinical examination settings for fetal and neonatal subjects. Therefore, the aforementioned approach would need to consider a significantly smaller number of repeats in order to perform a successful 3D reconstruction. In general, a reliable slice to volume reconstruction requires at least 8 scans that cover the brain region in two orthogonal directions. However, even this limited number of repeats requires a significant time investment during a clinical examination. In the next chapter, we investigate the methods which could address this limitation.

Chapter 6

Slice to volume reconstruction for sparsely sampled image space

6.1 Abstract

The goal of this study was to investigate reconstruction of 3D brain images of a chosen contrast where acquisition of multiple orthogonal slice stacks is prohibitive due to high sensitivity to motion and prolonged scanning. To achieve this, we built a conditional generative adversarial network that can synthesize a prior for a chosen target contrast to serve as a seed for Slice to Volume Reconstruction (SVR). Such a prior is synthesized from a source image, where a 3D reconstructed volume would already be available. As a result, individually scattered slices from a sparsely sampled image space were combined with the synthetic volume of the same contrast to reconstruct a fully sampled 3D image.

6.2 Introduction

Single shot acquisition is an effective technique for acquisition of motion free slices of the moving fetus that are free from artefacts (Gholipour et al. 2014). Further processing is required to

view this data as an anatomical 3D volume. A generalized approach to 3D reconstruction is an iterative framework that formulates slice mapping as an inverse problem. Transformation starts with a set of temporally contiguous stacks of slices that are progressively split into smaller temporal blocks and are individually realigned to a reference image, which is itself updated at the end of each iteration (Ferrante et al. 2017). A final volume is formed by super-resolution reconstruction of individually aligned slices in a spatially irregular mesh that densely covers the 3D volume of interest.

However, a substantial part of clinical scans is not fit for 3D reconstruction. Each reconstruction necessitates dense over sampling of the image space. At least 6 scans in 2 orthogonal directions are required at the very minimum and 8 scans in 3 orthogonal planes are advised for a reliable reconstruction (Ferrante et al. 2017). Due time constraints, such scenario is often not feasible and many scans are left for 2D assessment instead. Furthermore, in clinical investigations an absolute priority must be given to the examination of the underlying pathology, thus making any scientific protocols secondary and often further constrained in the time that is available to acquire sufficient data for a 3D reconstruction.

To address this problem we propose a novel framework that combines slice to volume registration with image to image translation to obtain a 3D reconstruction of the whole brain. At first, a synthetic volume is generated for the image contrast where only sparse data had been acquired (target contrast). This synthetic volume is formed by image translation from another contrast (source contrast) where a fully formed 3 dimensional structure is available and has been reconstructed from a densely sampled image space. Motion corrupted slices from the acquisition of the target contrast are combined with this synthetic volume to estimate their correct position with respect to the subject's 3D coordinate space via slice to volume reconstruction algorithm. The proposed method is built for reconstruction of 3D images from data with highly limited samples in the target contrast, where 3 or less repeats with only a single view are acquired.

This particular study has chosen to focus on reconstruction of 3D T2* weighted images. However, similar approaches may be applicable to other types of contrasts, such as T1, DTI or field maps. T2 scans were chosen as the source contrast for this study due to their availability in

clinical practice, where T2 acquisitions of the fetal brain in multiple orthogonal planes are typically treated as the base requirement, and are therefore the most frequently available to be 3D reconstructed with the classic approach where the image space is densely sampled.

6.3 Literature Review

6.3.1 Generative Algorithms

To build the prior, an image to image translation model was considered. Several machine learning based methods have been introduced in the past to tackle this, amongst them a random forest approach (Huynh et al. 2015), k-Nearest-Neighbour algorithm (Zhong et al. 2016) and variational autoencoders (Rezende et al. 2014). Recently, a newer type of generative models has been able to surpass the performance of the previous methods by modelling probability density of the data in an unsupervised manner – Generative Adversarial Networks (GAN) (Goodfellow et al. 2014). Just like the other deep generative models, GANs model the underlying data distribution via the principle of maximum likelihood. The likelihood estimate that describes the training data can be written as $\prod_{i=1}^{N} p_{model}\left(x_i;\theta\right)$, where N is the total number of training samples of x_i with models' probability density function p_{model} which is parametrized by θ. The maximum likelihood is estimated by:

$$\theta^* = \arg\max_{\theta} \prod_{i=1}^{N} p_{\text{model}}\left(x_i;\theta\right) = \arg\max_{\theta} \sum_{i=1}^{N} \log p_{\text{ model}}\ \left(x_i;\theta\right) \tag{6.1}$$

This can be further simplified by a property of $\arg\max f(m) = \arg\max \log f(m)$ and the log product rule $\log(mn) = \log(m) + \log(n)$. We can further rewrite this by noting that maximizing the likelihood is equivalent to minimizing the Kullbeck-Leibler divergence between the real distribution of the data and the distribution of the model:

$$\theta^* = \arg\min_{\theta} D_{\text{KL}}\left(p_{\text{data}}(x) \| p_{\text{model}}(x;\theta)\right) \tag{6.2}$$

In other other words, the GAN is trying to estimate the maximum likelihood of probability density distribution of the training set that is parametrized by θ_g, which is defined by the weights in its network.

In practice one does not have the probability density model of the data p_{data} and only has access to a limited number of samples x_i, such as a subset of MRI images.

Unlike other generative models, such as variational autoencoders or pixelCNNs, which explicitly set a model of the probability density function to maximize, GANs estimate p_{model} in an indirect manner. A basic GAN consists of two neural networks that are trained simultaneously. The generator network G is tasked with transforming a random noise input z into a sample x by drawing from $p_{model}(z)$. The discriminator network D attempts to estimate the likelihood that a given sample comes from $p_{model}(z)$ or from p_{data}. Such approach can be treated as a game between two competing players, where G is trying to minimize the difference between p_{model} and p_{data} while D is trying to maximize its ability to distinguish between samples coming from p_{model} and p_{data}. The objective of each network can be modelled as:

$$\begin{aligned} \mathcal{L}_D &= \max_D \mathbb{E}_{x\sim p_{data}(x)} \log D\left(x\right) + \mathbb{E}_{z\sim p_{model}(z)} \log\left(1 - D\left(G\left(z\right)\right)\right) \\ \mathcal{L}_G &= \min_G \mathbb{E}_{z\sim p_{model}(z)} \log\left(1 - D\left(G\left(z\right)\right)\right) \end{aligned} \tag{6.3}$$

In practice, the input to G for a basic GAN is noise drawn from a uniform distribution that is propagated through a series of convolutional layers to generate a synthetic image of the same dimensions as the training set images. The input to D is switched between an image from the training set or a generated image from G. The input to D is propagated through a series of convolutional layers with a densely connected layer at the end and a sigmoid activation function. In effect, D is a binary classifier that calculates the probability that a given input comes from p_{data} or from p_{model}. The discriminator and generator are trained in turns. The loss from training D is shared with G, thereby giving the generator an ability to sample probability density of the training set in an indirect manner and thereby train to estimate the maximum likelihood.

The nature of GAN formulation as a zero sum non-cooperative game between two players renders it as a saddle point optimization problem. As a consequence, GANs are exceptionally difficult to train and suffer from a number of pitfalls such as non-convergence, mode collapse, vanishing gradient problem and significant sensitivity to hyper-parameter selections (Mescheder et al. 2017). A large number of studies have attempted to tackle these issues. Notably, the original GAN paper by Goodfellow et al. (2014) proposed an alternative strategy to model the training criterion as minimization of the Jenson-Shannon divergence (JSD). Arjovsky et al. (2017) proposed new a loss function based on earth mover's distance criterion, called Wasserstein GAN (WGAN). Other formulations of loss functions, such as LSGAN (Mao et al. 2017), WGAN-GP (Gulrajani et al. 2017), and architecture, such as DCGAN (Radford et al. 2015), BEGAN (Berthelot et al. 2017), achieved partial success in resolving the aforementioned issues with GAN training. Nevertheless, a large scale study by Lucic et al. (2018) concluded that all of the aforementioned models reach similar, non statistically significant measure of quality for the generated images when tested across multiple datasets and multiple quantitative evaluation metrics. More importantly, the study also highlighted that thorough optimization of hyper-parameters is the key determining factor for a well trained GAN. A recent follow up study by Kurach et al. (2019) further noted that marginal improvements to sample quality can be achieved in the ResNet style architectures, which we review in the follow up section.

6.3.2 Conditional networks

This section focuses on GANs in the conditional setting such that the generated samples x are directed by a certain input y. Mathematically, the conditional training criterion can be re-written as:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}} \log D(x|y) + \mathbb{E}_{z\sim p_{model}(z)} \log\left(1 - D\left(G\left(z|y\right)\right)\right) \tag{6.4}$$

For image generation, the conditional input y is concatenated to x as input to G (Mirza et al. 2014). Isola et al. (2017) introduced a GAN tailored to conditional image priors, termed pix2pix.

A standard pix2pix implementation includes a U-net architecture (Ronneberger et al. 2015) that is built from stacked encoder-decoder convolutional layers in the generator, which ensures efficient compression of the data into a narrow feature space, and therefore further constrains the GAN during two-step minimax training process. Skip connections, originally introduced by the ResNet architecture (K. He et al. 2016), are added between same size encoder-decoder blocks to allow for low level information to pass through the network without weight multiplications, if necessary. The discriminator of pix2pix is formed from encoder only blocks that splits the original input space into patches of probability estimations for the authenticity of the input image, termed PatchGAN. Such approach encourages enhanced performance of the discriminator as the output is an average over a restricted set of receptive fields, each of which looks at high frequency components of a given image part.

Variants of pix2pix framework have been used for medical image translation from MR to MRA (Olut et al. 2018), PET to MR (Choi et al. 2018), PET to CT(Armanious et al. 2018), MR to CT (Emami et al. 2018), T1 to FLAIR (Yu et al. 2018), and T1 to T2 (Q. Yang et al. 2018) (Dar et al. 2019).

Despite the effectiveness of pix2pix framework for image to image translation tasks, many real world datasets, including those in medical imaging, do not have paired training data available. J.-Y. Zhu et al. (2017) addressed this issue with a new approach termed CycleGAN, which consists of two simultaneously trained concordant networks that attempt to find a representation of image $x \in X$ as $G_x(x) \rightarrow y$, while also taking $y \in Y$ and training a network for $G_y(y) \rightarrow x$. To regularize consistency between these networks an additional loss is shared between them, called cycle-consistency loss

$$\mathcal{L}_{\text{cyc}} = \mathbb{E}_{x \sim p_{\text{ data }}(x)} \|G_y(G_x(x)) - x\|_1 + \mathbb{E}_{y \sim p_{\text{ data }}(y)} \|G_x(G_y(y)) - y\|_1 \tag{6.5}$$

The networks are connected and form a backwards and forwards pass that allows for a two-way synthesis between images and thus stabilizes the network performance despite the lack of paired data. A complete objective function is given as combined loss of the forward GAN network,

which transforms x into y, backwards networks, and a cycle consistency loss

$$\mathcal{L}_{total} = \mathcal{L}_{\text{GANx}} + \mathcal{L}_{\text{GANy}} + \lambda \mathcal{L}_{\text{cyc}} \tag{6.6}$$

CycleGAN had been tested in medical imaging tasks such as synthesis between MR and CT (Wolterink et al. 2017), and T1 and T2 (Dar et al. 2019)(Welander et al. 2018).

6.3.3 Preserving high and low frequency information

Image to image translation GANs that rely solely on the adversarial loss for optimization of the weights do not achieve high fidelity of the reconstructed images as no constraints exist for preservation of the low frequency content between the images. Therefore, a large number of MR focused image translation GANs use some of form of regularization that forces the global structural consistency of the generated image (Creswell et al. 2018). The most common approach is to implement the adversarial loss with a weighted L1 loss over the pixel space of the image, such as a MAE.

Despite this, a highly weighted L1 loss tends to produce blurry results (Armanious et al. 2018). Moreover, for high resolution images, MAE fails to capture perceptual similarity between two images, particularly for small overall shifts between pixel space, that could be a byproduct of registration errors, for example. To counteract this, Armanious et al. (2018) used an additional regularizer that estimates the L1 loss of the extracted feature representations between the images, rather than the images themselves. It uses a common trick to build such an extractor by utilizing the discriminator architecture, which feeds into the regularized loss of the generator.

Further regularization may include an additional term known as 'style transfer' in the GAN literature. The goal of style transfer is to preserve fine grained high frequency information of the target image while transferring the low level information from the source. High frequency components are particularly important if the synthesized images will be further processed by other networks for registration and segmentation purposes. Similarly to the construction of the perceptual loss, high level features are extracted from the hidden layers of a deep CNN.

Armanious et al. (2018) use a VGG-19 network (Simonyan et al. 2014) to this purpose, which has been pre-trained on the ImageNet dataset (J. Deng et al. 2009). Despite the non medical nature of the original ImageNet dataset, it has been commonly used as a feature extractor for a variety of medical imaging applications. The key argument for its effectiveness is the considerable size of the ImageNet database on which the network was trained, which is orders of magnitude larger than any existing medical imaging repository. Due to this, a standard pre-trained VGG-19 is able to extract fine grained high level features from any image. Finally, to prevent overfitting on the dataset, additional regularization may be used. Q. Yang et al. (2018) uses several dropout layers for this.

Image translation for fetal MRI

The pretext to our study was the work of Salehi et al. (2018) at Boston Children's Hospital. Salehi et al. (2018) proposed a CNN based pose estimation network for arbitrary 3D rotation and translation of brain slices into a common learned atlas, tailored to neonatal and fetal applications. The network was trained on T2-weighted pediatric scans and have generalized poorly when trialed as a pose estimator on motion corrupted T1 slices. To circumvent this, in an auxiliary experiment, a T1 to T2 image translation GAN was built to synthesize T2 slices from corresponding T1 slices. Synthetic T2 slices were fed into the pose estimator to obtain transformation parameters for each slice in canonical 3D space. Translation of T1 images into synthetic T2 slices substantially reduced the error in estimation of slice position in canonical 3D space with respect to the ground truth, thereby demonstrating that it may be possible to apply a T2 trained pose estimation network to another contrast by using image to image translation techniques.

6.4 Method

A high level overview of the proposed approach is described in Figure 6.1. A high resolution T2 volume is constructed from densely sampled image space. A pre-trained conditional GAN is used

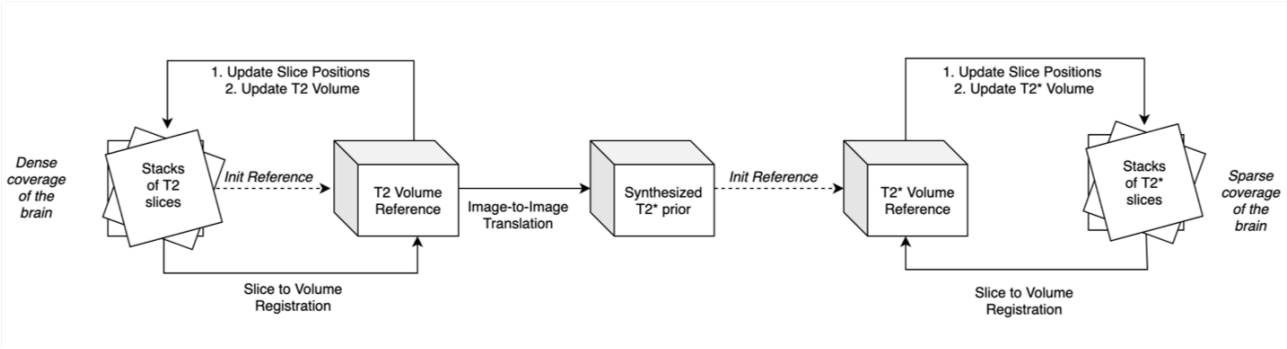

Figure 6.1: An overview of the proposed approach that includes a combination of slice to volume from real acquired samples and synthesized samples from another image contrast with sparse coverage of the brain.

to translate each slice in the T2 volume to a synthesized T2* slice. The slices are concatenated into a volume. Individual motion corrupted slices from various slices stacks are combined with the synthesized T2* volume by way of slice to volume registration.

6.4.1 Data

Data used in this work is taken from the neonatal subset of the Developing Human Connectome Project (dHCP) (Edwards et al. 2019), which contains 590 scans covering gestational ages between 20 and 45 weeks post conception from healthy neonatal or neonatal preterm subjects. Each scan has structural (T1 and T2), resting state functional and diffusion data. In this work, we select structural T2 and resting state fMRI T2*W images.

The images were obtained with a Philips 3T scanner using 32 channel neonatal coil. T2 images were acquired with a Turbo Spin Echo sequence with overlapping slices at 0.8 mm^3 resolution. The images are bias corrected, brain extracted and motion corrected using the work of Kuklisova-Murgasova et al. (2012) and Cordero-Grande et al. (2018), resulting in 0.5 mm^3 isotropic resolution. The final images are segmented by tissue type and label type using the Draw-EM algorithm (Makropoulos et al. 2014). T2*W images are acquired with a multi-shot multi-band EPI GRE sequence which is distortion corrected and motion corrected with final resolution at 2.16 mm^3 with FOV of 67x67 and 45 slices.

The following pre-processing steps were necessary in this study in order to make the data usable

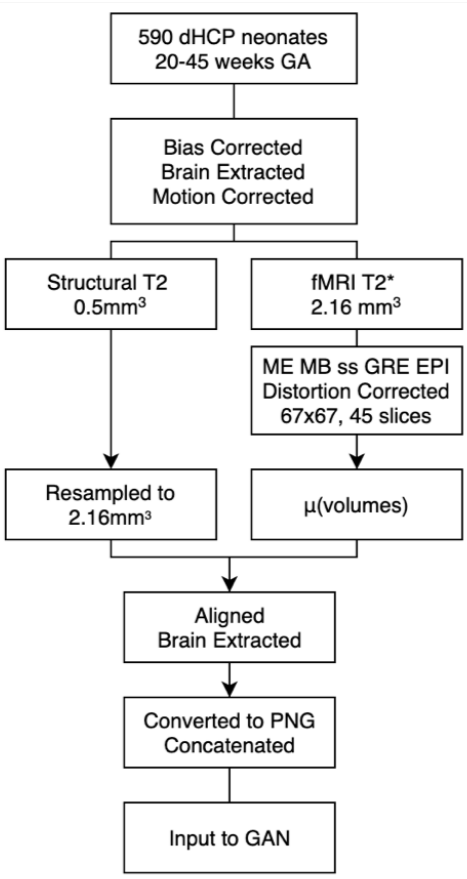

Figure 6.2: An overview of the pre-processing pipeline used to prepare the input images for ingestion by the image to image translation network.

for the target application (see Figure 6.2). T2 structural scans are aligned to T2* data using the affine matrix information and resliced to the same resolution of 2.16 mm^3. To remove any fluctuations due to the BOLD effect, an average image from the motion corrected fMRI data is used. Both images are masked to remove non brain tissue. The images are then converted into PNG files on a slice by slice basis, and concatenated together to form the input for the GAN based neural network. All the images that contain no data or the very edge of the brain are removed to ensure consisent input to the neural network, resulting in 37 slices for each subject. The training set is set to 550 subjects and test set is 40 subjects, which coincides with 20,350 and 1,480 images respectively.

The target dataset for this project are registered pairs of neonatal images from the Developing Human Connectome Project (dHCP). The use of neonatal scans, which are fully reconstructed with a matching template, allows us to quantify performance of the proposed image to image

translation and super resolution framework. Subsequent testing on the fetal part of the dataset will be performed when the data is fully released, which is currently expected to be in Q4 2019 - Q1 2020.

6.4.2 GAN Architecture

Following the arguments presented in the literature review section, our choice for the generative network was driven by the following considerations. First, we note from Lucic et al. (2018) that most variants of generative models achieve similar scores with sufficient amount of hyper-parameters optimization, and the only improvements are shown for networks that use U-net (Ronneberger et al. 2015) and resnet (K. He et al. 2016) architectural features. As a result, we select the base architecture introduced by Isola et al. (2017), which is also the most widely used GAN in medical imaging applications for image to image translation (Kazeminia et al. 2018; Yi et al. 2019). It is further noted that Isola et al. (2017) uses L1 regularization for its objective function to address the issue of high frequency preservation in the generated images. Our choice of the generative architecture is also driven by the results obtained by Salehi et al. (2018), which showed significant and almost equal performance for the CNN based pose estimator when using the same architecture for T2 to T1 translation in fetal images. We note further that our task has a smaller computational complexity in comparison to other image to image translation medical imaging tasks thanks to the fact that both fetal and neonatal images are significantly smaller than their adult counterparts, which may have had necessitated more complex or highly regularized networks with additional constraints. Finally, the last remark relates specifically to our chosen data. Unlike the T2 to T1 translation task, posed by Salehi et al. (2018), the data in this study is of much lower resolution due to the inherent low resolution of T2* weighted images. As a result, the computational complexity costs decreases significantly in this specific study. The combination of all these factors suggests that a standard pix2pix framework should deliver promising results as long as its hyperparameters are sufficiently tuned.

The algorithm adopts the work of Isola et al. (2017) for image to image translation with 16 encoder/decoder blocks in the generator and 5 encoder blocks in the discriminator, which follow

the U-net architecture that was originally designed for biomedical applications (Ronneberger et al. 2015) . The input image is fed through the first encoder block. Convolutional filters are progressively increased in sequence to the outputs of the previous layer until the data reaches the bottleneck layer which is a 1x1x512 latent vector, and is then progressively decoded back to the original resolution. Each encoder layer is connected to a corresponding decoder layer with skip connections.

The discriminator uses overlapping patches that are applied to the image input and pass through a sigmoid function to generate the probability that a given patch represents true data. The probabilities are averaged to produce a single likelihood that a given input image comes from real data distribution.

In our case, the generator receives as an input a T2 slice as a source image and outputs a synthetic T2* slice target. T2 slice is fed into the U-net based generator. The discriminator compares between real T2* slices and synthesized T2* slices to assess authenticity.

Under the assumption that images of the source and target domains are perfectly registered, a simple pixel-wise loss (MAE) was used in combination with the GAN loss in the objective function in a weighted fashion.

6.4.3 Motion simulation

To measure performance under controlled conditions, two approaches for motion simulation were designed.

Target Registration Error test

The first approach used a set of preselected transformation parameters that were applied iteratively to a subset of the subject data to each individual slice (see Figure 6.3). This enabled a test of consistency of performance of the same translation or rotation combinations across all slices and a measure of deviation by spatial location. Individually transformed slices were

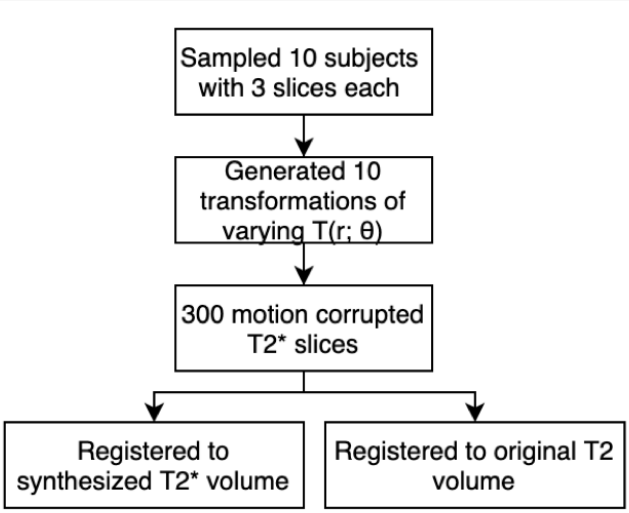

Figure 6.3: An overview of the processing steps used in the Target Registration Error test.

registered to two target volumes - the original T2 weighted volumes and the synthesized T2* weighted volumes. The aim of this test was to compute a target registration error (TRE) when registering a motion corrupted slice to a 3D volume. Target registration error (TRE) is computed for each slice to volume registration by comparing the output transformations with the original simulated translation and rotation vectors as ground truth. Formally, we note:

$$\mathrm{TRE}_i = \left\| P_i - T_{\mathrm{reg}}^{-1}\left(T_{\mathrm{ori}}\left(P_i\right)\right)\right\|^2$$

where P_i are the coordinates of the i-th voxel that is located on the vertex of a rectangle that covers the slice of interest, T_{reg} is the estimated rigid transformation, T_{ori} is the original rigid transformation applied to the given slice where P_i is located.

To generate motion, a subset of 10 subjects were randomly selected from the dHCP dataset. For each subject, a ground truth T2* weighted volume was available. A set of 10 translation and rotation vectors were applied to each volume on a varying scale. The first combination involved simple rotation in 2D plane, and the most complex included rotation and translation for all 6 degrees of freedom. The extent of motion was kept within a limit of motion that is typically observed in fetal MRI scans. For this, references were taken from graphs presented by (Ferrazzi et al. 2014; Jiang et al. 2007). Parameters used for simulating motion corrupted slices are listed in Table 6.1. Once the motion was simulated, the volumes were split into slices in

	Translation			Rotation		
	x	y	z	x	y	z
1	0	0	0	0	1	-5
2	0	0	0	-4	6	3
3	-2	3	0	0	0	0
4	-2	4.5	2.5	0	0	0
5	-2	-2.5	0	0	3	-2
6	2	-2.5	-1	-1	3	2
7	-2.5	2	0	-1	1	1
8	2.5	3.5	0	0	1	2
9	3	-2	-0.5	-3	2	-2
10	-1.5	2	-2.5	2	2	-2

Table 6.1: Motion simulation parameters for translation (in mm) and rotation (in degrees) for ten different scenarios.

the transverse plane. Three slices were selected from each slice stack - bottom, middle and top parts. In total, this approach generated 300 motion corrupted slices, which included 10 subjects, 10 parametrized transformations and 3 slices.

For each of the 10 subjects, a T2 volume and a corresponding synthesized T2* volume were also available. Each motion corrupted slice was therefore registered to the T2 and the synthesized T2* volume. Rigid registration and transformation was performed with IRTK software (Studholme et al. 1999). Registration was optimized for improved performance with the following parameters: 1 resolution level, 200 iterations, padding value of -1, NMI (T2 target volume) and CC (T2* target volume). All other parameters for rigid registration were kept as default.

Simulated Acquisition test

The second approach applied motion to each subject's T2* weighted volume randomly by simulating a fetus that is moving in a scanner while the data is being acquired sequentially along the slice selective gradient in the z-axis. Transformation vectors that were generated for each slice were capped at maximum translation and rotation values that were defined from observations in (Ferrazzi et al. 2014; Jiang et al. 2007). This approach was applied in all three orthogonal views of each T2* weighted volume multiple times. The result was a set of multiple motion corrupted slice stacks in three orthogonal views that represented individual repeated acquisitions of the fetal brain in the presence of motion. An overview of this Simulated

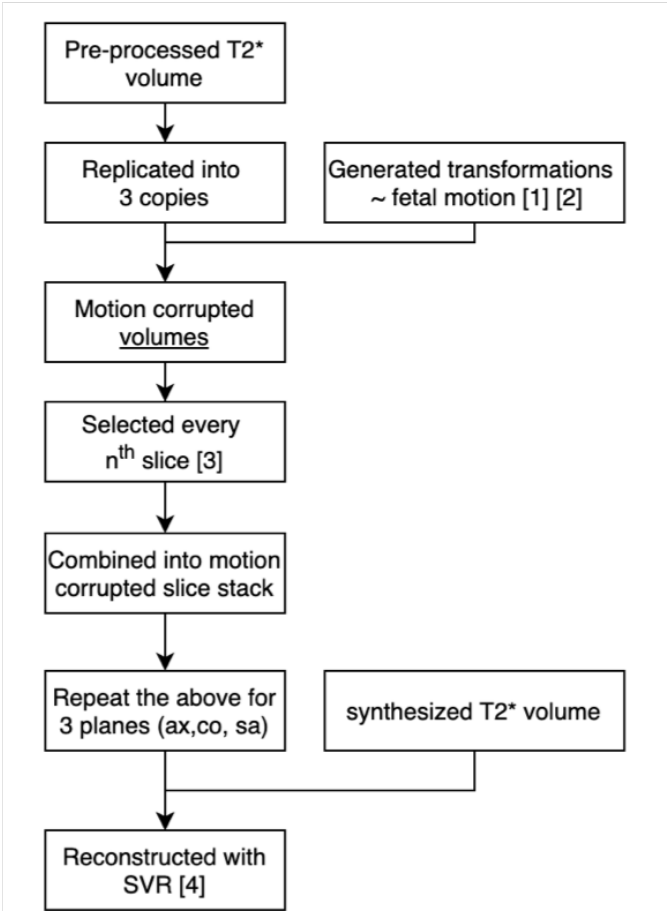

Figure 6.4: An overview of the processing steps in the Simulated Acquisition test. Simulated rigid transformations for each slice were not larger than the fetal motion parameters measured and reported by [1] Jiang et al. (2007) and [2] Ferrazzi et al. (2014)

Acquisition test framework is shown in Figure 6.4.

These stacks would be used in conjunction with a synthesized T2* volume during the reconstruction stage via SVR.

6.4.4 Reconstruction

Reconstruction was performed with the SVR algorithm presented by Kainz et al. (2015), described in sections 3.7 and 5.2.4. Individually synthesized T2* weighted slices were combined into a slice stack, thus forming a synthetic 3D volume of the brain. Additional slice stacks consisted from motion corrupted real T2* weighted slices that were constructed with motion simulation framework, as described in section 6.4.3. All of the data was fed through the SVR

pipeline.

6.5 Results

6.5.1 Network parameters

To optimize performance of the chosen architecture a variety of hyperparameters were evaluated by training the network and testing against the test set. The test set consisted of 40 subjects of varying gestational ages. To assess performance, 5 different image similarity metrics were considered for comparison between the synthesized T2* slice output from the generator and the corresponding ground truth T2* slice from the original data - mean squared error (MSE), normalized root mean squared error (NRMSE), structural similarity image metric (SSIM), peak signal to noise ratio (PSNR) and learned perceptual patch similarity metric (LPIPS) (R. Zhang et al. 2018). LPIPS is based on a pretrained VGG network for ImageNet classification and has shown to produce closer fidelity to human judgement than the classic metrics. To evaluate LPIPS against other metrics on MR image generation, 100 individual pairs of slices were assessed for quality by a human observer against 5 different variants of pix2pix models that were trained with varying hyperparameters. It was found that both LPIPS and SSIM performed the best out of all 5 metrics, and equally well between themselves, with respect to the human observation of fidelity between the images when selecting the image pair from the 5 variants of the pretrained pix2pix model. Therefore, it was decided to proceed with SSIM as the preferred metric due to its wider recognition within the literature. Next a much more comprehensive test was performed to tune the network across a variety of parameters.

To assess basic performance of the network, we tested three configurations for the number of filters in the first convolutional layer of the generator and discriminator networks. Some improvement was found for the increase in the number of filters of the discriminator network, as the network had larger capacity at both generating and discriminating synthesized images. Different types of GAN objective function were tested and it was found that cross-entropy

Hyperparameter	Value
Discriminator architecture	PatchGan
Generator architecture	U-net
GAN loss	Cross entropy
Optimizer learning rate	0.0002
Optimizer momentum	0.5
Layer normalization	Batch Norm
Regularization	Dropout
G filters in last layer	64
D filters in 1st layer	64
Batch size	1

Table 6.2: Final hyperparameters of the proposed neural network for synthesizing T2* slices from T2 slices.

objective function had performed the best out of all. PatchGAN design for the discriminator had shown improved performance over a simple 'pixel' level discriminator that takes the output over the whole image. A generator framework for a resnet 9 blocks had scored worse than a 256 U-net architecture. Comparison between batch normalization and instance normalization had yielded approximately similar results. The networks was also tuned with the respect to the hyperparameters of the adam optimizer - its momentum value and learning rate. The biggest performance improvement overall was seen for the size of the dataset, with maximal number of images scoring the best SSIM improvement out of all optimizations. This observation is in line with many prior works of literature which suggest that very large datasets help generalize better for most adversarial training tasks. This is further reinforced by the improved result yielded from adding random left-right flip during training to the training data.

The final training was run for 200 epochs with minibatch size of 1, adam optimizer with beta value of 0.05 and learning rate of 0.0002 (see Table 6.2). The weights were initialized using normal distribution and batch normalization was used between encoder/decoder blocks. The generator architecture consisted of a 256 U-net and the discriminator included a standard patchGAN of the selected framework. Additional left-right 'flips' were used to augment the data as per improved performance. The number of first and last filters set in the convolution layers were 64 for discriminator and generator respectively.

Figure 6.5 shows the generated synthetic T2* slices, which are compared to the ground truth T2*

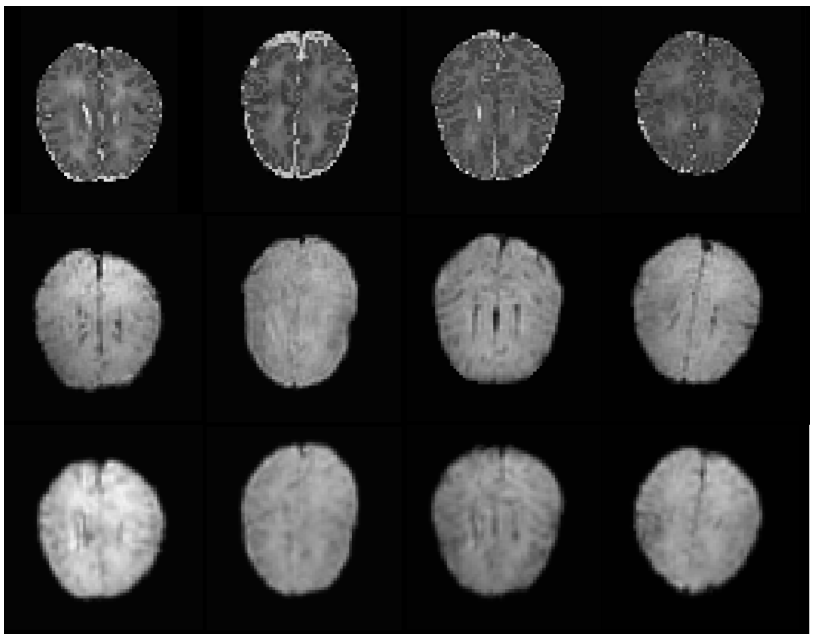

Figure 6.5: Examples of synthesized T2* slices from 4 different subjects. Top row: original T2 weighted slice. Middle row: original T2* weighted slice. Bottom row: synthesized T2* weighted slice from the T2 weighted slice.

weighted slices as well as the corresponding T2 slices that were used as a source for image for image translation. Note that the original T2* weighted data had a FOV of 67x67 with inplane resolution of 2.16 mm^2, thus the translated data has very low resolution also. Noticeably, the synthesized slices are blurrier than the ground truth T2* images. Figure 6.6 shows an example of the synthesized T2* slices concatenated together in two planes that are orthogonal to the plane of synthesis. An outline of the ground truth image is plotted over the generated image in yellow, and vice versa. The data shows some spatial inconsistency in the through slice direction as the loss function was not regularized in this direction. Additionally, synthesized slices are uncorrelated in contrast. However, this uncorrelation in contrast is automatically corrected via the SVR framework from Kainz et al. (2015), as a strong bias field effect.

6.5.2 Target Registration test

Table 6.3 shows mean registration accuracy between individual motion corrupted slices and two volumes that were used as targets during the registration. Results are split by category and are as follows. There were three different types of slices selected and transformed by the

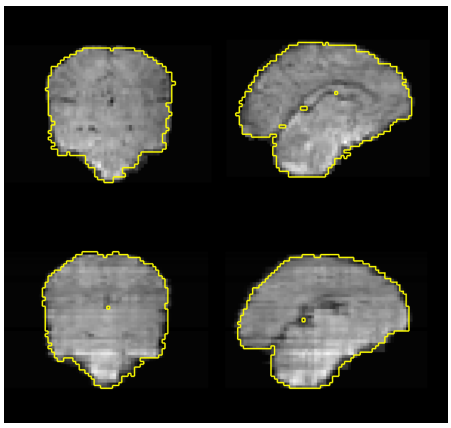

Figure 6.6: A visual comparison between T2* synthesized volume and the original ground truth T2* volume, shown in two planes that are orthogonal to the plane in which the T2* slices were synthesized from the T2 target slices. Top row: Original ground truth T2* volume. Yellow outline are the edges of the synthesized T2* volume. Bottom row: Synthesized T2* slices concatenated into a slice stack. Yellow outline are the edges of the ground truth T2* volume.

	All		by slice	by dof	by subject
Target	μ	σ	σ	σ	σ
T2	2.40	2.53	2.09	3.17	1.98
sT2*	3.05	3.19	2.61	4.03	2.36

Table 6.3: Target registration error (TRE) for motion simulated slices from 10 different subjects, across 3 locations in the brain and 10 different transformation parameters (denoted as dof). The first target volume for registration is the ground truth T2 volume. The second target volume is the synthesized T2* volume (sT2*). Mean and standard deviation are reported in mm. TRE was computed by comparing the estimated location of the slice with regards to its original location in the ground truth T2* volume.

same parameters across all subjects. Additionally, there were 10 different combinations of transformation parameters that simulated rigid motion with 6 degrees of freedom. Finally, there were 10 subjects that were tested. Notably, the standard deviation is highest amongst the transformation parameters group. Conversely, the variation between accuracy of registration between slice choices or subject choices is significantly smaller for both target volumes. The accuracy of the alignment was on average higher for T2 data, which is expected as the volume is fully consistent in 3D. In contrast, the synthetic T2* volumes that served as targets had on average 27% higher error than the T2 volume.

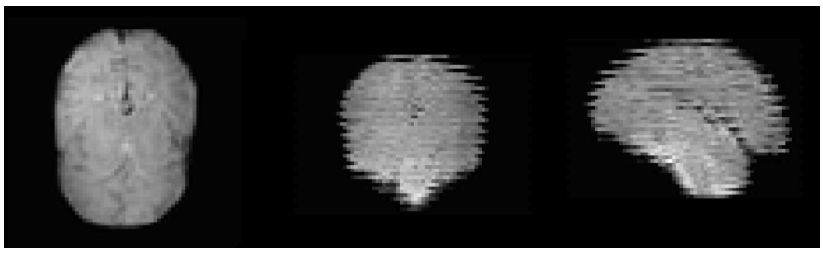

Figure 6.7: An example of a slice stack that has been motion corrupted with a simulation. The slices in this example were rotated by 1 degree and translated by 1mm in one orthogonal plane with a 3 slice interleaved slice acquisition model.

6.5.3 Simulated Acquisition test

Figure 6.7 shows an example of a T2* volume that has been processed with an algorithm that simulated fetal motion in the scanner. This process was repeated multiple times by adding a random gaussian noise to the generation procedure with varying strength of motion. Each repeat was treated as a separate acquisition and added as a stack of slices to a slice to volume reconstruction pipeline.

6.5.4 Reconstruction

To aid in evaluation, three reference reconstructions were built from motion corrupted slices. The first reference was constructed from 9 motion corrupted T2* weighted slice stacks "acquired" in three orthogonal planes with 3 repeats in each plane. Such setup is considered to yield the best accuracy of reconstruction in the presence of motion for a standard SVR pipeline. A second reference consistent of 3 slice stacks "acquired" in the transverse orientation and 3 slice stacks in the coronal orientation. This kind of acquisition constitutes the baseline requirement for fair accuracy of the SVR framework in the clinical setting. The third reference reconstruction was built from a sparsely sampled image space with only 3 repeats "acquired" in the transverse plane.

Mean registration accuracy of individual slices to the ground truth T2* volume for each of the reconstruction types across all subjects is listed in table 6.4. Despite that the initial motion corruption reached up to 5 mm in translation and 5 degrees in rotation, the registered accuracy

	μ	σ
9 stacks	1.82	0.93
6 stacks	2.06	0.97
3 stacks + GAN	2.54	1.74
3 stacks	2.89	1.21

Table 6.4: Mean target registration accuracy (TRE) calculated over all motion corrupted slices that were registered back to the ground truth T2* volume via Slice to Volume Reconstruction (SVR) algorithm. The comparison is made between different types of input to the SVR algorithm. Mean and standard deviation are reported in mm.

for reference reconstructions was within the range of the original resolution for densely sampled data (less than 2.15mm). While registration accuracy for sparsely sampled image space was significantly lower, the result has shown a 20% improvement when the reconstruction included a synthesized T2* volume. A visual example of the reconstruction results is shown in Figure 6.8. It is noted that while the structural fidelity is well preserved in general, some blurring is present, even in the reference reconstructions. The primary causes behind this is likely to arise from a relatively coarse resolution of the data and poor contrast. As was noted in section 3.7, Jiang et al. (2007) were able to achieve a higher quality reconstruction with less overall blurring than present in Rousseau et al. (2006) by acquiring significantly thinner slices of 1mm. We further note from Figure 6.5 that the T2* weighted data has a poorer tissue contrast in comparison to the T2 weighted data, which would play an effect on the performance of the registration in the presence of large motion.

Reference reconstructions from densely sampled image space that consisted from 9 and 6 slice stacks respectively produce similar results, as shown in Figure 6.8. The performance is significantly reduced for a reconstruction with 3 slice stacks in a single plane of view. By using a synthesized T2* volume as an additional slice stack, the reconstruction is noticeably improved. In particular, the transverse view shows that a correct placement of the slice had been made, which indicates better structural coherence. However, some variation is present in the transverse view of the reconstruction that used a synthesized T2* volume. For example, the frontal left region of the brain in the transverse view has "less dephasing" present. This is likely to arise from spatial inconsistency between the synthesized T2* volume and the ground truth T2* volume.

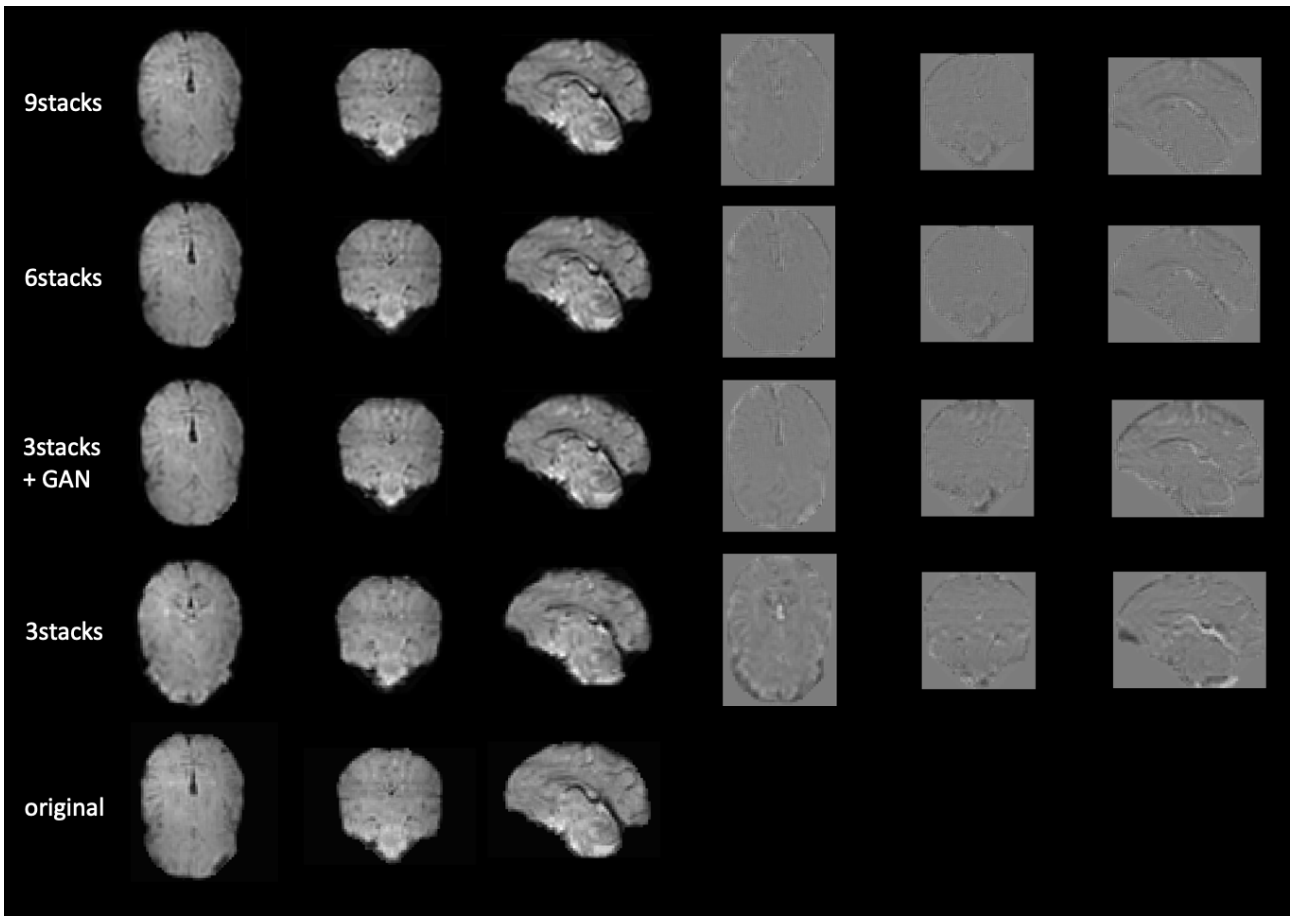

Figure 6.8: Slice to volume reconstruction results for a single subject with varying levels of image data available to each reconstruction. Results are presented as a view in 3 orthogonal planes. A difference image is generated by comparing a given reconstruction result with the original T2* volume prior to motion simulation.

The spatial consistency of the synthesized volume can be improved by regularizing the objective loss function of the GAN in the through-slice direction. This could be made possible by performing image to image translation in 3D. An alternate strategy may be to condition each slice by its neighbours or by a segmentation mask that is derived from the T2 and T2* data. Furthermore, the synthesized slices are noticeably blurrier than the ground truth image. An improved performance over finer details can be reached by using an additional regularization, such as the style transfer loss presented by Armanious et al. (2018). Alternatively, the network may be trained with a multi-resolution approach, termed pix2pixHD, such as the one presented in T.-C. Wang et al. (2018). A multi-resolution approach is likely to yield better fidelity of the translated images overall, while yielding high resolution.

6.6 Discussion and Conclusion

This study proposed a new framework for reconstruction of 3D brain anatomy from a moving fetus where the acquired data is highly sparse. To achieve this, we utilize the prior knowledge that is available in the form of a consistent 3D brain image from another contrast from the same subject. A synthetic image is generated from the prior by using a pre-trained conditional generative adversarial network. The synthetic volume is combined with individually scattered motion corrupted slices and fed through a slice to volume reconstruction framework to build a representation of 3D brain anatomy.

To the best of our knowledge, this is the first method that attempts to perform slice to volume reconstruction with highly sparse data. The current state-of-the-art methods for slice to volume reconstruction require dense and highly oversampled image data. This creates substantial limitations in the clinical pipeline for what types of images can be reconstructed in 3D due to the length of the required acquisitions. As a result, most scans are left for 2D assessment. While the outcomes of this study are not yet comparable to clinically acceptable reconstruction quality, the proposed framework is the first step towards that goal.

Our choice of the generative adversarial network, pix2pix, was partially motivated by its wide and successful adoption within the medical image community (Kazeminia et al. 2018; Yi et al. 2019) as well as the relative computational simplicity of the chosen task due to the low resolution and high similarity between T2 and T2* contrasts. However, the performance of pix2pix comes short at generating slices with sufficient spatial consistency in the through-slice direction. There are several strategies that may be used to tackle this issue. The first may look into generation of whole 3D structures. The second may consider conditional input that is constrained by the neighbouring slices. The third may choose to constrain the generated slices by using a mask or a segmentation that's derived from T2 and T2* weighted data at training time. We discuss the benefits and constraints of these approaches below as methods that may be utilized in future study to improve our results.

6.6.1 3D GANs

Despite the inherent 3D nature of MR, CT and PET imaging data, an overwhelming majority of applications of deep learning in medical imaging choose to model the data on a slice by slice basis. The core challenge lies in the computational complexity that's associated with the added dimension, which leads to gradient saturation as well as significantly prolonged training time. The gradient saturation is particularly problematic as the GAN essentially stops improving itself after a certain number of epochs and does not yield good results. However, exclusion of the 3D spatial relationships in GANs often leads to discontinuities across slices as they are concatenated post analysis, as seen in this study also. This problem remains to be one of the key challenges in medical imaging community. Only three studies to date (as of September 2019) have experimented with 3D structures for GANs for image to image translation in medical imaging.

The first study by Yu et al. (2018), introduced a 3D conditional GAN with U-net architecture for the generator and a CNN based discriminator. Its output is a 128x128x128 image, that's used for T1 to FLAIR translation to aid in segmentation of T1 images. The network was regularized by a weighted combination of images from the target training patches with the synthetically generated target image from the generator. The contribution of target patches is based on their L2 norm similarity to the synthetic sample. The study explicitly stated that such fusion introduces artifacts into the generated image but quantitatively improved the resulting segmentation.

Conversely, Nie et al. (2017) proposed a Fully Connected Network (FCN) with added residual connections for the generator and a CNN for its discriminator to perform MR to CT image translations. In contrast to its CNN counterparts, FCNs yield better structural coherence of the patched inputs, but may result in larger blur artifacts. To counteract this, Nie et al. (2018) proposed a gradient based image difference regularizer for the generator. The network is trained in a patch to patch manner, thereby reducing the 3D space substantially and the corresponding computational complexity. The individual outputs of small synthetic patches are concatenated at the end of each pipe in a weighed fashion to force consistency.

Recently,Z. Zhang et al. (2018) adopted the CycleGAN approach to model between cardiovascular MRI and CT images at 3D scale, with an added shape-consistency loss to enforce robust anatomical structures during synthesis. The architecture includes an additional 3D convolutional neural network for volume based segmentation, which is trained together with the adversarial generators. Unlike the previous two methods, this work did not perform generation on small 3D patches that were extracted from source images. Instead, it had introduced a number of substantial modifications to the standard CycleGAN architecture that allowed it to simplify the computational complexity.

While it can be argued that the data in this project is lower resolution than most medical imaging tasks, its computational complexity in 3D would still be substantial. For example, a 3D fetal brain with 40 slices and 67x67 in plane resolution yields 180,000 components. In contrast, a high resolution 2D MR image is typically reconstructed with a 256x256 matrix, thereby yielding 65,000 components. The context of 3D modelling, a fetal brain is therefore a challenge that is similar to generation of high resolution images. For example, a 512x512 image has a vector size of 260,000.

6.6.2 High resolution synthesis

Until recently, generation of high resolution images with GANs has been a prohibitive endeavour due to the extreme dimensionality that is associated with such images. The key progress in tackling this has been made in the work of Karras et al. (2017), termed Progressively Grown GANs (PG GAN). It introduced an incremental multi-resolution training procedure, which forces the network to optimize its parameters for generation of low resolution images with a few layers first, and progressively adds more layers and more layers to the generator and discriminator over time, while also increasing the resolution of the supplied image set. Such approach managed to successfully train an image generator to a record resolution of 1024x1024 images and has been subsequently applied to six studies in the medical domain (as of the current date). These studies used PG GAN primarily for data augmentation in high resolution mammography (Korkinof et al. 2018), Alzheimer's condition(Bowles et al. 2018a), brain tumors

(Han et al. 2019), vascular retinopathy (Beers et al. 2018), T1,T2 images (Bowles et al. 2018b) and skin lesions baur2018melanogans. The highest resolution was achieved by Korkinof et al. (2018) with 1280x1024 2D images.

Unfortunately, neither the PG GAN architecture nor its follow up study, termed StyleGAN (Karras et al. 2019), were built for conditional image generation to tackle image translation tasks, nor 3D image generation. Therefore, substantial efforts would be required to train and test such a network. Nevertheless, this may be a promising and valuable undertaking, and could well serve the purpose for the extension of the current study. I intend to trial this path in the follow up works.

6.6.3 Spatially constrained GANs

To reduce the computational complexity of modelling consistency in the through-slice dimension for the entire height of the slice stack, it may be possible to just consider the neighbouring slices as regularizers for a given target slice. Such approach was used by Dar et al. (2019) to perform image to image translation between T1 and T2 data, which showed improved structural consistency of the translated target slice.

Alternatively, generation can be constrained in-plane by forcing the generator to adhere to a certain structure. This can be served in a form of a mask or a more detailed segmentation of brain structures. Such approach has been tested by Z. Zhang et al. (2018), who added an additional cost to the objective function based on shape consistency. An alternative approach may be to concatenate the segmentation mask of the target image in the paired data to the source image on input. This may be a more appropriate method given that the T2* weighted data is generally constrained due to the signal loss on the tissue air interfaces. By feeding this information to the network through the segmentation mask, it may force it to indirectly learn this contrast based relationship.

In summary, there are a breadth of exciting new approaches that can be utilized to extend the results of this study to generate high resolution high fidelity consistent 3D volumes of the fetal

brain anatomy with fewer data. There is an important clinical need to reduce the number of scans required for such a 3D reconstruction. The author will explore these new approaches in his future work.

Chapter 7

Conclusion

7.1 Summary of Thesis Achievements

The aim of the research presented in this thesis was to develop motion robust methods for quantitative T2* relaxometry for the developing brain. Our approach to this field of study considered the challenges that lie on both sides of the MR imaging spectrum - acquisition and reconstruction. This view had enabled us to develop a holistic understanding of the processes that are involved in fetal and neonatal brain imaging, not only from a research perspective, but also from a perspective of clinical limitations that may be imposed on the imaging process in the day to day scanning environment.

7.1.1 Acquisition

Due to maternal respiration and unpredictable motion of the fetus, in-utero MR imaging is highly constrained by the acquisition methods that can attain artifact free fetal anatomy. Conventional methods for quantitative T2* relaxometry that provide high resolution densely sampled T2* decay estimates are not feasible in such environment. Therefore, at the beginning of our research, we focused on developing a method for ultra fast scanning that will be robust to motion artifacts. We proposed and tested the use of multi-echo single shot gradient echo echo planar imaging

(ME ssGRE EPI) sequence with the added acceleration from SENSE reconstruction and partial Fourier acquisition, as an effective method for achieving this goal.

The sequence was tested for accuracy and precision by a way of comparison with a golden reference scan for T2* relaxometry, a multi echo gradient recalled acquisition. The tests were performed on adult subjects and T2* estimates were calculated within the limits that were achievable for the short T2* values of the matured adult brain. The ME ssGRE EPI pipeline was optimized to achieve a balance between satisfactory resolution, T2* signal decay, acoustic noise and clinically acceptable scan times.

For our next study, we transported this approach to a 3T Phillips Achieva scanner. Due to the higher magnetic field strength, further considerations had to be given for the acoustic noise that was generated due to the stronger gradients. Clinical guidelines for acoustic noise had prompted us to reduce the slew rates on the acquisition sequence.

As the T2* relaxation times in the newborn brain are similar to those of fetal brain, and are therefore considerably higher than in adult subjects, ME ssGRE EPI sequence could also be used to acquire sufficiently dense samples of the decaying T2* signal in the newborn subject group. While motion is a significantly lesser challenge in the neonatal domain, the benefit of this approach is the speed of acquisition. Similarly to the fetal brain imaging, neonatal scanning time is constrained. Therefore, fast MR imaging sequences are not only convenient but also essential in many clinical day to day scenarios. The approach that we proposed can therefore be used practically within standard clinical setting in most cases.

The proposed acquisition yielded a sequence of slices along the T2* decaying signal, sampled within a period of $<$500ms. Each individual slice, at an individual echo time, was acquired within a single k-space traversal, and thus represented 43-50ms in time. While most of the fetal motion would be excluded within this timeframe, small misalignments between slices at different echo times were occasionally visible in single digit percentages of the total slices acquired. A substantial majority of these misalignments occured as 2D rotations in plane, which could be registered with via 2D rigid registration scheme.

Conversely, significant motion was observed within the spatially neighbouring slices. This manifested in stacks of slices that did not represent a unique 3D brain anatomy. In the first chapter of our thesis, the T2* measurements were therefore performed manually on a slice by slice basis, under the guidance of an expert radiologist. This procedure is time consuming. Consequently, it had been considered that the next valuable step would be to develop an automated method for the study of T2* values in the developing brain.

7.1.2 Reconstruction

To achieve automated segmentation of the key regions of interest, we fused a collection of individual methods for fetal and neonatal brain analysis. To accurately assign the labels to each voxel, an anatomical 3D representation of the brain is required for each subject. Thus, a slice to volume reconstruction was necessary. However, the use of ME ssGRE EPI sequence for this purpose was limited in the following ways. Firstly, the acquisition yielded a relatively low resolution. While the resolution could be increased, this would lengthen the inter-echo timing and therefore result in fewer and sparser samples along the decaying T2* signal, which would in turn reduce the accuracy of the T2* estimates. Secondly, the proposed approach included only 3 acquisitions of the fetal brain anatomy in the transverse plane, thereby making it insufficient for a robust slice to volume reconstruction. A denser sampling of the fetal brain anatomy with multiple orthogonal views would require a prolonged sequence, which would place further limitations on its practical usage.

To overcome the aforementioned challenges, an alternative reconstruction pipeline was considered, which would make use of auxiliary data. Structural T2 weighted scans are a primary imaging contrast for assessment of fetal and neonatal brain abnormalities and therefore form part of the standard clinical protocol. In a day to day scenario, it is a requirement for a radiologist to acquire sufficient amount of T2 weighted data to aid in slice to volume reconstruction of a high resolution 3D representation of the T2 image of the brain, which is held as a priority above all other scans. As a result, it is always possible to obtain a slice to volume reconstructed 3D brain T2W image for a subject that has been scanned. Our approach was therefore to perform 3D

reconstruction on the T2W image, segment it with help of an atlas, and propagate the labels into the T2* weighed space. This framework enabled estimation of individual voxels labels for the entirety of the reconstructed T2* maps, and thus completed our goal of performing automated T2* mapping.

The motivation for our next study was to improve the limitation posed by slice to volume reconstruction algorithm, which required dense oversampling of the image space in multiple orthogonal planes in order to deliver high quality 3D reconstructions of the fetal brain anatomy. This limitation had been expressly relevant to our previous studies, and was therefore a logical extension to the overall goal of this thesis. To solve this challenge, we chose to focus on deep learning as a tool that had achieved transformational advances in the medical image analysis in the last few years. To tackle sparsity in the image domain, a conditional generative adversarial model was trained between T2 and T2* weighted data to enable image to image translation between these contrasts. By using the knowledge that is embedded in the data from a very large training set from the developing Human Connectome Project, and the prior information that comes from the auxiliary contrast of the same subject, our algorithm was able to generate high fidelity T2* slices of the whole brain for each subject. This synthetic volume could then be added to the slice to volume reconstruction pipeline together with the limited number of motion corrupted slices to enable high resolution reconstruction of 3D brain images.

7.1.3 Clinical Significance

The findings presented in this thesis showed a significant difference for T2* relaxation times between fetal and preterm neonatal subjects. While the fetal brain had shown consistent trend in the decrease of T2* values with gestational age in frontal and occipital white matter, no such trend was detected in preterm neonatal studies in the works of this thesis and neither in the works of other literature. Further analysis revealed high variability in T2* values for both fetal and preterm neonatal subjects, which indicates that the development rate and microstructure in the developing brain is highly varied and thus normalizes at a much later stage during infancy. However, by comparing variability and trends with studies that reported T2 values in the

preterm brain, we also noted that the variability in preterm brain could be attributed to iron deficiency. At this point, further research is necessary to make a conclusive remark on iron deficiency and would be of valuable contribution to scientific community. Finally, no significant difference had been found in T2* values of clinical and control fetal brain subjects.

7.2 Applications

The research presented in this thesis forms part of the first extensive study of quantitative T2* values in the fetal and neonatal brain. The measurements reported in our study can be used as a reference for clinicians and future research endeavours.

T2* relaxometry have previously been shown to be a marker for myelin microstructure and a guide for orientation and size of myelin fibers. The methods developed in this thesis can be used in conjunction with other sequences, such as diffusion tensor imaging, to provide adjacent measurements for greater reliability. A range of prior works have used these two mechanisms in tandem in the adult brain studies.

Furthermore, quantitative T2* mapping is a primary method for measurement of in-vivo iron concentration via MR imaging. It is routinely used in study of cardiac and iron thalassemia, and in the studies of neurodegenerative diseases in the old age, which are linked to changes in iron concentration. Notably, iron deficiency is also a prevalent issue in preterm birth, affecting anywhere between 25-85% of the infants depending on the socio-economic and geographical factors. More importantly, reduced iron metabolism has been shown to have an effect on cognitive development in the later infancy. The work presented in this thesis can be used as a primer towards the development of reliable iron biomarkers in fetal and neonatal subjects. However, it is noted here that T2* relaxometry is sensitive to both myelin changes and iron concentration, thereby requiring disentanglement of both measures as these will occur simultaneously in the developing brain. This could be achieved by joint imaging with methods that are sensitive to myelin changes but are not sensitive to iron concentration.

Finally, the T2* values reported in this thesis can be a valuable indicator for optimization of

BOLD fMRI acquisition. Conventionally, BOLD fMRI protocols strive to equate the selected echo time (TE) with the T2* value of the key area of interest for which the study is being designed for. Such arrangement achieves optimal contrast between BOLD activations, which are highly sensitive to noise. In instances where multiple regions with highly different values of T2* relaxation time are being considered, the knowledge of the reported values in this thesis could aid in selecting the optimal parameters for a dual-echo, or even a multi-echo acquisition.

Notably, our study indicated a high variability in T2* values between individual subjects, both in fetal and neonatal cohorts. This variation may yield a significant influence on the accuracy of fMRI studies from subject to subject. An alternate approach may be to acquire a reference T2* map of the brain prior to the fMRI scan, and vary the echo time parameter in the fMRI protocol accordingly for each patient dynamically.

7.3 Future Work

There exist several exciting directions in which the work in this thesis could be extended to.

7.3.1 SMS

From the perspective of acquisition, one of the most promising recent techniques for fast MR imaging has been the introduction of simultaneous multi-slice imaging (Barth et al. 2016). It could be argued that its impact on the community is following a similar trajectory to that of parallel imaging with SENSE when it was first published in 1999. A large number of centers had been integrating this technique into the clinical pipeline and it forms part of the largest MR studies in the recent years such as the HCP and dHCP projects (Edwards et al. 2019; Ugurbil et al. 2013). Moreover, this technique is particularly useful for fetal imaging where speed of acquisition is directly linked to quality of the collected images. Very recently, SMS techniques have also been used to constrain the search space for iterative registration of slice to volume reconstruction (Marami et al. 2019), which has shown to improve the quality of the

reconstruction. Integrating the proposed acquisition approach with SMS would therefore play a vital role in making the acquisition faster and thereby eminently more valuable as a tool that could be used for T2* mapping within regular clinical practice.

7.3.2 T2 relaxometry

The second valuable direction for the future application of the methods introduced in this thesis is to apply them to the study of T2 relaxometry of the fetal brain. Unlike T2* constants, T2 constants are not confounded by local field inhomogeneities and are governed solely by spin-spin relaxation processes. This yields T2 relaxometry as a valuable quantitative tool for the study of water compartmentalization and, as a consequence, a direct measure of myelination that is not confounded by heme or non-heme iron. Furthermore, the use of T2 relaxometry may be particularly relevant to the study of diffuse excessive high signal intensity (DEHSI), which is a phenomenon that is frequently observed in T2 weighted imaging of preterm born neonates and has been linked to be a predictor of pathology at later stages of infancy (Parikh et al. 2013).

The methods developed in this thesis are capable to be transported to the domain of T2 mapping. In fact, during the experimental work conducted on this thesis, an early attempt was made by the author to perform T2 relaxometry on the fetal brain. This involved acquisition of single slice multi-echo single shot fast spin echo echo planar images. However, initial tests had exhibited a number of artifacts, the cause of which was hard to disentangle and would require substantial effort in a separate individual project. Once these would be resolved, additional efforts would be necessary to correct the signal from stimulated echoes, which arise due to imperfect refocusing pulses in multi echo spin echo sequences (Hennig 1991). A set of distinct approaches exist that aim to correct for this effect. The approach that would be most appropriate is the work of Lebel et al. (2010), which completes the fitting model by considering the extended phase graph (EPG) algorithm and the knowledge of the slice profile. Alternatively, the work of Ben-Eliezer et al. (2015) introduced and idea of creating a dictionary of all possible echo-modulated trajectories that a multi-echo spin echo signal can take. Such trajectories can be simulated via Bloch equations, and the recorded signal would then be matched with a corresponding dictionary

entry to uncover the true T2 constants.

Furthermore, we note that the use of highly accelerated multi-echo single shot EPI fast spin echo sequences is prone to very high specific absorption rate (SAR), which requires practical changes to the sequence (Allison et al. 2015). Specifically, a common strategy to reduce SAR is to implement a much lower flip angle, e.g. 130°, which in turn yields much more pronounced stimulated echoes (Loening et al. 2015). Variable refocusing pulses is an effective remedy for this problem, particularly for very long echo trains that would be required for multi-echo imaging in T2 relaxometry (Busse et al. 2006).

Alternatively, a different sequence type may be trialled altogether to obtain the T2 values, such as the balanced steady-state free precession (bSSFP). bSSFP is specifically known for very low SAR (Schieda et al. 2017) and is widely used in fetal imaging (Malamateniou et al. 2013). Deoni et al. (2005) introduced a bSSFP based method, termed DESPOT2, which calculates the transverse T2 relaxation times from the knowledge of TR, TE, T1 and flip angle. One downside of this approach is that it requires an additional acquisition, such as the DESPOT1 method, to obtain the T1 values. Furthermore, acquisition methods used in DESPOT1 and DESPOT2, are motion sensitive and time intensive, and would therefore require further exploration if they were to be considered for fetal MR imaging.

7.3.3 Multi-contrast SVR for sparse data

The work described in this section will be undertaken by the author at his next research role at Boston Children's Hospital. The first direction of the research will be the inclusion of 3D spatial relationships into the conditional generative adversarial network presented in Chapter 6 to lesser the discontinuities across the slices as they are concatenated and used in slice to volume reconstruction. A number of methods that could achieve this have been described in section 6.6, which include 3D GANs, through-slice regularization via neighbouring slices and in-plane regularization via segmentation masks. Once sufficient consistency would have been achieved and tested on the fetal data, an attempt to apply this approach to other image contrasts, such as DTI, field maps and T1 acquisitions, will be a highly valuable next step. In general, it takes

a great amount of effort to acquire images of satisfactory quality during a fetal clinical MRI examination, which often results in multiple re-runs. With strict time constraints, as well as the multitude of risks of artifacts that may arise due to motion, it is often difficult to obtain sufficient amount of data for 3D reconstruction. Pushing the limit of SVR methods in the direction of lesser requirements on the density of the acquired data would enable faster examinations and more thorough analysis by availability of 3D structures from multiple contrasts.

Bibliography

Alansary, A. et al. (2017). "PVR: Patch-to-volume reconstruction for large area motion correction of fetal MRI". In: *IEEE transactions on medical imaging* 36.10, p. 2031.

Allison, J. and Yanasak, N. (2015). "What MRI sequences produce the highest specific absorption rate (SAR), and is there something we should be doing to reduce the SAR during standard examinations?" In: *American Journal of Roentgenology* 205.2, W140–W140.

Arjovsky, M., Chintala, S., and Bottou, L. (2017). "Wasserstein gan". In: *arXiv preprint arXiv:1701.07875*.

Armanious, K., Jiang, C., Fischer, M., Küstner, T., Nikolaou, K., Gatidis, S., and Yang, B. (2018). "MedGAN: Medical image translation using GANs". In: *arXiv preprint arXiv:1806.06397*.

Ashburner, J. and Friston, K. J. (2007). "Rigid body registration". In: *Statistical parametric mapping: The analysis of functional brain images*, pp. 49–62.

Barth, M., Breuer, F., Koopmans, P., Norris, D., and Poser, B. (2016). "Simultaneous multislice (SMS) imaging techniques". fr. In: *Magnetic Resonance in Medicine* 75.1, pp. 63–81.

Beers, A., Brown, J., Chang, K., Campbell, J. P., Ostmo, S., Chiang, M. F., and Kalpathy-Cramer, J. (2018). "High-resolution medical image synthesis using progressively grown generative adversarial networks". In: *arXiv preprint arXiv:1805.03144*.

Ben-Eliezer, N., Sodickson, D. K., and Block, K. T. (2015). "Rapid and accurate T2 mapping from multi–spin-echo data using Bloch-simulation-based reconstruction". In: *Magnetic resonance in medicine* 73.2, pp. 809–817.

Bender, B. and Klose, U. (2010). "The in vivo influence of white matter fiber orientation towards B0 on T2* in the human brain". In: *NMR in biomedicine* 23.9, pp. 1071–1076.

Berthelot, D., Schumm, T., and Metz, L. (2017). "Began: Boundary equilibrium generative adversarial networks". In: *arXiv preprint arXiv:1703.10717*.

Blazejewska, A. I., Seshamani, S., McKown, S., Caucutt, J., Dighe, M., Gatenby, C., and Studholme, C. (2015). "Full 3D mapping of T2* relaxation times from mid to late gestation of the normal fetal brain". In: *Proceedings of ISMRM*.

Blazejewska, A. I., Seshamani, S., McKown, S. K., Caucutt, J. S., Dighe, M., Gatenby, C., and Studholme, C. (2017). "3D in utero quantification of T2* relaxation times in human fetal brain tissues for age optimized structural and functional MRI". In: *Magnetic resonance in medicine* 78.3, pp. 909–916.

Bloch, F. (1946). "Nuclear induction". In: *Physical review* 70.7-8, p. 460.

Bloembergen, N., Purcell, E. M., and Pound, R. V. (1948). "Relaxation effects in nuclear magnetic resonance absorption". In: *Physical review* 73.7, p. 679.

Bonel, H., Frei, K., Raio, L., Meyer-Wittkopf, M., Remonda, L., and Wiest, R. (2008). "Prospective navigator-echo-based real-time triggering of fetal head movement for the reduction of artifacts". In: *European radiology* 18.4, p. 822.

Bowles, C., Gunn, R., Hammers, A., and Rueckert, D. (2018a). "Modelling the progression of Alzheimer's disease in MRI using generative adversarial networks". In: *Medical Imaging 2018: Image Processing*. Vol. 10574. International Society for Optics and Photonics, 105741K.

Bowles, C. et al. (2018b). "GAN augmentation: augmenting training data using generative adversarial networks". In: *arXiv preprint arXiv:1810.10863*.

Brown, R. W., Cheng, Y.-C. N., Haacke, E. M., Thompson, M. R., and Venkatesan, R. (2014). *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons.

Busse, R. F., Hariharan, H., Vu, A., and Brittain, J. H. (2006). "Fast spin echo sequences with very long echo trains: design of variable refocusing flip angle schedules and generation of clinical T2 contrast". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 55.5, pp. 1030–1037.

Bydder, G. and Young, I. (1985). "MR imaging: clinical use of the inversion recovery sequence". In:

Cameron, I. (2013). "Techniques of Fat Suppression". In: *MR Phys Ottawa Hosp* 501.

Carneiro, A., Vilela, G., De Araujo, D., and Baffa, O. (2006). "MRI relaxometry, methods and applications". en. In: *Brazilian journal of physics* 36.1, pp. 9–15.

Chavhan, G. B., Babyn, P. S., Thomas, B., Shroff, M. M., and Haacke, E. M. (2009). "Principles, techniques, and applications of T2*-based MR imaging and its special applications". In: *Radiographics* 29.5, pp. 1433–1449.

Cherubini, A., Péran, P., Hagberg, G. E., Varsi, A. E., Luccichenti, G., Caltagirone, C., Sabatini, U., and Spalletta, G. (2009). "Characterization of white matter fiber bundles with T relaxometry and diffusion tensor imaging". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 61.5, pp. 1066–1072.

Cho, Z. and Ro, Y. (1992). "Reduction of susceptibility artifact in gradient-echo imaging". In: *Magnetic resonance in medicine* 23.1, pp. 193–200.

Choi, H. and Lee, D. S. (2018). "Generation of structural MR images from amyloid PET: Application to MR-less quantification". In: *Journal of Nuclear Medicine* 59.7, pp. 1111–1117.

Chung, H. (2019). *Magnetic Resonance Imaging: Principles And Its Applications, Lecture Notes: MRI Image Artifacts and their Remedies*. URL: `https://mrilab.org/courses.html`. (accessed: 01.08.2019).

Cohen-Adad, J. (2014). "What can we learn from T2* maps of the cortex?" en. In: *Neuroimage* 93, pp. 189–200.

Conolly, S., Nishimura, D., Macovski, A., and Glover, G. (1988). "Variable-rate selective excitation". In: *Journal of Magnetic Resonance (1969)* 78.3, pp. 440–458.

Cordero-Grande, L., Hughes, E. J., Hutter, J., Price, A. N., and Hajnal, J. V. (2018). "Three-dimensional motion corrected sensitivity encoding reconstruction for multi-shot multi-slice MRI: application to neonatal brain imaging". In: *Magnetic resonance in medicine* 79.3, pp. 1365–1376.

Counsell, S., Kennea, N., Herlihy, A., Allsop, J., Harrison, M., Cowan, F., and Hajnal, J. (2003). "T2 relaxation values in the developing preterm brain". en. In: *American Journal of Neuroradiology* 24.8, pp. 1654–60.

Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., and Bharath, A. A. (2018). "Generative adversarial networks: An overview". In: *IEEE Signal Processing Magazine* 35.1, pp. 53–65.

Czervionke, L. F., Czervionke, J. M., Daniels, D. L., and Haughton, V. M. (1988). "Characteristic features of MR truncation artifacts". In: *American journal of roentgenology* 151.6, pp. 1219–1228.

Czervionke, L., Daniels, D., Wehrli, F., Mark, L., Hendrix, L., Strandt, J., Williams, A., and Haughton, V. (1988). "Magnetic susceptibility artifacts in gradient-recalled echo MR imaging." In: *American journal of neuroradiology* 9.6, pp. 1149–1155.

Dahnke, H. and Schaeffter, T. (2005). "Limits of detection of SPIO at 3.0 T using T2* relaxometry". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 53.5, pp. 1202–1206.

Dar, S. U., Yurt, M., Karacan, L., Erdem, A., Erdem, E., and Çukur, T. (2019). "Image synthesis in multi-contrast MRI with conditional generative adversarial networks". In: *IEEE transactions on medical imaging*.

Deistung, A., Schäfer, A., Schweser, F., Biedermann, U., Turner, R., and Reichenbach, J. R. (2013). "Toward in vivo histology: a comparison of quantitative susceptibility mapping (QSM) with magnitude-, phase-, and R2*-imaging at ultra-high magnetic field strength". In: *Neuroimage* 65, pp. 299–314.

Delfaut, E. M., Beltran, J., Johnson, G., Rousseau, J., Marchandise, X., and Cotten, A. (1999). "Fat suppression in MR imaging: techniques and pitfalls". In: *Radiographics* 19.2, pp. 373–382.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). "ImageNet: A Large-Scale Hierarchical Image Database". In: *CVPR09*.

Deng, W., Yang, C., Alagappan, V., Wald, L. L., Boada, F. E., and Stenger, V. A. (2009). "Simultaneous z-shim method for reducing susceptibility artifacts with multiple transmitters". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 61.2, pp. 255–259.

Deoni, S. (2010). "Quantitative Relaxometry of the Brain". en. In: *Topics in Magnetic Resonance Imaging*. Vol. 21. 2, pp. 101–113.

Deoni, S., Dean, D., O'Muircheartaigh, J., Dirks, H., and Jerskey, B. (2012). "Investigating white matter development in infancy and early childhood using myelin water faction and relaxation time mapping". en. In: *NeuroImage* 63.3, pp. 1038–1053.

Deoni, S., Peters, T., and Rutt, B. (2005). "High-resolution T1 and T2 mapping of the brain in a clinically acceptable time with DESPOT1 and DESPOT2". en. In: *Magnetic Resonance in Medicine* 53.1, pp. 237–241.

Dobbing, J. and Sands, J. (1973). "Quantitative growth and development of human brain". en. In: *Archives in disease in childhood*. Vol. 48, pp. 757–767.

Dubois, J., Dehaene-Lambertz, G., Kulikova, S., Poupon, C., Hüppi, P. S., and Hertz-Pannier, L. (2014). "The early development of brain white matter: a review of imaging studies in fetuses, newborns and infants". In: *neuroscience* 276, pp. 48–71.

Edwards, D., Hajnal, J., Rueckert, D., and Smith, S. (2019). *Developing Human Connectome Project*. URL: `http://www.developingconnectome.org`. (accessed: 01.08.2019).

Emami, H., Dong, M., Nejad-Davarani, S. P., and Glide-Hurst, C. K. (2018). "Generating synthetic CTs from magnetic resonance images using generative adversarial networks". In: *Medical physics* 45.8, pp. 3627–3636.

Ertl-Wagner, B., Lienemann, A., Strauss, A., and Reiser, M. F. (2002). "Fetal magnetic resonance imaging: indications, technique, anatomical considerations and a review of fetal abnormalities". In: *European radiology* 12.8, pp. 1931–1940.

Feinberg, D., Moeller, S., Smith, E., Auerbach, S., Ramanna, S., Gunther, M., and Glasser, M. (2010a). "Multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging". en. In: *PloS* one, pp. 5–12.

Feinberg, D., Moeller, S., Smith, S., Auerbach, E., Ramanna, S., Glasser, M., Miller, K., Ugurbil, K., and Yacoub, E. (2010b). "Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging". en. In: *PloS* 5.12, p. 15710.

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

Felmlee, J. P. and Ehman, R. L. (1987). "Spatial presaturation: a method for suppressing flow artifacts and improving depiction of vascular anatomy in MR imaging." In: *Radiology* 164.2, pp. 559–564.

Feng, Y., He, T., Gatehouse, P. D., Li, X., Harith Alam, M., Pennell, D. J., Chen, W., and Firmin, D. N. (2013). "Improved MRI R2* relaxometry of iron-loaded liver with noise correction". In: *Magnetic resonance in medicine* 70.6, pp. 1765–1774.

Feng, Y. et al. (2014). "Improved pixel-by-pixel MRI R2* relaxometry by nonlocal means". In: *Magnetic resonance in medicine* 72.1, pp. 260–268.

Fernandez-Seara, M. and Wehrli, F. (2000). "Postprocessing technique to correct for background gradients in image-based R* 2 measurements". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 44.3, pp. 358–366.

Ferrante, E. and Paragios, N. (2017). "Slice-to-volume medical image registration: A survey". In: *Medical image analysis* 39, pp. 101–123.

Ferrazzi, G. et al. (2014). "Resting state fMRI in the moving fetus: a robust framework for motion, bias field and spin history correction". In: *Neuroimage* 101, pp. 555–568.

Frahm, J., Merboldt, K.-D., and Hänicke, W. (1988). "Direct FLASH MR imaging of magnetic field inhomogeneities by gradient compensation". In: *Magnetic resonance in medicine* 6.4, pp. 474–480.

Fukunaga, M. et al. (2010). "Layer-specific variation of iron content in cerebral cortex as a source of MRI contrast". In: *Proceedings of the National Academy of Sciences* 107.8, pp. 3834–3839.

Gati, J. S., Menon, R. S., Uğurbil, K., and Rutt, B. K. (1997). "Experimental determination of the BOLD field strength dependence in vessels and tissue". In: *Magnetic resonance in medicine* 38.2, pp. 296–302.

Georgieff, M. K. and Innis, S. M. (2005). "Controversial nutrients that potentially affect preterm neurodevelopment: essential fatty acids and iron". In: *Pediatric Research* 57.5 Part 2, 99R.

Gholipour, A. et al. (2014). "Fetal MRI: a technical update with educational aspirations". In: *Concepts in Magnetic Resonance Part A* 43.6, pp. 237–266.

Gholipour, A. et al. (2017). "A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth". In: *Scientific reports* 7.1, p. 476.

Gholipour, A., Estroff, J. A., and Warfield, S. K. (2010). "Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI". In: *IEEE transactions on medical imaging* 29.10, pp. 1739–1758.

Goitein, O. et al. (2013). "Fetal liver T2* values, defining a standardized scale". it. In: *Journal of Magnetic Resonance Imaging* 38.6, pp. 1342–1345.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). "Generative adversarial nets". In: *Advances in neural information processing systems*, pp. 2672–2680.

Gousias, I., Edwards, D., Rutherford, M., Counsell, S., Hajnal, J., Rueckert, D., and Hammers, A. (2012). "Magnetic Resonance Imaging of the newborn brain, Manual segmentation of labelled atlases in term-born and preterm infants". fr. In: *NeuroImage* 62.3, pp. 1499–1509.

Gousias, I. S., Hammers, A., Counsell, S. J., Srinivasan, L., Rutherford, M. A., Heckemann, R. A., Hajnal, J. V., Rueckert, D., and Edwards, A. D. (2013). "Magnetic resonance imaging of the newborn brain: automatic segmentation of brain images into 50 anatomical regions". In: *PloS one* 8.4, e59990.

Graham, S. J., Stanchev, P. L., and Bronskill, M. J. (1996). "Criteria for analysis of multicomponent tissue T2 relaxation data". In: *Magnetic Resonance in Medicine* 35.3, pp. 370–378.

Griffiths, P. and Crossman, A. (1993). "Distribution of iron in the basal ganglia and neocortex in postmortem tissue in Parkinson's disease and Alzheimer's disease". In: *Dementia and geriatric cognitive disorders* 4.2, pp. 61–65.

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. (2017). "Improved training of wasserstein gans". In: *Advances in neural information processing systems*, pp. 5767–5777.

Haase, A., Frahm, J., Matthaei, D., Hanicke, W., and Merboldt, K. (1986). "FLASH imaging, rapid NMR imaging using low flip angle pulses". es. In: *Journal of Magnetic Resonance* 62.2, pp. 258–266.

Hagemeier, J., Yeh, E. A., Brown, M. H., Bergsland, N., Dwyer, M. G., Carl, E., Weinstock-Guttman, B., and Zivadinov, R. (2013). "Iron content of the pulvinar nucleus of the thalamus is increased in adolescent multiple sclerosis". In: *Multiple Sclerosis Journal* 19.5, pp. 567–576.

Halliday, H. L., Lappin, T. R., and McClure, G. (1984). "Iron status of the preterm infant during the first year of life". In: *Neonatology* 45.5, pp. 228–235.

Han, C., Rundo, L., Araki, R., Furukawa, Y., Mauri, G., Nakayama, H., and Hayashi, H. (2019). "Infinite Brain MR Images: PGGAN-based Data Augmentation for Tumor Detection". In: *arXiv preprint arXiv:1903.12564*.

Hayat, T., Nihat, A., Martinez-Biarge, M., McGuinness, A., Allsop, J., Hajnal, J., and Rutherford, M. (2011). "Optimization and initial experience of a multisection balanced steady-state free precession cine sequence for the assessment of fetal behavior in utero". en. In: *American Journal of Neuroradiology* 32.2, pp. 331–338.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.

He, T., Gatehouse, P. D., Smith, G. C., Mohiaddin, R. H., Pennell, D. J., and Firmin, D. N. (2008). "Myocardial T measurements in iron-overloaded thalassemia: An in vivo study to investigate optimal methods of quantification". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 60.5, pp. 1082–1089.

He, T., Zhang, J., Carpenter, J.-P., Feng, Y., Smith, G. C., Pennell, D. J., and Firmin, D. N. (2013). "Automated truncation method for myocardial T2* measurement in thalassemia". In: *Journal of magnetic resonance imaging* 37.2, pp. 479–483.

Hennig, J. (1991). "Echoes—how to generate, recognize, use or avoid them in MR-imaging sequences. Part I: Fundamental and not so fundamental properties of spin echoes". In: *Concepts in Magnetic Resonance* 3.3, pp. 125–143.

Hernando, D., Vigen, K. K., Shimakawa, A., and Reeder, S. B. (2012). "R mapping in the presence of macroscopic B0 field variations". In: *Magnetic resonance in medicine* 68.3, pp. 830–840.

Holland, P. and Welsch, R. (1977). "Robust regression using iteratively reweighted least-squares". en. In: *Communications in Statistics, Theory and Methods, A6*, pp. 813–827.

Hood, M. N., Ho, V. B., Smirniotopoulos, J. G., and Szumowski, J. (1999). "Chemical shift: the artifact and clinical tool revisited". In: *Radiographics* 19.2, pp. 357–371.

Hou, B., Alansary, A., McDonagh, S., Davidson, A., Rutherford, M., Hajnal, J. V., Rueckert, D., Glocker, B., and Kainz, B. (2017). "Predicting slice-to-volume transformation in presence of arbitrary subject motion". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 296–304.

Huynh, T., Gao, Y., Kang, J., Wang, L., Zhang, P., Lian, J., and Shen, D. (2015). "Estimating CT image from MRI data using structured random forest and auto-context model". In: *IEEE transactions on medical imaging* 35.1, pp. 174–183.

Hwang, D., Kim, D., and Du, Y. (2010a). "In vivo multi-slice mapping of myelin water content using T2* decay". en. In: *NeuroImage* 52.1, pp. 198–204.

Hwang, D., Kim, D.-H., and Du, Y. P. (2010b). "In vivo multi-slice mapping of myelin water content using T2* decay". In: *NeuroImage* 52.1, pp. 198–204.

Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2017). "Image-to-image translation with conditional adversarial networks". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134.

Jezzard, P. (2006). "Shim Coil Design, Limitations and Implications". en. In: *International Society of Magnetic Resonance in Medicine (ISMRM) Annual Meeting*.

Jezzard, P. and Balaban, R. (1995). "Correction for geometric distortion in echo planar images from B0 field variations". en. In: *Magnetic Resonance in Medicine* 34.1, pp. 65–73.

Jezzard, P., Barnett, A. S., and Pierpaoli, C. (1998). "Characterization of and correction for eddy current artifacts in echo planar diffusion imaging". In: *Magnetic resonance in medicine* 39.5, pp. 801–812.

Jezzard, P. and Clare, S. (1999). "Sources of distortion in functional MRI data". In: *Human brain mapping* 8.2-3, pp. 80–85.

Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., and Hajnal, J. (2007). "MRI of moving subjects using multislice snapshot images with volume reconstruction (SVR), application to fetal, neonatal, and adult brain studies". en. In: *IEEE Transactions on Medical Imaging* 26.7, pp. 967–80.

Juchem, C. (2013). “Shimming: Fields, Coils & Control”. In:

Kainz, B. et al. (2015). “Fast volume reconstruction from motion corrupted stacks of 2D slices”. In: *IEEE transactions on medical imaging* 34.9, pp. 1901–1913.

Karras, T., Aila, T., Laine, S., and Lehtinen, J. (2017). “Progressive growing of gans for improved quality, stability, and variation”. In: *arXiv preprint arXiv:1710.10196*.

Karras, T., Laine, S., and Aila, T. (2019). “A style-based generator architecture for generative adversarial networks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410.

Kazeminia, S., Baur, C., Kuijper, A., Ginneken, B. van, Navab, N., Albarqouni, S., and Mukhopadhyay, A. (2018). “GANs for medical image analysis”. In: *arXiv preprint arXiv:1809.06222*.

Kim, K., Habas, P. A., Rajagopalan, V., Scott, J. A., Corbett-Detig, J. M., Rousseau, F., Barkovich, A. J., Glenn, O. A., and Studholme, C. (2011). “Bias field inconsistency correction of motion-scattered multislice MRI for improved 3D image reconstruction”. In: *IEEE transactions on medical imaging* 30.9, pp. 1704–1712.

Kim, K., Habas, P. A., Rousseau, F., Glenn, O. A., Barkovich, A. J., and Studholme, C. (2009). “Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation”. In: *IEEE transactions on medical imaging* 29.1, pp. 146–158.

Kinney, H. C., Brody, B. A., Kloman, A. S., and Gilles, F. H. (1988). “Sequence of central nervous system myelination in human infancy: II. Patterns of myelination in autopsied infants”. In: *Journal of Neuropathology & Experimental Neurology* 47.3, pp. 217–234.

Koopmans, P. J., Boyacioğlu, R., Barth, M., and Norris, D. G. (2013). “Simultaneous multislice inversion contrast imaging using power independent of the number of slices (PINS) and delays alternating with nutation for tailored excitation (DANTE) radio frequency pulses”. In: *Magnetic resonance in medicine* 69.6, pp. 1670–1676.

Korkinof, D., Rijken, T., O’Neill, M., Yearsley, J., Harvey, H., and Glocker, B. (2018). “High-resolution mammogram synthesis using progressive generative adversarial networks”. In: *arXiv preprint arXiv:1807.03401*.

Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M. A., Hajnal, J. V., and Schnabel, J. A. (2012). "Reconstruction of fetal brain MRI with intensity matching and complete outlier removal". In: *Medical image analysis* 16.8, pp. 1550–1564.

Kurach, K., Lučić, M., Zhai, X., Michalski, M., and Gelly, S. (2019). "A Large-Scale Study on Regularization and Normalization in GANs". In: *International Conference on Machine Learning*, pp. 3581–3590.

Langkammer, C., Krebs, N., Goessler, W., Scheurer, E., Ebner, F., Yen, K., Fazekas, F., and Ropele, S. (2010). "Quantitative MR imaging of brain iron: a postmortem validation study". In: *Radiology* 257.2, pp. 455–462.

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

Larkman, D. J. and Nunes, R. G. (2007). "Parallel magnetic resonance imaging". In: *Physics in Medicine & Biology* 52.7, R15.

Laule, C. and Moore, G. W. (2018). "Myelin water imaging to detect demyelination and remyelination and its validation in pathology". In: *Brain Pathology* 28.5, pp. 750–764.

Lauterbur, P. C. et al. (1973). "Image formation by induced local interactions: examples employing nuclear magnetic resonance". In:

Lebel, R. M. and Wilman, A. H. (2010). "Transverse relaxometry with stimulated echo compensation". In: *Magnetic resonance in medicine* 64.4, pp. 1005–1014.

Lee, W., Donner, E., Nossin-Manor, R., Whyte, H., Sled, J., and Taylor, M. (2012). "Visual functional Magnetic Resonance Imaging of preterm infants". fr. In: *Developmental Medicine and Child Neurology* 54.8, pp. 724–9.

Levenberg, K. (1944). "A method for the solution of certain non-linear problems in least squares". en. In: *Quarterly of Applied Mathematics* 2, pp. 164–168.

Levine, D., Stroustrup-Smith, A., Barbaras, L., Gonsalves, A., and Frank, M. (n.d.). *Fetal MRI Appendix C: Artifacts*. URL: `http://radnet.bidmc.harvard.edu/fetalatlas/appendicies/Artifacts/Artifacts.html`. (accessed: 01.08.2019).

Li, T.-Q., Gelderen, P. van, Merkle, H., Talagala, L., Koretsky, A. P., and Duyn, J. (2006). "Extensive heterogeneity in white matter intensity in high-resolution T2*-weighted MRI of the human brain at 7.0 T". In: *Neuroimage* 32.3, pp. 1032–1040.

Li, T.-Q., Yao, B., Gelderen, P. van, Merkle, H., Dodd, S., Talagala, L., Koretsky, A. P., and Duyn, J. (2009). "Characterization of T2* heterogeneity in human brain white matter". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 62.6, pp. 1652–1657.

Liu, G. and Ogawa, S. (2006). "EPI image reconstruction with correction of distortion and signal losses". In: *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* 24.3, pp. 683–689.

Loening, A. M., Saranathan, M., Ruangwattanapaisarn, N., Litwiller, D. V., Shimakawa, A., and Vasanawala, S. S. (2015). "Increased speed and image quality in single-shot fast spin echo imaging via variable refocusing flip angles". In: *Journal of Magnetic Resonance Imaging* 42.6, pp. 1747–1758.

Lozoff, B. and Georgieff, M. K. (2006). "Iron deficiency and brain development". In: *Seminars in pediatric neurology*. Vol. 13. 3. Elsevier, pp. 158–165.

Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. (2018). "Are gans created equal? a large-scale study". In: *Advances in neural information processing systems*, pp. 700–709.

Lundström, U., Siimes, M. A., and Dallman, P. R. (1977). "At what age does iron supplementation become necessary in low-birth-weight infants?" In: *The Journal of pediatrics* 91.6, pp. 878–883.

Majumdar, A. and Ward, R. K. (2012). "Iterative estimation of MRI sensitivity maps and image based on sense reconstruction method (isense)". In: *Concepts in Magnetic Resonance Part A* 40.6, pp. 269–280.

Makropoulos, A., Gousias, I. S., Ledig, C., Aljabar, P., Serag, A., Hajnal, J. V., Edwards, A. D., Counsell, S. J., and Rueckert, D. (2014). "Automatic whole brain MRI segmentation of the developing neonatal brain". In: *IEEE transactions on medical imaging* 33.9, pp. 1818–1831.

Malamateniou, C., McGuinness, A., Allsop, J., O'Regan, D., Rutherford, M., and Hajnal, J. (2011). "Snapshot inversion recovery, an optimized single-shot T1-weighted inversion-recovery sequence for improved fetal brain anatomic delineation". en. In: *Radiology* 28.1, pp. 229–235.

Malamateniou, C. et al. (2013). "Motion-compensation techniques in neonatal and fetal MR imaging". In: *American Journal of Neuroradiology* 34.6, pp. 1124–1136.

Mansfield, P. and Maudsley, A. A. (1977). "Medical imaging by NMR". In: *The British journal of radiology* 50.591, pp. 188–194.

Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., and Paul Smolley, S. (2017). "Least squares generative adversarial networks". In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2794–2802.

Marami, B., Scherrer, B., Khan, S., Afacan, O., Prabhu, S. P., Sahin, M., Warfield, S. K., and Gholipour, A. (2019). "Motion-robust diffusion compartment imaging using simultaneous multi-slice acquisition". In: *Magnetic resonance in medicine* 81.5, pp. 3314–3329.

Marro, K., Otto, R., Kolokythas, O., Shimamura, A., Sanders, J. E., McDonald, G. B., and Friedman, S. D. (2011). "A simulation-based comparison of two methods for determining relaxation rates from relaxometry images". In: *Magnetic resonance imaging* 29.4, pp. 497–506.

Mescheder, L., Nowozin, S., and Geiger, A. (2017). "The numerics of gans". In: *Advances in Neural Information Processing Systems*, pp. 1825–1835.

Miller, D. J. et al. (2012). "Prolonged myelination in human neocortical evolution". In: *Proceedings of the National Academy of Sciences* 109.41, pp. 16480–16485.

Mirza, M. and Osindero, S. (2014). "Conditional generative adversarial nets". In: *arXiv preprint arXiv:1411.1784*.

Mitsumori, F., Watanabe, H., and Takaya, N. (2009). "Estimation of brain iron concentration in vivo using a linear relationship between regional iron and apparent transverse relaxation rate of the tissue water at 4.7 T". In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 62.5, pp. 1326–1330.

Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., and Uğurbil, K. (2010). "Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel

imaging with application to high spatial and temporal whole-brain fMRI". In: *Magnetic resonance in medicine* 63.5, pp. 1144–1153.

Nie, D., Trullo, R., Lian, J., Petitjean, C., Ruan, S., Wang, Q., and Shen, D. (2017). "Medical image synthesis with context-aware generative adversarial networks". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 417–425.

Nie, D., Trullo, R., Lian, J., Wang, L., Petitjean, C., Ruan, S., Wang, Q., and Shen, D. (2018). "Medical image synthesis with deep convolutional adversarial networks". In: *IEEE Transactions on Biomedical Engineering* 65.12, pp. 2720–2730.

OECD (2019). "Magnetic resonance imaging (MRI) exams (indicator)". In: (Accessed on 01 June 2020).

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

Ogawa, S., Lee, T.-M., Nayak, A. S., and Glynn, P. (1990). "Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields". In: *Magnetic resonance in medicine* 14.1, pp. 68–78.

Oh, C., Hilal, S., and Cho, Z. (1988). "Selective partial inversion recovery (SPIR) in steady state for selective saturation Magnetic Resonance Imaging". en. In: *(MRI). Proceedings of the 7th Annual Meeting, International Society for Magnetic Resonance in Medicine*.

Olut, S., Sahin, Y. H., Demir, U., and Unal, G. (2018). "Generative adversarial training for mra image synthesis using multi-contrast mri". In: *International Workshop on PRedictive Intelligence In MEdicine*. Springer, pp. 147–154.

Parikh, N. A., He, L., Bonfante-Mejia, E., Hochhauser, L., Wilder, P. E., Burson, K., and Kaur, S. (2013). "Automatically quantified diffuse excessive high signal intensity on MRI predicts cognitive development in preterm infants". In: *Pediatric neurology* 49.6, pp. 424–430.

Pennock, J. (1992). *Patient preparation, safety and hazards in imaging infants and children*. URL: `http://www.mrineonatalbrain.com/ch01-01.php`. (accessed: 01.08.2019).

Peran, P., Hagberg, G., Luccichenti, G., Cherubini, A., Brainovich, V., Celsis, P., Caltagirone, C., and Sabatini, U. (2007). "Voxel-based analysis of R2* maps in the healthy human brain". en. In: *Journal of Magnetic Resonance Imaging* 26.6, pp. 1413–20.

Peters, A. M., Brookes, M. J., Hoogenraad, F. G., Gowland, P. A., Francis, S. T., Morris, P. G., and Bowtell, R. (2007). "T2* measurements in human brain at 1.5, 3 and 7 T". In: *Magnetic resonance imaging* 25.6, pp. 748–753.

Pitt, D., Boster, A., Pei, W., Wohleb, E., Jasne, A., Zachariah, C. R., Rammohan, K., Knopp, M. V., and Schmalbrock, P. (2010). "IMaging cortical lesions in multiple sclerosis with ultra–high-field magnetic resonance imaging". In: *Archives of neurology* 67.7, pp. 812–818.

Posse, S. et al. (1999). "Enhancement of BOLD-contrast sensitivity by single-shot multi-echo functional MR imaging". en. In: *Magnetic Resonance in Medicine* 42.1, pp. 87–97.

Poustchi-Amin, M., Mirowitz, S., Brown, J., McKinstry, R., and Li, T. (2001a). "Principles and applications of echo-planar imaging, a review for the general radiologist". en. In: *Radiographics* 21.3, pp. 767–779.

Poustchi-Amin, M., Mirowitz, S. A., Brown, J. J., McKinstry, R. C., and Li, T. (2001b). "Principles and applications of echo-planar imaging: a review for the general radiologist". In: *Radiographics* 21.3, pp. 767–779.

Poynton, C., Jenkinson, M., Whalen, S., Golby, A. J., and Wells, W. (2008). "Fieldmap-free retrospective registration and distortion correction for epi-based functional imaging". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 271–279.

Price, A., Cordero-Grande, L., Malik, S., Abaei, M., Arichi, T., Hughes, E., Rueckert, D., Edwards, A., and Hajnal, J. (2015). "Accelerated Neonatal fMRI using Multiband EPI". In: *Proceedings of the 23rd annual meeting of ISMRM*.

Price, A., Hutter, J., Buhrer, L., Cordero Grande, L., Texeira, R., Hughes, E., Rutherford, M., and Hajnal, J. (2016a). "Multiband TSE Imaging of the Fetal Brain at 3T". en. In: *Proceedings of the 24th annual meeting of ISMRM*.

Price, A., Hutter, J., Cordero Grande, L., Hughes, E., Pegoretti, K., Gaspar, A., McCabe, L., Rutherford, M., and Hajnal, J. (2016b). "Multiplexed Multiband EPI for Increased Efficiency in Fetal FMRI and DMRI". en. In: *Proceedings of the 24th annual meeting of ISMRM*.

Pruessmann, K., Weiger, M., Scheidegger, M., and Boesiger, P. (1999). "SENSE, sensitivity encoding for fast MRI". en. In: *Magnetic Resonance in Medicine* 42.5, pp. 952–62.

Radford, A., Metz, L., and Chintala, S. (2015). "Unsupervised representation learning with deep convolutional generative adversarial networks". In: *arXiv preprint arXiv:1511.06434*.

Rao, R. and Georgieff, M. K. (2007). "Iron in fetal and neonatal nutrition". In: *Seminars in fetal and neonatal medicine*. Vol. 12. 1. Elsevier, pp. 54–63.

— (2009). "Iron therapy for preterm infants". In: *Clinics in perinatology* 36.1, pp. 27–42.

Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). "Stochastic backpropagation and approximate inference in deep generative models". In: *arXiv preprint arXiv:1401.4082*.

Rivkin, M. et al. (2004). "Prolonged T*2 values in newborn versus adult brain, Implications for fMRI studies of newborns". fr. In: *Magnetic Resonance in Medicine* 51.6, pp. 1287–91.

Rodrigue, K. M., Haacke, E. M., and Raz, N. (2011). "Differential effects of age and history of hypertension on regional brain volumes and iron". In: *Neuroimage* 54.2, pp. 750–759.

Ronneberger, O., Fischer, P., and Brox, T. (2015). "U-net: Convolutional networks for biomedical image segmentation". In: *International Conference on Medical image computing and computer-assisted intervention*. Springer, pp. 234–241.

Rouphael, T. J. (2009). *RF and digital signal processing for software-defined radio: a multi-standard multi-mode approach*. Newnes.

Rousseau, F., Glenn, O., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D., Barkovich, J., and Studholme, C. (2006). "Registration-based approach for reconstruction of high-resolution in utero fetal MR brain images". en. In: *Academic Radiology* 13.9, pp. 1072–1081.

Rousseau, F., Kim, K., Studholme, C., Koob, M., and Dietemann, J.-L. (2010). "On super-resolution for fetal brain MRI". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 355–362.

Saleem, S. (2014). "Fetal MRI, An approach to practice, A review". en. In: *Journal of Advanced Research* 5.5, pp. 507–523.

Salehi, S. S. M., Khan, S., Erdogmus, D., and Gholipour, A. (2018). "Real-time deep pose estimation with geodesic loss for image-to-template rigid registration". In: *IEEE transactions on medical imaging* 38.2, pp. 470–481.

Schieda, N., Isupov, I., Chung, A., Coffey, N., and Avruch, L. (2017). "Practical applications of balanced steady-state free-precession (bSSFP) imaging in the abdomen and pelvis". In: *Journal of Magnetic Resonance Imaging* 45.1, pp. 11–20.

Seip, M. and Halvorsen, S. (1956). "Erythrocyte Production and Iron Stores in Premature Infants During the First Months of Life The Anemia of Prematurity-Etiology, Pathogenesis, Iron Requirement". In: *Acta Paediatrica* 45.6, pp. 600–617.

Seshamani, S., Gatenby, C., Fogtmann, M., Cheng, X., Dighe, M., and Studholme, C. (2013a). "Combining R2* mapping and slice registration for fMRI analysis of moving subjects". en. In: *Proceedings of the Annual Meeting of International Society of Magnetic Resonance in Medicine*.

Seshamani, S., Blazejewska, A. I., Gatenby, C., Mckown, S., Caucutt, J., Dighe, M., and Studholme, C. (2015a). "Robust R*2 map estimation from motion scattered slices for fetal fMRI". In: *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*. IEEE, pp. 845–848.

Seshamani, S., Blazejewska, A. I., Gatenby, C., Mckown, S., Caucutt, J., Dighe, M., and Studholme, C. (2015b). "Comparing consistency of R2* and T2*-weighted BOLD analysis of resting state fetal fMRI". In: *Medical Imaging 2015: Biomedical Applications in Molecular, Structural, and Functional Imaging*. Vol. 9417. International Society for Optics and Photonics, 94170N.

Seshamani, S., Fogtmann, M., Cheng, X., Thomason, M., Gatenby, C., and Studholme, C. (2013b). "Cascaded slice to volume registration for moving fetal FMRI". In: *2013 IEEE 10th International Symposium on Biomedical Imaging*. IEEE, pp. 796–799.

Setsompop, K., Gagoski, B., Polimeni, J., Witzel, T., Wedeen, V., and Wald, L. (2012). "Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty". en. In: *Magnetic Resonance in Medicine* 67.5, pp. 1210–1224.

Siemonsen, S., Finsterbusch, J., Matschke, J., Lorenzen, A., Ding, X.-Q., and Fiehler, J. (2008). "Age-dependent normal values of T2* and T2' in brain parenchyma". In: *American Journal of Neuroradiology* 29.5, pp. 950–955.

Simonyan, K. and Zisserman, A. (2014). "Very deep convolutional networks for large-scale image recognition". In: *arXiv preprint arXiv:1409.1556*.

Smith, S., Andersson, J., Auerbach, E., Beckmann, C., Bijsterbosch, J., Douaud, G., and Duff, E. (2013). "Resting-state fMRI in the human connectome project". en. In: *NeuroImage* 80, pp. 144–68.

Smith, S. M. (2002). "Fast robust automated brain extraction". In: *Human brain mapping* 17.3, pp. 143–155.

Smith, T. B. and Nayak, K. S. (2010). "MRI artifacts and correction strategies". In: *Imaging in Medicine* 2.4, pp. 445–457.

Song, S., Zheng, Y., and He, Y. (2017). "A review of Methods for Bias Correction in Medical Images". In: *Biomedical Engineering Review* 1.1.

Speck, O. and Hennig, J. (1998). "Functional imaging by I0- and T2*-parameter mapping using multi-image EPI". en. In: *Magnetic Resonance in Medicine* 40.2, pp. 243–8.

Srinivasan, L., Dutta, R., Counsell, S., Allsop, J., Boardman, J., Rutherford, M., and Edwards, D. (2007). "Quantification of deep gray matter in preterm infants at term-equivalent age using manual volumetry of 3-tesla magnetic resonance images". fr. In: *Pediatrics* 119.4, pp. 759–65.

Stephenson, E., Nathoo, N., Mahjoub, Y., Dunn, J. F., and Yong, V. W. (2014). "Iron in multiple sclerosis: roles in neurodegeneration and repair". In: *Nature Reviews Neurology* 10.8, p. 459.

Straughen, J., Trudeau, S., and Misra, V. (2013). "Changes in adipose tissue distribution during pregnancy in overweight and obese compared with normal weight women". In: *Nutrition & diabetes* 3.8, e84.

Studholme, C., Hill, D., and Hawkes, D. (1999). "An overlap invariant entropy measure of 3D medical image alignment". en. In: *Pattern Recognition* 32.1, pp. 71–86.

Studholme, C. (2011). "Mapping fetal brain development in utero using magnetic resonance imaging: the Big Bang of brain mapping". In: *Annual review of biomedical engineering* 13, pp. 345–368.

Thompson, R. A. and Nelson, C. A. (2001). "Developmental science and the media: Early brain development." In: *American Psychologist* 56.1, p. 5.

Todorich, B., Pasquini, J. M., Garcia, C. I., Paez, P. M., and Connor, J. R. (2009). "Oligodendrocytes and myelination: the role of iron". In: *Glia* 57.5, pp. 467–478.

Tourbier, S., Hagmann, P., Meuli, R., and Cuadra, M. (2018). "Slice-by-Slice Intensity Inhomogeneity Correction from a Globally-Estimated Bias Fieldfor Fetal MRI Reconstruction". In: *Proceedings of the Swiss Congress of Radiology*.

Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., and Gee, J. C. (2010). "N4ITK: improved N3 bias correction". In: *IEEE transactions on medical imaging* 29.6, p. 1310.

Ugurbil, K., Xu, J., Auerbach, E., Moeller, S., Vu, A., Duarte-Carvajalino, J., and Lenglet, C. (2013). "Pushing spatial and temporal resolution for functional and diffusion MRI in the human connectome project". en. In: *NeuroImage* 80, pp. 80–104.

Uğurbil, K. et al. (2013). "Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project". In: *Neuroimage* 80, pp. 80–104.

Van Gelderen, P., De Zwart, J. A., Lee, J., Sati, P., Reich, D. S., and Duyn, J. H. (2012). "Nonexponential T2* decay in white matter". In: *Magnetic resonance in medicine* 67.1, pp. 110–117.

Vasylechko, S., Malamateniou, C., Nunes, R. G., Fox, M., Allsop, J., Rutherford, M., Rueckert, D., and Hajnal, J. V. (2013). "T2* Measurement of Fetal Brain Using a Motion Tolerant Method". In: *Proceedings of ISMRM*.

— (2015). "T2* relaxometry of fetal brain at 1.5 Tesla using a motion tolerant method". In: *Magnetic resonance in medicine* 73.5, pp. 1795–1802.

Wachowicz, K. (2014). "Evaluation of active and passive shimming in magnetic resonance imaging". In: *Research and Reports in Nuclear Medicine* 4, p. 1.

Wang, C., Zhang, X., Liu, X., He, T., Chen, W., Feng, Q., and Feng, Y. (2018). "Improved liver mapping by pixel-wise curve fitting with adaptive neighborhood regularization". In: *Magnetic resonance in medicine* 80.2, pp. 792–801.

Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., Tao, A., Kautz, J., and Catanzaro, B. (2018). "High-resolution image synthesis and semantic manipulation with conditional gans". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8798–8807.

Welander, P., Karlsson, S., and Eklund, A. (2018). "Generative adversarial networks for image-to-image translation on multi-contrast MR images-A comparison of CycleGAN and UNIT". In: *arXiv preprint arXiv:1806.07777*.

Wolterink, J. M., Dinkla, A. M., Savenije, M. H., Seevinck, P. R., Berg, C. A. van den, and Išgum, I. (2017). "Deep MR to CT synthesis using unpaired data". In: *International Workshop on Simulation and Synthesis in Medical Imaging*. Springer, pp. 14–23.

Wu, B., Li, W., Avram, A., Gho, S., and Liu, C. (2012). "Fast and tissue-optimized mapping of magnetic susceptibility and T2* with multi-echo and multi-shot spirals". en. In: *NeuroImage* 59.1, pp. 297–305.

Yanasak, N. E. and Kelly, M. J. (2014). "MR imaging artifacts and parallel imaging techniques with calibration scanning: a new twist on old problems". In: *Radiographics* 34.2, pp. 532–548.

Yang, Q., Li, N., Zhao, Z., Fan, X., Chang, E.-C., Xu, Y., et al. (2018). "MRI image-to-image translation for cross-modality image registration and segmentation". In: *arXiv preprint arXiv:1801.06940*.

Yao, B., Li, T.-Q., Gelderen, P. van, Shmueli, K., Zwart, J. A. de, and Duyn, J. H. (2009). "Susceptibility contrast in high field MRI of human brain as a function of tissue iron content". In: *Neuroimage* 44.4, pp. 1259–1266.

Yi, X., Walia, E., and Babyn, P. (2019). "Generative adversarial network in medical imaging: A review". In: *Medical Image Analysis*, p. 101552.

Yu, B., Zhou, L., Wang, L., Fripp, J., and Bourgeat, P. (2018). "3D cGAN based cross-modality MR image synthesis for brain tumor segmentation". In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, pp. 626–630.

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. (2018). "The unreasonable effectiveness of deep features as a perceptual metric". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595.

Zhang, Z., Yang, L., and Zheng, Y. (2018). "Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9242–9251.

Zhong, L., Lin, L., Lu, Z., Wu, Y., Lu, Z., Huang, M., Yang, W., and Feng, Q. (2016). "Predict CT image from MRI data using KNN-regression with learned local descriptors". In: *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*. IEEE, pp. 743–746.

Zhu, C., Wang, Y., Wu, Y., and Wang, X. (2019). "Iron metabolism and brain development in premature infants". In: *Frontiers in physiology* 10, p. 463.

Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. (2017). "Unpaired image-to-image translation using cycle-consistent adversarial networks". In: *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232.