Noise reduction during diffusion tensor imaging of infants

by

Johannes Petrus Jordaan
JRDJOH011

In fulfilment of the requirements for the degree of Master of Science (Medicine) in Biomedical Engineering

Faculty of Health Sciences
University of Cape Town
South Africa
February 2019

Supervisors:
Prof. Ernesta Meintjes
Department of Human Biology
University of Cape Town

Dr. André van der Kouwe
A.A. Martinos Center for Biomedical Imaging

Declaration of Authorship

I, JP Jordaan, hereby declare that the work on which this dissertation/thesis is based is my original work (except where acknowledgements indicate otherwise) and that neither the whole work nor any part of it has been, is being, or is to be submitted for another degree in this or any other university.

I empower the university to reproduce for the purpose of research either the whole or any portion of the contents in any manner whatsoever.

Signature: Signed by candidate

Date: 29/09/2019

Dedication

I dedicate this work to my wife, Elaine, for her endless love and unconditional support.

Acknowledgements

I would like to sincerely thank both my supervisors for without them none of this would have been possible. To Prof. Ernesta Meintjes, who granted me the extraordinary opportunity to work on this project and whose dedication and guidance allowed this project to reach conclusion. To Dr. André van der Kouwe for his deep technical insight and support as well as his personal support during my phenomenal trip to the Martinos Centre for Biomedical Imaging in Boston, Massachusetts, USA. I would also like to thank Adam van Niekerk and Stephen Jermy for their technical advice and long discussions regarding MRI-related things. Lastly, I would like to thank my family for all their love and support, I am truly grateful.

This work was funded by grants from the National Institutes of Health (R01 HD085813), the South African National Research Chairs Initiative, and the University of Cape Town through the Max & Lillie Sonnenberg Scholarship.

UNIVERSITY OF CAPE TOWN

Abstract

Faculty of Health Sciences
Department of Human Biology

Master of Science (Medicine) in Biomedical Engineering

Noise reduction during diffusion tensor imaging of infants

by JP Jordaan

Acoustic noise produced during echo planar imaging (EPI) has been known to reach excessive levels. In addition to causing general patient discomfort and anxiety, this level of noise makes the scanning of young children and infants particularly difficult. Infants are typically scanned while sleeping to minimise motion as they cannot ethically be sedated for research purposes. The extreme noise during MRI acquisitions often cause them to wake before the end of the scanning session. This problem is exacerbated by particularly noisy acquisitions, such as the single shot echo planar diffusion tensor imaging sequence. The main aim of this project was to reduce the noise of this particular acquisition specifically for the scanning of newborn infants.

Acoustic noise during MRI acquisitions mainly originates from mechanical vibrations in the gradient coil assemblies due to interactions between the rapidly changing currents applied to the coils and the main static field. A transfer function relating the output acoustic noise spectrum to the gradient excitation input spectrum was developed and used to identify resonant peaks which would amplify coinciding gradient waveform harmonics. In addition to resonant peaks, the transfer function showed significant amplification of frequencies above 1 kHz.

In this work, noise reduction was achieved by implementing digital low-pass filters to reduce high-frequency harmonics of the standard trapezoidal gradient waveforms, focusing on the EPI readout portion of a diffusion tensor imaging (DTI) sequence. For comparison purposes, an EPI readout using sinusoidal frequency encoding waveforms and a constant phase encoding blip was also implemented. In addition to reducing produced noise, a passive noise reduction enclosure was built from open cell polyurethane foam mounted in a PVC frame to surround the sleeping infant and act as an acoustic insulation box. Lastly, the effectiveness of introducing pink noise from an external source to mask the abrupt changes in scanner noise, was also investigated.

The altered k-space trajectories due to the modifications made to the EPI readout gradient waveforms were corrected through a custom one-dimensional regridding procedure applied along the frequency encoding axis in k-space.

Noise reduction was measured with an Optimic 1155 optical microphone from Optoacoustics, attached on top of a cylindrical water phantom inside a 16 channel

infant head coil in the isocenter along the Z-direction, facing the bore in the right-left direction (similar to the orientation that the ears of a sleeping infant would be).

Signal-to-noise ratio (SNR) and fractional anisotropy (FA) within the corpus callosum (CC) were compared for images acquired using the standard and modified (filtered and sinusoidal readouts) DTI sequences, the latter each for regridding kernel window sizes of 2 and 4, respectively. The acoustic noise spectra of the filtered and sinusoidal EPI sequences demonstrated a significant reduction in EPI harmonics compared to the standard sequence, but very little difference between each other. Without the foam enclosure, the filtered acquisition with filtered crushers reduced peak sound pressure levels (SPL) by 3.4 and 4 dBA for strong and no fat suppression, respectively, and A-weighted equivalent continuous sound levels ($L_{A,eq}$) by 2.5 and 2.8 dBA, respectively. Adding the foam enclosure increased peak SPL reduction to 4.8 dBA with fat suppression and 7 dBA without. The sinusoidal sequence performing similarly or marginally (no more than 0.5%) worse than the filtered on all outcomes.

SNR measurements in the CC were higher for all volumes of the filtered acquisition compared to the standard, while those of the sinusoidal were similar or slightly lower compared to the standard acquisition. FA values in the CC of the sinusoidal and filtered acquisitions did not differ from those of the standard acquisition (pairwise student's t-test, all p's >0.2). For the 16 channel head coil, image reconstruction time increased by only 45 seconds for a regridding kernel width $W = 2$.

Filtering gradient waveforms is an effective technique to reduce acoustic noise during DTI without increasing acquisition time, reducing image quality, or altering FA measures. The proposed method has the potential to be generalized to most gradient waveforms across a variety of sequences. With the addition of the passive noise reduction enclosure, the combined noise reduction could greatly reduce infant anxiety and startling, leading to an increase in the number of infants in whom the acquisition protocol is completed.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Problem identification

Acoustic noise produced during echo planar imaging (EPI) have been measured to reach 132 dBA (Foster et al., 2000). In addition to causing patient discomfort and anxiety, high levels of acoustic noise have been shown to influence results from functional magnetic resonance imaging (MRI) (Zhang, Zhu, and Chen, 2005; Tomasi et al., 2005). This level of noise makes the scanning of young children and infants particularly difficult. Infants are typically scanned while sleeping to minimise noise as they cannot ethically be sedated for research purposes. The extreme noise during MRI acquisitions often cause them to wake before the end of the scanning session. This problem is exacerbated by particularly noisy acquisitions, such as single shot echo planar diffusion tensor imaging (DTI) sequences. The main aim of this project was to reduce the noise of this particular acquisition specifically for the scanning of newborn infants.

1.2 Significance and Rationale

Acoustic noise produced during MRI acquisitions mainly originate from the gradient coil assemblies due to interactions between the rapidly changing currents applied to the coils and the main static magnetic field. The coil assemblies experience time-varying Lorentz forces resulting in mechanical vibrations that cause the acoustic noise heard during scans. The resulting noise is highly dependent on the scanner's geometry and material properties. The noise produced during EPI sequences is especially loud due to the fast switching of the polarity of the readout gradient.

Various hardware modifications have been proposed to reduce the scanner acoustic noise. Edelstein et al. (2002) proposed physical changes to the structure of the scanner including placement of the gradient assembly in a vibrationally isolated vacuum, the use of a low-eddy-current radio frequency (RF) coil, and a non-conducting inner bore cryostat. As a follow-on, they investigated effectiveness of a passive copper shield fitted to the outside of the gradient coil assembly to reduce mechanical power deposition in the warm bore (Edelstein et al., 2005). The design of new gradient coils less susceptible to the vibrations caused by Lorentz forces (Forbes et al., 2007) and the development of a mathematical model for imaging in non-linear magnetic fields, which could potentially reduce vibration in the gradient coils compared to homogeneous linear fields (Jackson et al., 2011), have also been investigated. Ma-

jor limitations of these hardware methods are the cost and complexity involved in hardware modifications.

Although some active noise cancellation (ANC) earphones are commercially available, such as the OptoACTIVE^{TM} active noise control earphones from Optoacoustics, they are not suitable for use on the small heads of infants and are expensive.

Other than earmuffs and earplugs, passive methods involve using foam to acoustically insulate the head and / or body of the patient (Ravicz and Melcher, 2001; Moelker, Vogel, and Pattynama, 2003; Nordell et al., 2009). For example, Nordell et al. (2009) achieved a peak noise reduction of 22 dBA during a DTI spin echo EPI acquisition by covering the inside of the bore with a 104 mm thick polyurethane foam insert covered by a 4 mm vinyl lining.

Software approaches (Hennel, Girard, and Loenneker, 1999; Hennel, 2001; Schmitter et al., 2008; Kannan et al., 2011; Hutter et al., 2018) have largely been based on the independent linear response theory by Hedeen and Edelstein (1997) who recognised the strong correlation between the spectrum of a gradient waveform and the spectrum of the resulting acoustic noise. The authors developed a scanner specific frequency response function (FRF) that relates the spectrum of the gradient coil input waveform to the acoustic noise output. Figure 1.1 shows an example of such a frequency response transfer function for the 3 T Siemens Skyra MRI (Erlangen, Germany) used in this study. This work paved the way for subsequent approaches that attempt to minimise gradient waveform spectral content and to confine the remaining components to local minima in the FRF.

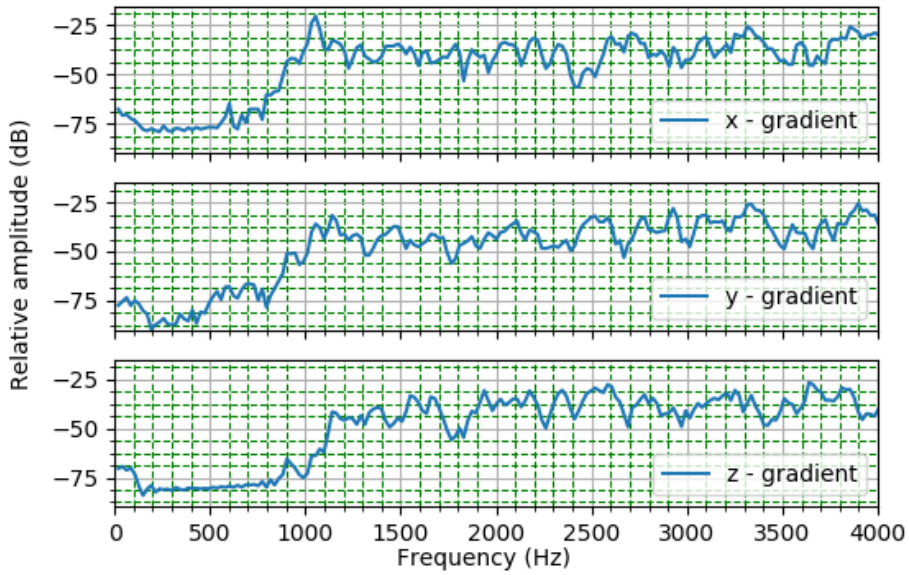

Figure 1.1: Frequency response function for x, y, and z gradient coil assemblies in a 3 T Siemens Skyra (Erlangen, Germany) scanner acquired using a protocol similar to Wu et al. (2014).

"Soft" pulses use long sinusoidal ramps to limit the pulse spectrum to lower frequencies where acoustic noise is less. Hedeen and Edelstein (1997) using this method,

which is restricted to slow sequences, achieved scanner noise below the ambient noise of the air circulation system in the room. Subsequently, Hennel (2001) demonstrated noise reduction of up to 40 dBA (A-weighted decibel scale) in fast gradient echo (FLASH) and spin echo (RARE) sequences by changing the readout gradient waveform to an entirely sinusoidal shape with no plateau and extending the phase encoding pulses throughout the entire readout period. The resulting curved k-space trajectories and non-uniform sampling densities were corrected using a 1D regridding procedure.

A similar approach has been applied to the EPI readout portion of a functional MRI (fMRI) sequence where the traditional trapezoidal readout gradients were replaced with a sinusoid and a constant phase encoding blip was implemented (Schmitter et al., 2008). By additionally replacing the readout pre-winder of moment half the k-space width by a sinusoidal waveform with a linearly increasing amplitude and a total moment equal to that of the original pre-winder, and appending a sinusoidal wave with linearly decreasing amplitude to the end of the EPI readout to prevent the abrupt ending of the waveform, a microphone placed on top of a water phantom inside a head coil measured noise reduction of 22.5 dBA. Signal-to-noise-ratio (SNR) for the modified sequence was approximately 15% lower than with the standard sequence.

Another approach aims to avoid the machine specific resonance peaks contributing most to the noise power by using longer echo spacings in EPI and longer repetition times (TR) in other sequences. This method has been shown to reduce sound pressure levels (SPL) of various sequences by 3.9 dB on average (Smink, Plattel, and Harvey, 2007).

More recently, Hutter et al. (2018) developed a quieter EPI sequence dubbed QuEPI for fetal fMRI and diffusion scans by implementing a sinusoidal EPI readout and constant phase encoding blip with the addition of overlapping smoothed CAIPIRINHA ("controlled aliasing in parallel imaging results in higher acceleration") blips and merged crusher gradients. The EPI fundamental frequency was tuned to coincide with a trough in the FRF to further reduce noise output. For DTI acquisition, compared to a standard sequence with EPI echo spacing 0.57 ms, QuEPI with echo spacing 0.99 ms reduced acoustic noise output by 9 dBA (from 112.6 dBA to 103.6 dBA). Compared to the scanner optimized EPI sequence, *in vivo* images acquired using QuEPI demonstrated a slight reduction in SNR (19.7 to 18.2).

Although sinusoidal ramps (Hennel, Girard, and Loenneker, 1999) and sinusoidal readout pulses (Hennel, 2001; Schmitter et al., 2008; Hutter et al., 2018) previously implemented achieved effective noise reduction, these methods lead to increased acquisition times and k-space complexities, respectively.

In the present work, because of the high cost of hardware implementations, we sought to find a software-based solution that would reduce the spectral content of the gradient waveforms and avoid resonance peaks in the FRFs, without some of the limitations of previous methods. Since the EPI readout portion, which involves rapid switching of the polarity of the frequency encoding readout gradients as well

as sharp, impulse-like phase encoding blips, generates a significant amount of high frequency noise compared to the rest of the sequence, reducing noise associated with this portion of the sequence was the main focus of this study. In contrast to previous studies that implemented low-pass filters by replacing different parts of the sequence with sinusoidal components and/or by using analytical methods (Hennel, Girard, and Loenneker, 1999; Hennel, 2001; Schmitter et al., 2008; Hutter et al., 2018), we propose a solution using modern digital filters that can easily be integrated into any sequence using time-optimised trapezoidal pulses. This would reduce implementation complexities by avoiding the need to redesign individual waveforms.

In addition to the sequence modifications, we built an acoustic noise insulating enclosure and tried to reduce the "startle response" that results from the abrupt increase in noise when scanning starts by continuously piping pink noise into the scanner room via the built-in scanner speaker.

1.3 Aim

The aim of this project was to reduce the acoustic noise produced during diffusion tensor imaging of sleeping infants while maintaining image quality and keeping the required scan time to a minimum.

1.4 Objectives

- To develop a software-based method to reduce the acoustic noise produced during magnetic resonance diffusion tensor imaging while keeping the scan time to a minimum.
- To validate the acoustic noise reduction from the new pulse sequence against the previously used sequence by comparing sound pressure level (SPL) measurements of the two sequences.
- To ensure no significant reduction in image quality or change in diffusion parameters of the new sequence compared to the original sequence.

1.5 Background

1.5.1 Basics of Magnetic Resonance Imaging (MRI)

Magnetic Resonance Imaging (MRI) is a non-ionising imaging modality that uses of a large magnetic field and radio frequency (RF) pulses to create a signal that is manipulated using gradient fields to form an image. The RF pulses are designed to excite hydrogen protons, allowing images to be taken of living tissue due to its abundance in water and fat in the body.

In the scanner, a large static magnetic field ($\mathbf{B_0}$) is generated using a superconductive coil submerged in liquid helium. For clinical MRI, B_0 values range from 1.5 to 3 Tesla (T). In addition to the main magnet, three mutually orthogonal (X, Y, Z), spatially varying magnetic fields are superimposed on $\mathbf{B_0}$, causing the field strength to vary with position in the scanner bore.

Hydrogen protons are constantly spinning, causing a small magnetic field called a magnetic dipole moment. In the presence of the large $\mathbf{B_0}$ magnetic field, these dipole moments experience a torque that causes them to precess around $\mathbf{B_0}$ (Figure 1.2) at the Larmor frequency (ω_L), given by

$$\omega_L = \gamma B_0, \tag{1.1}$$

where γ is the gyromagnetic ratio, namely 42.6 MHz/T for H^1. The effect of millions of precessing dipole moments is a net magnetization, $\mathbf{M_0}$, aligned parallel to the main magnetic field.

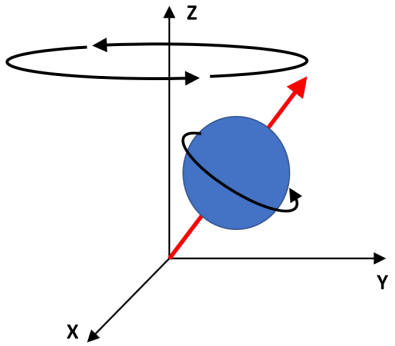

Figure 1.2: Representation of a proton precessing about the z-axis due to the presence of an external magnetic field along same direction.

Applying an RF pulse at the Larmor frequency of the hydrogen protons in a coil perpendicular to the main magnetic field, establishes an oscillating magnetic field perpendicular to the main magnetic field that tips $\mathbf{M_0}$ towards the transverse plane orthogonal to the longitudinal Z-axis. Upon termination of the RF pulse, the dipole moments again precess around $\mathbf{B_0}$ and return to their equilibrium distribution with $\mathbf{M_0}$ pointing along the main magnetic field. The flux induced in the receiver coils by the transitory transverse component of $\mathbf{M_0}$ **(Mxy)** is the MRI signal and is termed a free induction decay (FID, Figure 1.3).

The return of the magnetization to its original value along the main magnetic field comprises two distinct and independent processes, called T1 and T2 relaxation. T1 relaxation, also known as longitudinal or spin-lattice relaxation, involves the release of the energy absorbed during RF excitation back to the lattice and determines the rate at which the Z-component of the magnetization re-grows to its original maximum value along B0. T2 relaxation, also termed spin-spin dephasing, results from the fact that every individual hydrogen nucleus experiences a slightly different net magnetic field due to its unique chemical environment causing it to precess at a unique frequency and a collection of nuclei to gradually get more and more out of phase with each other. The rate at which this dephasing occurs determines the rate at which the transverse component of the magnetization, which is measured by the receiver coil, is lost. Excitation, T1 and T2 relaxation are depicted in Figure 1.4.

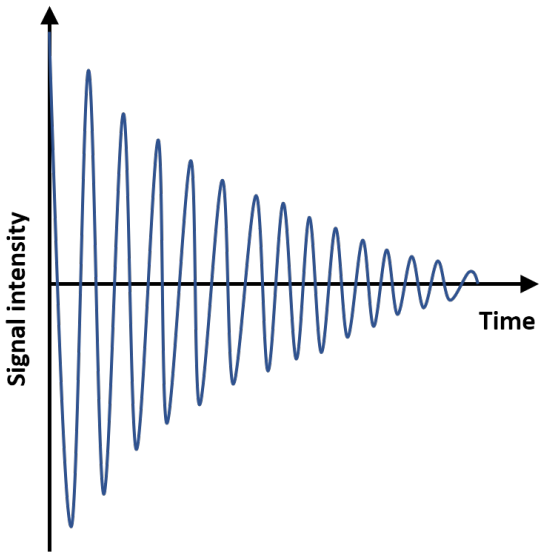

Figure 1.3: Example of a free induction decay (FID) signal directly after a 90° RF pulse.

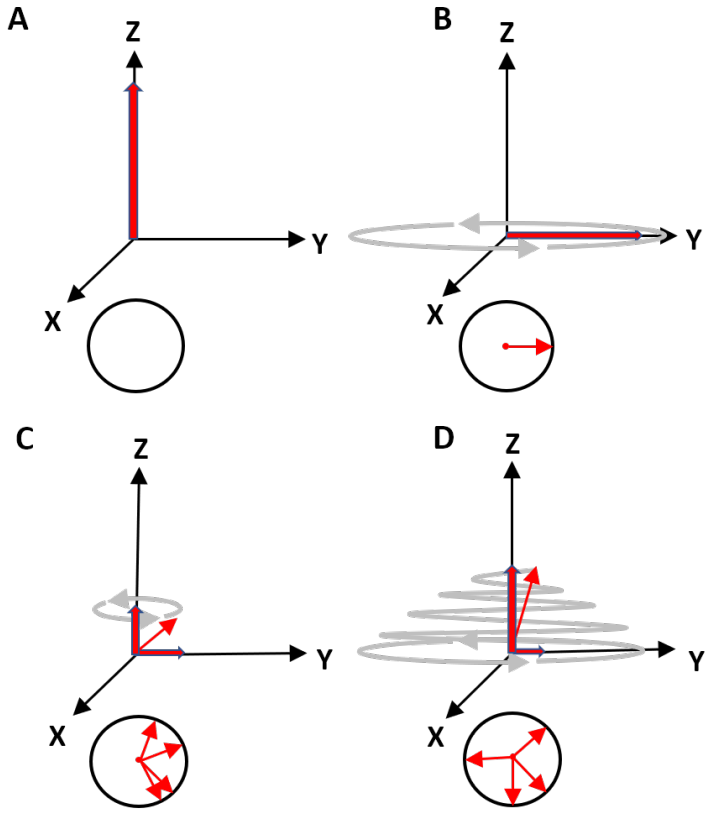

Figure 1.4: Recovery after transverse magnetization due to a 90° RF pulse. Before the RF pulse the net magnetisation of the protons is in the direction of the B0 field along the Z-axis (A). After the 90° RF pulse the net magnetization is in the X-Y plane, with all the proton spins in phase with one another (B). Some of the spins begin to dephase (T2 relaxation) and revert back to alignment with the B_0 field (T1 relaxation), with the net magnetization precessing about, and returning to alignment with the Z-axis (C and D).

In spin echo (SE) sequences the echo of the FID signal is measured rather than the FID signal itself. The echo is formed by applying a 180° RF pulse to the dephased transverse magnetised proton vectors at a specified time, which reverses the precession direction causing the transverse vectors to rephase after that time again. The time after the RF pulse till when the echo is formed is termed the echo time (Figure 1.5).

To selectively excite the protons in a given slice, a linearly varying gradient field (G_S) is superimposed onto the B_0 field in the slice-selection direction (Figure 1.6), which is perpendicular to the desired imaging slice. This gradient causes protons at different positions along the gradient direction to precess at different Larmor frequencies. By applying a 90° RF pulse with a small bandwidth during the time when this gradient is applied, only protons with Larmor frequencies around the centre frequency of the 90° RF pulse are excited.

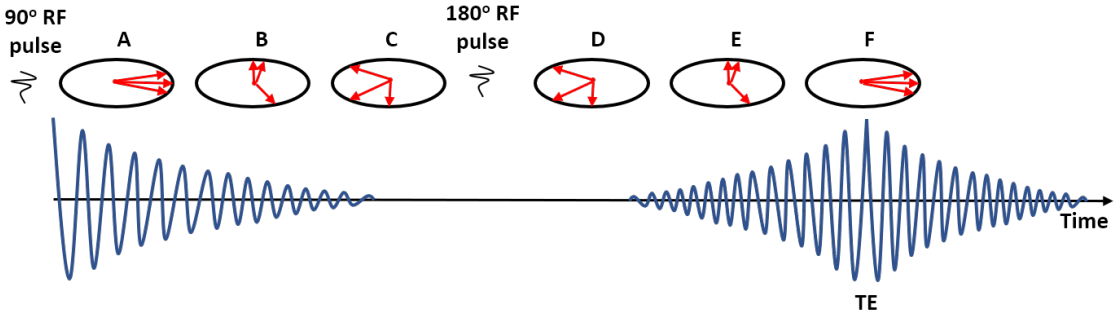

Figure 1.5: For spin echo, the T2 relaxation in reversed by applying a 180° RF pulse at half the echo time (TE) which allows the spins to realign and form an echo of the free induction decay (FID) signal at TE.

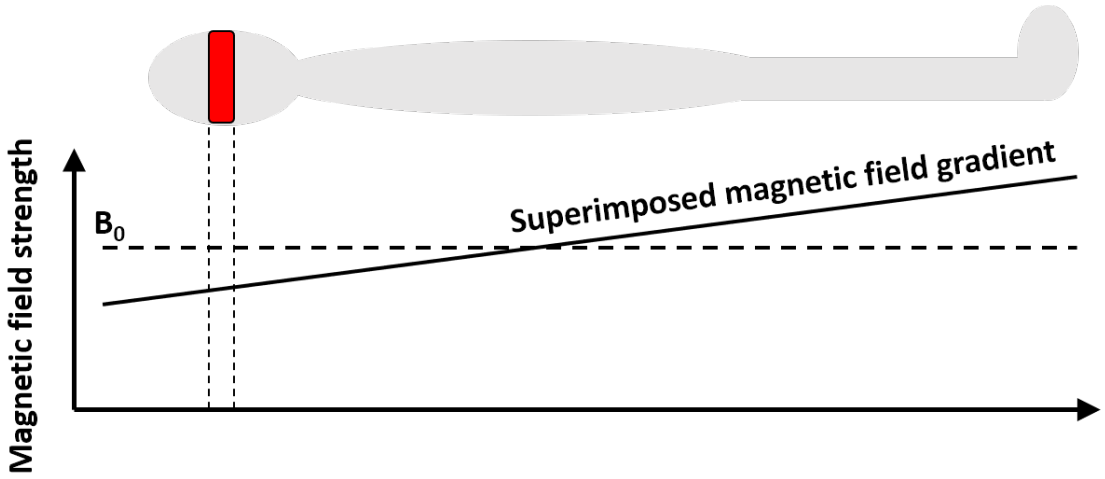

Figure 1.6: Slice selection is performed by superimposing a gradient field onto the B_0 field which causes the protons to precess at different frequencies along that direction. This allows for the excitation of a specific slice by applying an RF pulse with a small bandwidth centred in the region of interest (red).

Spatial encoding within an excited slice is achieved using orthogonal frequency (G_F) and phase (G_F) encoding gradients, respectively. After an RF pulse, all protons within the excited slice precess in phase with one another. By applying a magnetic gradient along the phase encoding axis, the protons spin either faster or slower, depending on the magnetic gradient at each location. After the gradient is removed, all the excited protons precess at the same rate but are now out of phase with respect to one another, depending on their position along the phase encoding direction. This gradient is applied before the measurement of the echo takes place. The frequency encoding gradient, which is applied during the measurement of the echo, causes protons to precess at different frequencies relative to their position along the frequency encoding gradient. This process is depicted in Figure 1.7.

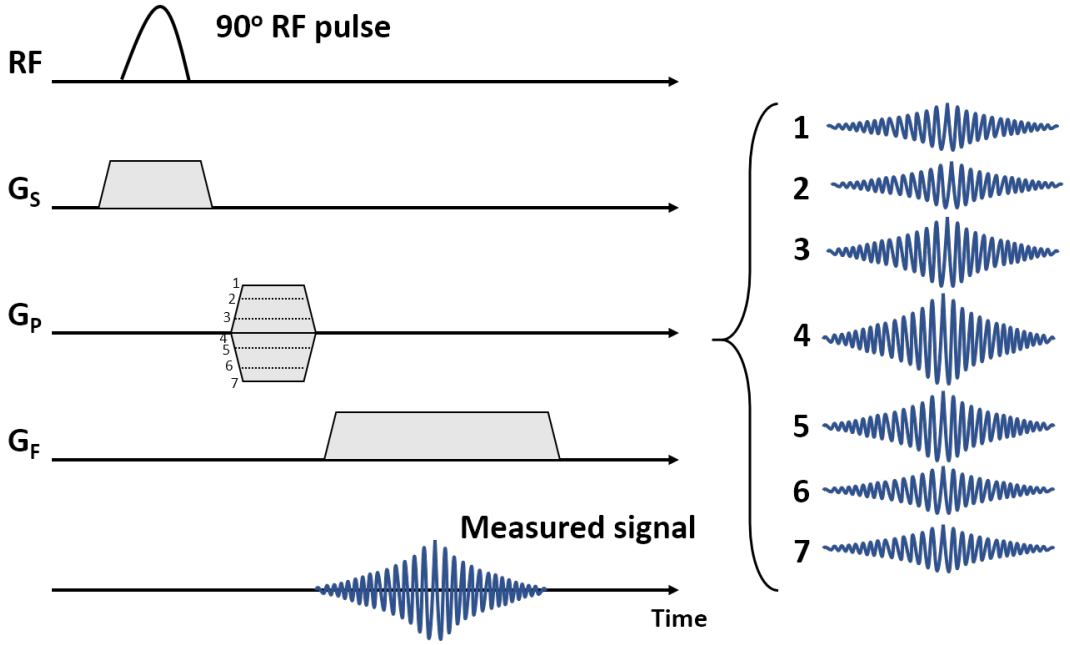

Figure 1.7: Simplified diagram of a spin echo sequence. After a specific slice has been excited using an RF pulse and the slice selection gradient (G_S), the measured echoes are spatially encoded by altering the size of the phase encoding gradient (G_F) in subsequent measurements and applying a frequency encoding gradient (G_F) during the measurement of the echoes.

The spatially encoded information is recorded into a matrix called k-space, with each row containing a recorded echo for a different phase encoding gradient. Thus k-space is an array of numbers that are frequency encoded in the X-axis, phase encoded in the Y-axis, and the intensity of each data point represents the relative contribution of that data point's frequency and phase to the final image. A typical representation of k-space is shown in Figure 1.8 where each point in k-space contains frequency and phase information about every pixel in the final image. The image is produced by performing a 2D Fourier transform on the k-space matrix which requires that the data points in the matrix be equally spaced in a rectangular grid. If for some reason the sampling in k-space is not equally spaced, some sort of interpolation must be implemented to correct this before the 2D Fourier transform can be applied. This is called regridding.

Instead of acquiring one line of k-space at a time, echo planar imaging (EPI) allows the acquisition of an entire slice, or 2D k-space matrix, with a single 90° RF excitation (hence single shot EPI). This is done by applying strong frequency encoding gradients of alternating polarity, separated by low-magnitude phase encoding gradients. This process is illustrated in Figure 1.9.

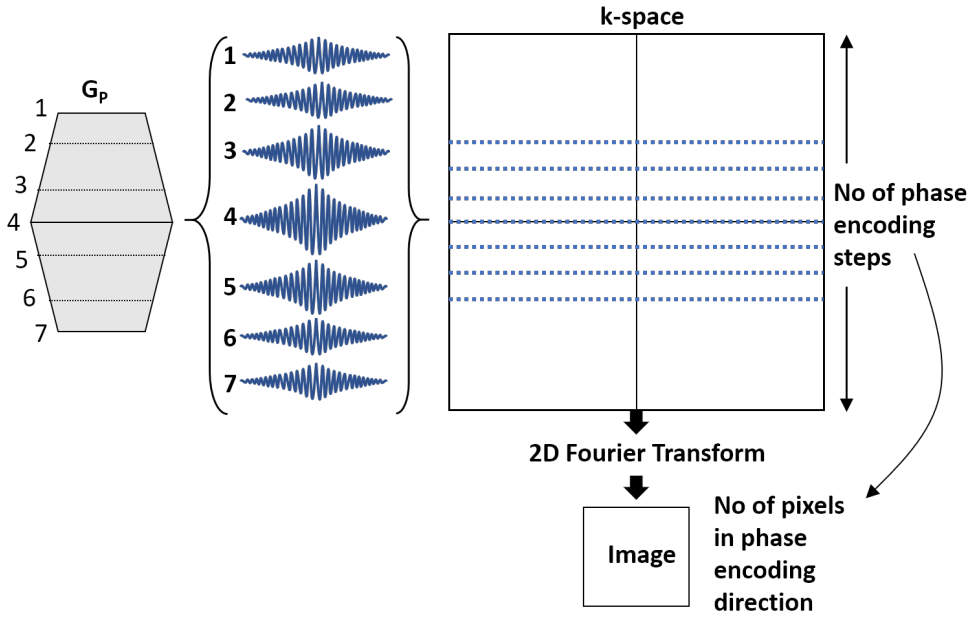

Figure 1.8: Each spin echo measurement (Figure 1.7) fills one line of k-space. By applying a 2D Fourier Transform on the equally spaced samples in the k-space matrix, the final image is acquired.

1.5.2 Diffusion tensor imaging (DTI)

Diffusion-weighted imaging (DWI) is an MRI modality that allows for the observation of molecule diffusion rates and directions within the architecture of different tissues. This modality is particularly useful for imaging white matter (WM) in the brain since diffusion is less restricted along the axons of neurons than perpendicular to it, making diffusion in white matter more anisotropic than in grey matter or cerebrospinal fluid (CSF). Isotropic diffusion means the extent of diffusion is equal in all directions. Conversely, anisotropic diffusion refers to directional diffusion similar to water in a pipe where higher levels of diffusion can be observed along the pipe than in directions orthogonal to the pipe.

For spin echo DWI, large diffusion-sensitizing gradient fields are applied before and after the 180° inversion pulse in the spin echo EPI sequence in Figure 1.10. The signal from static protons at TE is unaffected by the diffusion gradients as the second diffusion gradient simply reverses the phase accumulation caused by the first. For diffusing protons, however, the phase accumulation is not reversed due to the movement of the protons to different positions between the first and second diffusion gradient, causing a loss of signal. The strength of the diffusion effects is determined

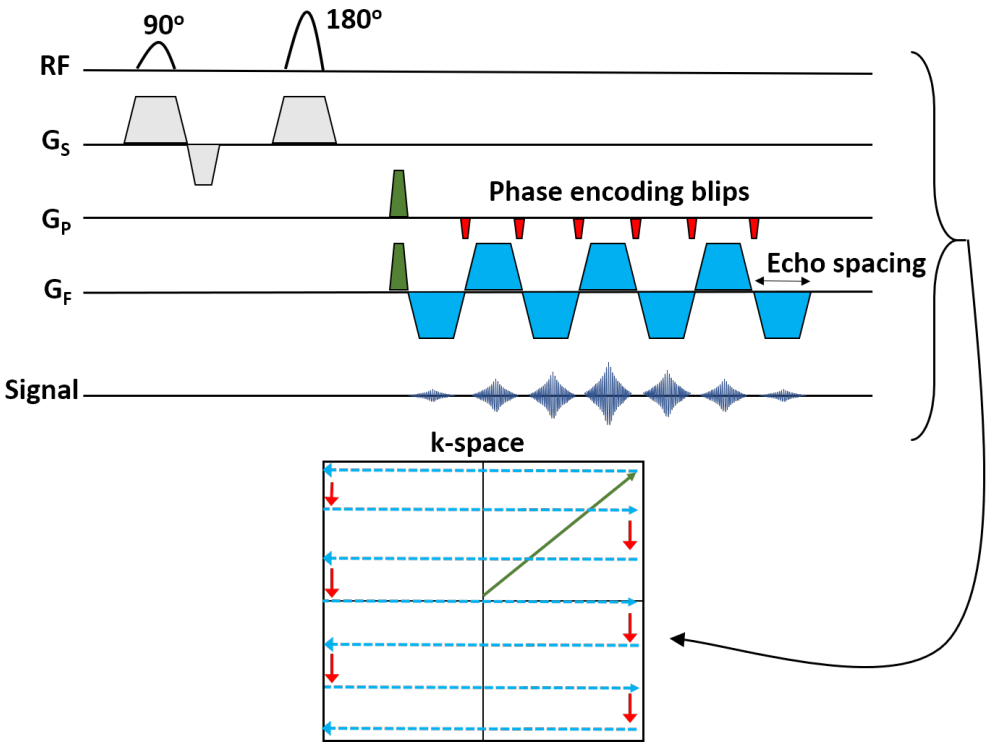

Figure 1.9: Spin echo (SE) echo planar imaging (EPI) is an accelerated version of the simple spin echo sequence which allows the acquisition of multiple k-space lines from one 90° RF excitation. The implementation of the phase encoding blips (G_P) and the alternating polarity frequency encoding gradients (G_F) places the measured echoes in a zig-zag like trajectory in k-space.

by a parameter called the b-value, which depends on the strength, duration, and spacing of the diffusion gradients. Higher b-values values cause stronger diffusion weighting and a b-value of zero (b0) corresponds to a purely T2-weighted image. Diffusion-weighted images can be sensitive to diffusion in different directions by altering the orientation of the diffusion gradients using different combinations of the gradient coils (X, Y, Z).

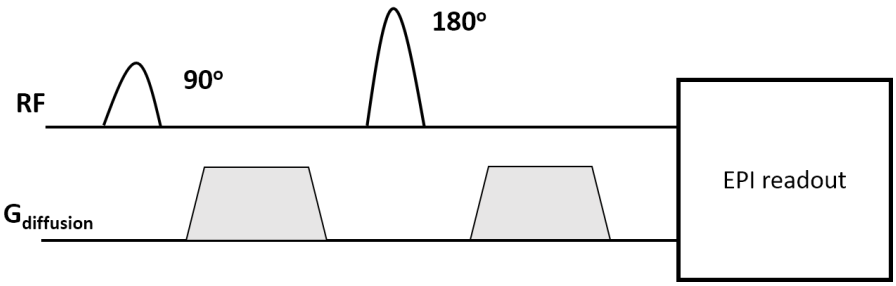

Figure 1.10: Basic spin echo diffusion-weighted imaging (DWI) sequence with echo planar imaging readout. Diffusion-sensitizing gradients are applied before and after the 180° RF inversion pulse.

While DWI refers to a specific type of image contrast, DTI refers to a method of modelling the DWI data. The diffusion tensor D is a 3x3 matrix that describes diffusion rates in a combination of directions and is given by

$$D = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix}, \tag{1.2}$$

where the elements D_{xx}, D_{yy}, and D_{zz} represent diffusion along the principal axes of the scanner (X, Y, Z). The tensor can be uniquely solved by acquiring at least one baseline b0 dataset where the diffusion gradients are turned off and at least six datasets with non-zero b-values (b=1000 in this project) and the diffusion-sensitizing gradients applied in different directions. The tensor is visualized using an ellipsoid with the main axis parallel to the principal diffusion direction within each voxel (3-dimensional pixel used to represent a volume). Figure 1.11 depicts an ellipsoid with the principal axis aligned with neuron axons, showing anisotropic diffusion along the axons.

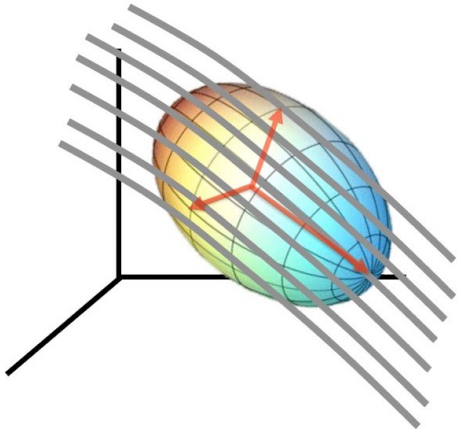

Figure 1.11: Ellipsoid generated from the diffusion tensor matrix in Equation 1.2 showing anisotropic diffusion along the main axis of the ellipsoid in the direction of neuronal axons (grey lines).

From the tensor, fractional anisotropy (FA) maps can be constructed. FA is a scalar measure of the directionality of diffusion within a voxel. FA values range between 0 and 1 where 0 corresponds to pure isotropic diffusion (sphere) and 1 to pure anisotropic diffusion. FA maps are thus grey-scale images with bright areas representing highly anisotropic diffusion (Figure 1.12 (Left)). Coloured FA maps add information regarding the principal direction of the tensor ellipsoid in each voxel in addition to the level of anisotropy (brightness) by assigning colour to the voxel depending on the direction of the principal axis of the tensor ellipsoid (Figure 1.12 (Right)).

1.5.3 Frequency response function (FRF)

The FRF is based on the independent linear response theory developed by Hedeen and Edelstein (1997) and validated by Wu et al. (2014), who experimentally confirmed model homogeneity and superposition errors for each gradient axis (X, Y, Z) less than 3%. Total noise predictions for all three gradient axes combined were accurate to within 4%.

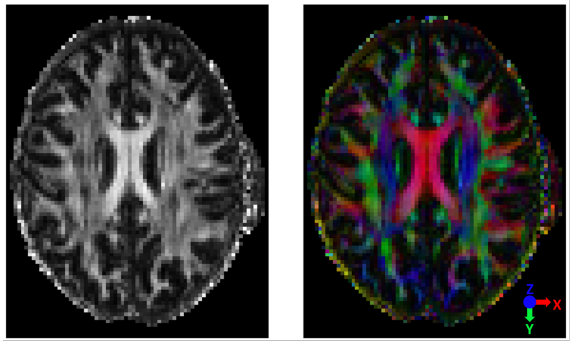

Figure 1.12: Fractional anisotropy (FA) maps showing white matter tracts in a human brain. The brightest part in the center of the brain (left) is the corpus callosum which consists almost entirely of white matter and exhibits highly anisotropic diffusion. The coloured FA map (right) indicates the principal direction of diffusion in each voxel according to the key on the bottom right.

If the physical structure of the MRI scanner is treated as a linear electromechanical system with an arbitrary gradient input $g(t)$, the output acoustic noise function $p(t)$ is given by the time domain convolution of the input $g(t)$ and the system impulse response $h(t)$,

$$p(t) = \int_0^\infty h(\tau)g(t-\tau)d\tau. \tag{1.3}$$

Since convolution of two signals in the time domain is equivalent to the multiplication of the same two signals in the frequency domain, Equation 1.3 becomes

$$P(f) = H(f)G(f). \tag{1.4}$$

Rearranging Equation 1.4 gives the system frequency response function $H(f)$ as a complex valued ratio of the output acoustic spectrum, $P(f)$, to the gradient excitation input spectrum, $G(f)$ (Hedeen and Edelstein, 1997).

$$H(f) = \frac{P(f)}{G(f)} \tag{1.5}$$

Since a unit impulse is physically unrealisable, white noise, band-limited to the maximum slew-rate achievable by a specific scanner, could be used as the gradient excitation input. The system FRF can then be computed from the Fourier transform of the recorded acoustic noise. Note, however, that the FRF will only be valid for the unaltered frequency band of the white noise.

1.5.4 Acoustics

The range of sound pressures levels perceivable by a young, healthy adult ear is extremely large, with the lowest detectable pressure difference being $p_0 = 20\,\mu Pa$ at 4 kHz. Sound pressure level, SPL or (L_p), in decibel (dB) is a measure of sound pressure p relative to the lowest perceivable level p_0, as given in Equation 1.6.

$$SPL = 10log_{10}\frac{p_{rms}^2}{p_0^2} = 20log_{10}\frac{p_{rms}}{p_0} = 20log_{10}p_{rms} - 20log_{10}p_0 \tag{1.6}$$

For convenience, Equation 1.6 can be expressed as

$$SPL = 20log_{10}p_{rms} + 94dB. \tag{1.7}$$

The sensitivity of the human ear differs depending on frequency. Frequency weighting networks "weight" the contributions of different frequencies to the overall SPL to account for the differing sensitivities of the human ear to different frequencies in the audible range. The "A" weighting network is an internationally standardised weighting network with its characteristics specified in the IEC 60651. Due to the popularity of this network in the literature, all noise measurements in this project were reported in dBA. Figure 1.13 shows frequency weighting curves for the A-weighting network along with the B, and C-weighting networks.

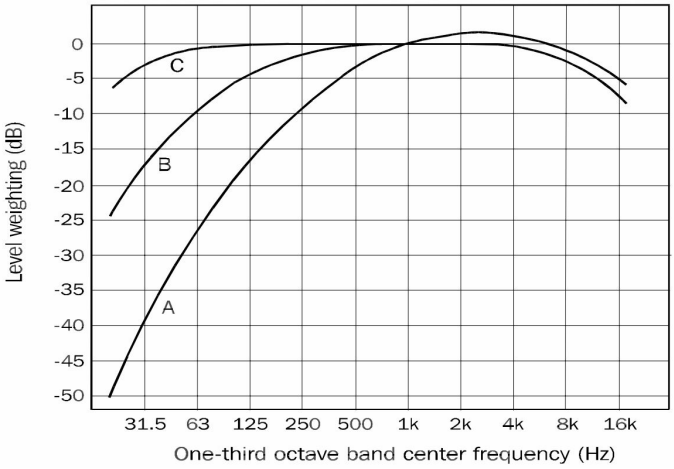

Figure 1.13: Frequency weighting characteristics for A, B, and C networks (Berenice Goelzer, Colin H. Hansen, 2001).

Typical sound pressure levels produced by various noise sources are listed in Figure 1.14.

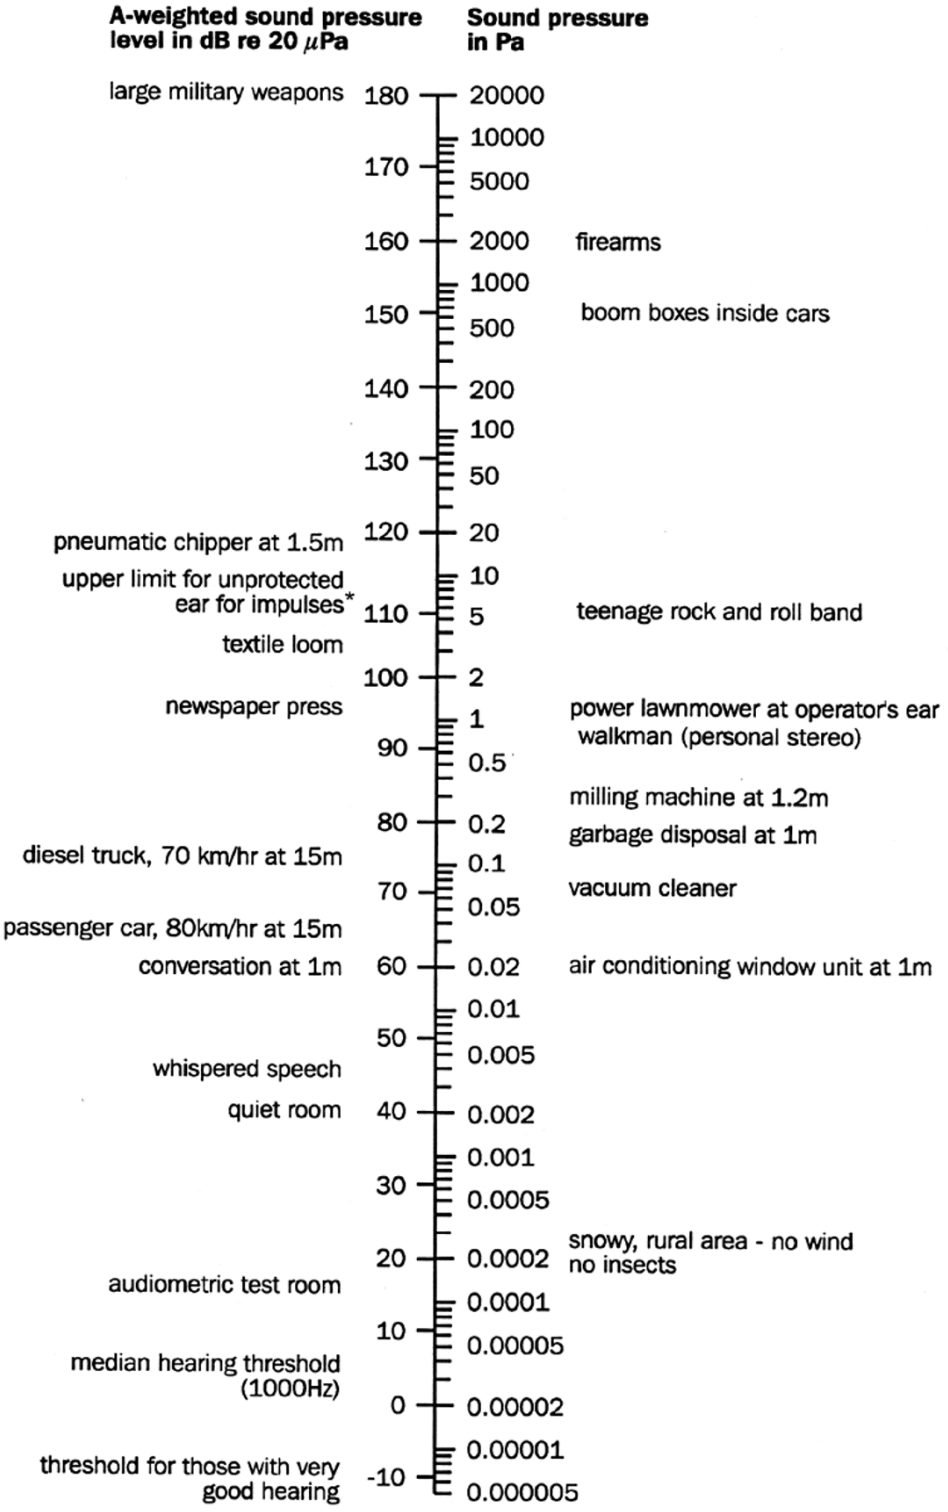

Figure 1.14: Sound levels produced by typical noise sources (Berenice Goelzer, Colin H. Hansen, 2001).

Chapter 2

Methods

2.1 Noise reduction strategies

The main approach to reducing acoustic noise exposure to the sleeping infants was to reduce the noise produced during acquisitions by minimising the spectral content of pulses applied to gradient coils, focusing specifically on the EPI readout portion of the DTI sequence. In addition to reducing produced noise, a passive noise reduction enclosure was built to surround the sleeping infant and act as an acoustic isolation box. Lastly, the introduction of noise from an external source to mask abrupt changes in scanner noise, was also investigated.

Experiments conducted for this project were performed on a 3 T Siemens Skyra scanner (Erlangen, Germany) according to protocols approved by the Faculty of Health Sciences Human Research Ethics Committee. The coordinate system for the imaging gradients X, Y, Z correspond to right-left, anterior-posterior, and superior-inferior directions, respectively.

2.1.1 Minimising gradient waveform spectral content

Two methods for minimising high frequency spectral content were implemented:

1. A method proposed by Schmitter et al. (2008) where the trapezoidal time optimized pulses in the EPI frequency encoding readout train are replaced with a sinusoidal wave and the set of phase encoding blips with one constant continuous gradient;

2. A new method using a second order Butterworth infinite impulse response (IIR) low-pass filter (LPF) to filter the gradient pulse waveforms, attenuating rather than eliminating the harmonics of traditional trapezoidal waveforms.

Filter coefficients were generated using Python's signal processing toolbox from SciPy.

To preserve pulse width, and thus time efficiency, pulses to be filtered were recreated as three identical pulses with alternating polarity, as shown in Figure 2.1. The alternating polarity waveform was then filtered both forward and backward in time to eliminate phase distortion. Finally, the middle portion of the pulse was used as

the filtered version of the original. Figure 2.1 shows the original and filtered pulse dimensions, and the alternating polarity three pulse arrangement.

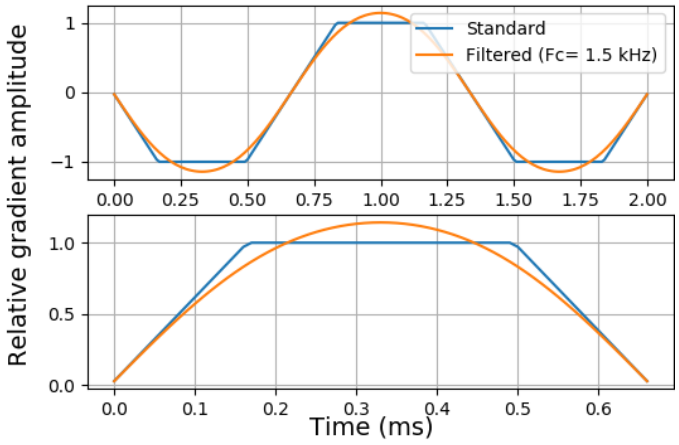

Figure 2.1: Top: Alternating polarity three pulse arrangement; Bottom: Middle portion of the waveform that is used as the filtered pulse.

Figure 2.2 shows pulse sequence simulations of the standard time optimized trapezoidal EPI readout waveforms, filtered trapezoidal waveforms using filters with fc = 1.5 kHz and 3 kHz respectively, as well as the sinusoidal waveforms. The frequency responses of the simulated EPI readout pulses in Figure 2.2 are shown in Figure 2.3. Table 2.1 gives the peak amplitudes for the resulting waveforms that maintain the original pulse moment as well as the maximum slew rate required to realise these waveforms.

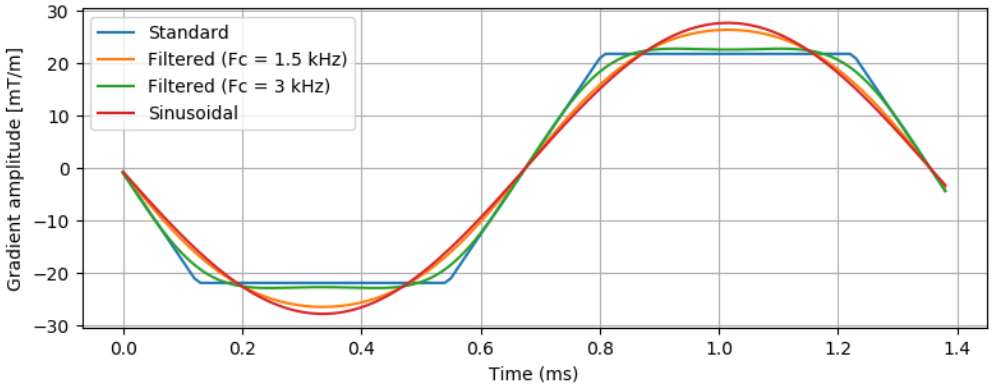

Figure 2.2: Simulated EPI frequency readout waveforms where the standard trapezoidal waveform (blue) was filtered using cut-off frequencies of 1.5 kHz (orange) and 3 kHz (green), respectively, along with a sinusoidal implementation (red). All four waveforms provide identical moments.

Pulse sequence modifications

In this project, we implemented both sinusoidal and digitally filtered frequency encoding waveforms in the a_ep2d_diff Siemens DTI pulse sequence.

In the case of the filtered implementation, the frequency encoding readout, as well as both the pre-winders for the frequency and phase encoding gradients were filtered

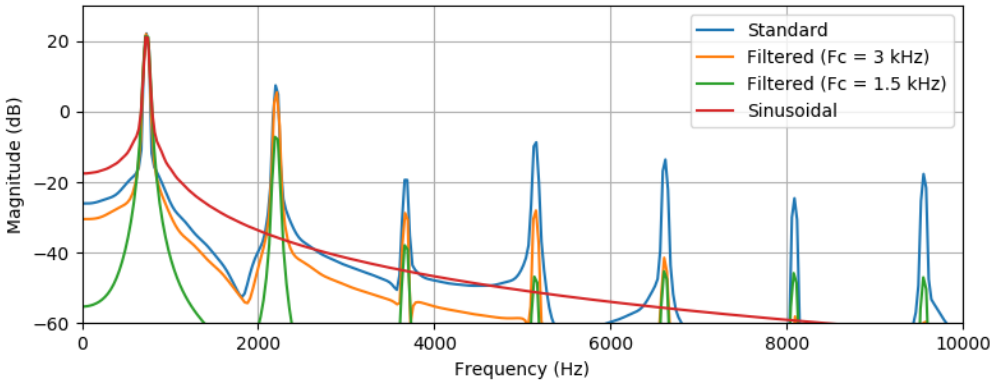

Figure 2.3: Frequency spectrum for each of the scanner simulated EPI readout waveforms shown in Figure 2.2.

Table 2.1: Simulated peak amplitude and slew rate requirements for different EPI frequency readout pulse shapes for an echo spacing of 0.68ms.

EPI RO Mode	Max Amplitude [mT/m]	Max slew rate [mT/m/ms]
Trapezoidal	21.48	165.25
Filtered (Fc=1.5kHz)	25.99	137.16
Filtered (Fc=3kHz)	22.45	172.37
Sinusoidal	27.28	126.08

using a second order Butterworth low-pass filter with frequency cut-off Fc = 1.5 kHz (top row of Figure 2.4). Fc was chosen to ensure attenuation of all harmonics of the frequency readout pulse train for echo spacings up to 1 ms. Instead of the constant continuous phase encoding gradient proposed by Schmitter et al. (2008), which causes substantial deviation from the standard k-space trajectory in the phase direction (Figure 2.6), the base time of the phase encoding blip was increased by 300 μs, as shown in the bottom plot of Figure 2.4. This causes considerably less phase deviation from the standard trajectory while significantly reducing the spectral content of the pulse compared to the standard blip. The frequency responses for the simulated waveforms in Figure 2.4 are shown in Figure 2.5.

For the sinusoidal implementation, the frequency encoding gradients were replaced by sinusoidal waveforms, the phase encoding blip with one constant continuous gradient pulse as proposed by Schmitter et al. (2008), and the frequency and phase encoding pre-winders by filtered waveforms using the same digital filter as in the filtered implementation.

For both the sinusoidal and filtered implementations, the crushers around the refocusing RF pulses, were also filtered using the same digital filter. Amplitudes of all filtered waveforms were adjusted to ensure the original gradient pulse moments were preserved.

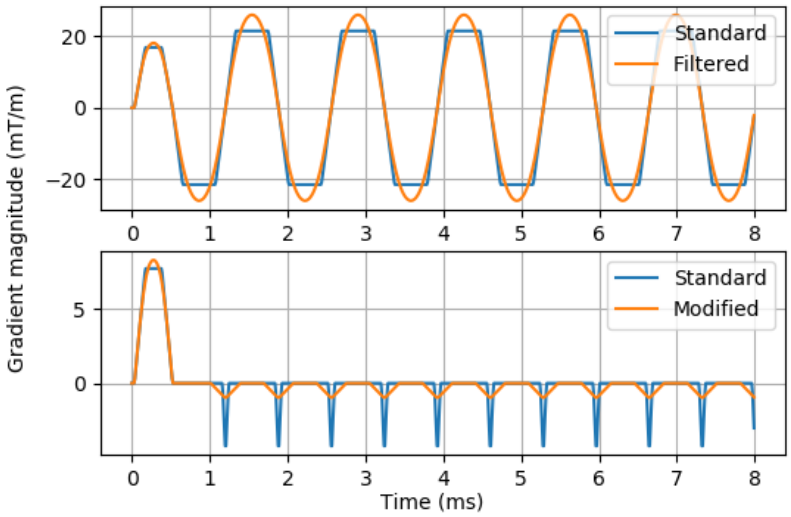

Figure 2.4: Simulated segment of the standard and filtered EPI frequency encoding (top) and phase encoding (bottom) readout. The low-pass filter cut-off frequency Fc was 1.5 kHz.

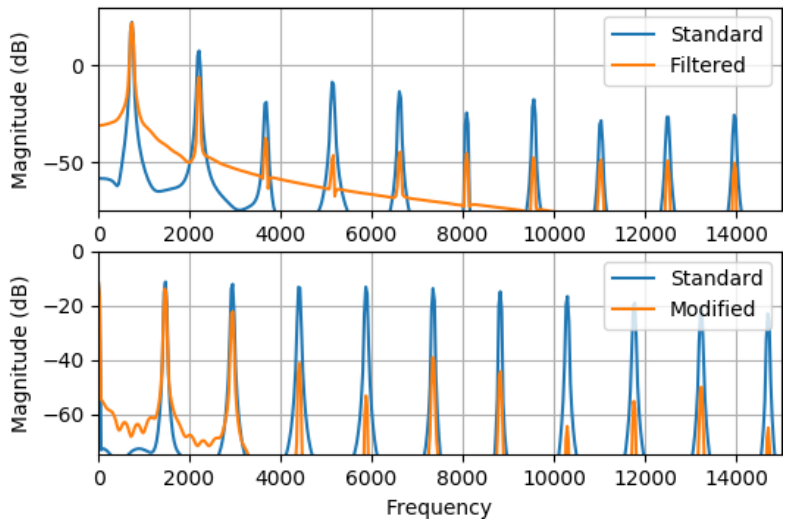

Figure 2.5: Frequency spectrum for the standard and filtered EPI frequency encoding waveforms (top), and standard and widened phase encoding blips (bottom) shown in Figure 2.4. The low-pass filter cut-off frequency Fc was 1.5 kHz.

K-space regridding

Since the modifications to the EPI readout pulse train alters the k-space trajectories, one dimensional regridding was applied to the raw data along the frequency encoding axis only. No regridding was performed in the phase encoding direction. Figure 2.6 shows the simulated partial k-space trajectories for EPI readouts using standard trapezoidal frequency encoding pulses with phase encoding blips, filtered frequency encoding pulses and widened phase encoding blips, and sinusoidal frequency encoding pulses with continuous constant phase encoding, respectively. Due to ramp sampling, the relative sampling density differs for the different waveforms as shown in the bottom plot of Figure 2.6.

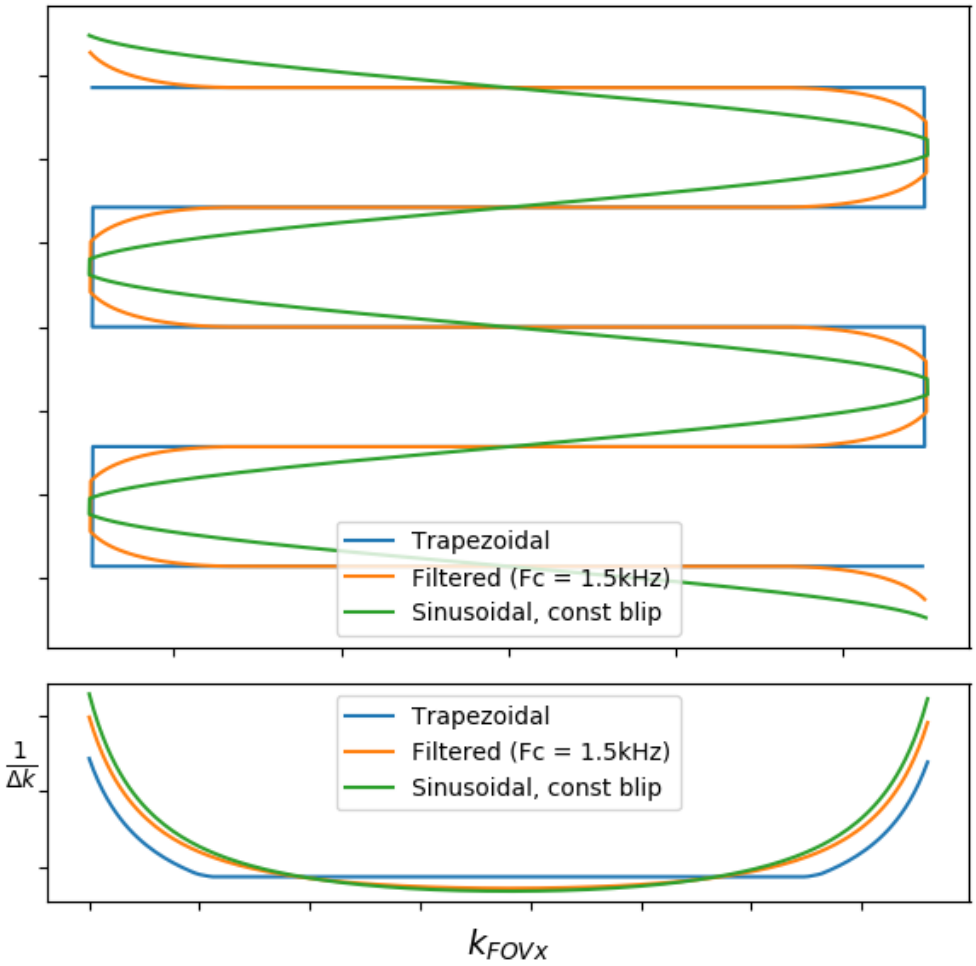

Figure 2.6: (Top) Simulated partial k-space trajectories for EPI readouts implemented using (i) standard trapezoidal frequency encoding pulses and phase encoding blips, (ii) filtered pulses and widened phase encode blips, and (iii) sinusoidal pulses and a continuous constant phase encoding gradient. (Bottom) Plots showing the differing relative sampling densities along the frequency encode direction arising from ramp sampling for each of the trajectories.

The raw data were acquired line by line and regridding was done using a Kaiser-Bessel window convolution function. The Kaiser-Bessel convolving kernel has superior image side lobe suppression compared to other common window functions like Hamming and Blackman windows (Jackson et al., 1991) while still being computationally efficient. The general regridding algorithm is given by

$$S^c(m\Delta k) = \Sigma_j S(k_j) g(m\Delta k - k_j) \Delta k_j^{(s)} \tag{2.1}$$

where $S^c(m\Delta k)$ is the convolved data, resampled at equidistant points $m\Delta k$ and m is an integer indicating the m^{th} data point. $S(k_j)$ represents the non-uniformly sampled data.

The convolution kernel $g(m\Delta k - k_j)$ in Equation 2.1 is given by

$$g(k) = \frac{1}{W} \times I_o \left(\beta \left(1 - \frac{2k}{W} \right)^2 \right) \times RECT \left(\frac{2k}{W} \right) \tag{2.2}$$

in which

$$RECT(u) = \begin{cases} 1 & \text{if } |u| < 0.5 \\ 0 & \text{otherwise} \end{cases} \tag{2.3}$$

$k = m\Delta k - k_j$ and W is the chosen kernel window size in units of rectilinear k-space samples. Although larger window sizes increase side lobe suppression in the center of the image FOV (Jackson et al., 1991)], more multiplications are required.

The number of multiplications required for the convolution in 1D regridding in Equation 2.1 is given by WN_s where N_s is the number of input data samples and can become substantial for large window sizes. This can cause delays in image reconstruction time. The values for β in Equation 2.2 are given in Table 2.2 (Jackson et al., 1991).

Table 2.2: Parameter values for the convolution function given in Equation 2.2 (Jackson et al., 1991).

Window width (W)	β
2	2.3934
4	5.7567

In Equation 2.2, I_o is the zero-order modified Bessel function of the first kind and was implemented using the approximation from Salahat et al. (2013) given in Equation 2.4 where α_i and β_i are real values given in Table 2.3. For all values of z in Equation 2.4 the maximum absolute error is of the order 10^{-4} (Salahat et al., 2013).

$$I_0(z) \approx \sum_{i=1}^{4} \alpha_i e^{\beta_i z} \tag{2.4}$$

Table 2.3: Parameter values for the zero-order modified Bessel function approximation given in Equation 2.4.

$0 \leq z \leq 11.5$		$11.5 < z \leq 20$		$20 < z \leq 37.25$		$37.25 < z$	
i	α_i	β_i	α_i	β_i	α_i	β_i	α_i	β_i
1	0.1682	0.7536	0.2667	0.4710	0.1121	0.9807	2.41e-9	1.144
2	0.1472	0.9739	0.4916	-163.4	0.1055	0.8672	0.06745	0.995
3	0.445	-0..715	0.111	0.9852	-0.00018	1.0795	0.05471	0.5686
4	0.2382	0.2343	0.1304	0.8554	0.00326	1.0385	0.07869	0.946

Density compensation was performed by weighting each value based on it's proximity to adjacent samples according to

$$\Delta k_j^{(s)} \approx |k_j - k_{j-1}|. \tag{2.5}$$

This was necessary to compensate for the increasing sample density closer to the edges of k-space in the frequency encoding direction arising from the altered frequency encoding waveforms.

2.1.2 Frequency response function (FRF)

We measured the FRF for each gradient coil, respectively, by recording the acoustic noise resulting from 10 different, uniformly distributed random noise gradient waveforms, band-limited to 6 kHz (Wu et al., 2014). In this way noise from the switching of the frequency encoding gradient was limited by ensuring that the EPI fundamental frequency ($\frac{1}{2*echo\,spacing}$) coincides with a trough in the FRF. Figure 2.7 shows an example of such a pulse and Figure 2.8 shows the measured FRFs for each gradient.

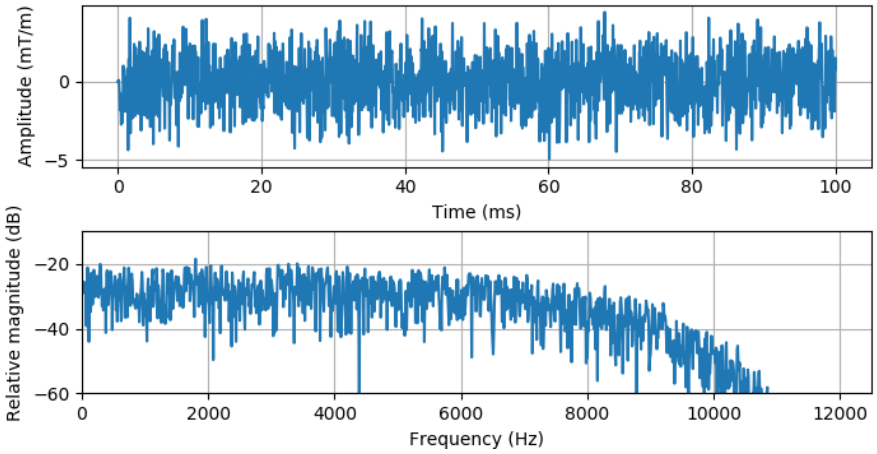

Figure 2.7: An example of a uniformly distributed random noise gradient waveform, band-limited to 6 kHz.

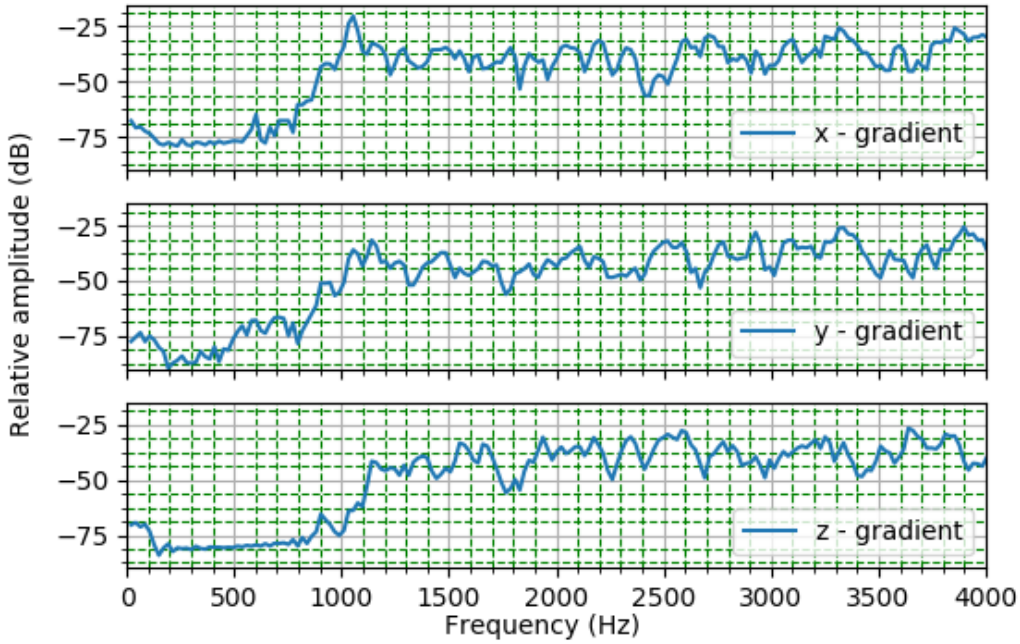

Figure 2.8: Measured frequency response function for X,Y, and Z gradient coil assemblies, respectively.

Acquisition protocol

Both phantom and *in vivo* scans were conducted using each of the sequences (standard, filtered, sinusoidal) to compare acoustic noise and image quality. For phantom scans a shorter protocol was used. Table 2.4 gives the DTI protocols used for phantom (short) and *in vivo* (long) acquisitions, respectively.

Table 2.4: Acquisition protocols used to measure noise and compare image quality in each of the sequences. The short protocol was used for phantom scans and the long protocol to scan a healthy adult volunteer. Bold highlights parameters that were changed in the long protocol.

Parameter	Short protocol	Long protocol	Unit
Number of slices	30	**84**	
TR	4700	**13000**	ms
TE	80/81*	80	ms
FOV (read)	244	244	mm
Base resolution	120	120	Px
GRAPPA acceleration factor	2	2	
Fat suppression	None / Strong	**Strong**	
Diffusion directions	6	**30**	
b-values	0	**0 (5 averages)**	
	1000	1000	
Bandwidth	1665 / 1544 *	1665	Hz/Px
TA	0:45	**8:14**	minutes

* For the filtered and sinusoidal acquisitions with the short protocol, the echo spacing was increased to 0.73 ms to coincide with a trough in the FRF, changing TE to 81 ms and bandwidth to 1544 Hz/Px.

The EPI echo spacing of 0.68 ms for the standard protocol wasn't optimal with regards acoustic noise due to the location of the fundamental frequency and the first two harmonics on the frequency response functions of the gradient coils shown in Figure 2.8. Since an echo spacing of 0.73 ms would produce less noise while not requiring a larger TR, this value was used for the filtered and sinusoidal phantom acquisitions. This required a different bandwidth to the standard acquisition, namely 1544 Hz/Px, and resulted in a tiny (1.1 ms) increase in TE. The acquisition time was identical for all three sequences.

2.1.3 Acoustic insulation enclosure

PROCELL® open cell polyurethane foam was used as the insulation material for constructing the sides of the enclosure. A simple frame was built using 40 mm PVC pipes to hold the foam in place. A 1 x 1 m foam plate, 50 mm thick was then simply bent to form a semi-circle and placed inside the frame to form the enclosure. The distal end was closed with an extra piece of foam while the feet-end was left open to allow the radiographer to see the infant during acquisitions. The prototype used for the acoustic measurements is shown in Figure 2.9. The microphone was placed on a water phantom inside the infant head coil, close to where the infants ear would be situated, facing to the left.

Once the sleeping infant has been successfully placed in the infant head coil on the scanner bed, the foam enclosure can simply be placed over the sleeping infant after

which scanning can commence as normal.

Figure 2.9: Prototype acoustic isolation box used for testing possible noise reduction to sleeping infants during acquisitions.

2.1.4 Masking scanner noise

To improve noise continuity between acquisitions and reduce the startle effect to the infant at the start of acquisitions, pink noise was played in the scanner room via the scanner speakers for the duration of the scan session. Scanner noise was recorded with and without the added pink noise to show the effectiveness of the proposed method to reduce acoustic contrast.

2.2 Acoustic noise recording

Scanner noise was measured with an Optimic 1155 optical microphone from Optoacoustics, attached to the top of a water phantom inside a 16 channel infant head coil in the isocenter along the Z-axis, facing the bore in the right-left direction. The technical specifications for the optical microphone are given in Table 2.5.

Table 2.5: Relevant specifications for Optimic 1155 optical microphone.

Polar pattern	Omidirectional
Frequency Response	10-15000 Hz
Conformity 10-10000Hz	$\pm$ 1 dB
Conformity 10000-15000Hz	$\pm$ 3 dB
Sensitivity	15mV/Pa $\pm$ 10% at 1kHz
Equivalent Self-Noise	$\leq$ 45 dBA SPL
Total Harmonic Distortion	$\leq$ 1% at 100dB SPL
Maximum Acoustic Pressure	130 dB SPL

2.2.1 Audio recording calibration

The analogue output of the microphone was recorded using the line input from a Dell Latitude E5570 laptop which digitized the analogue data using the on-board sound card analogue-to-digital converter (ADC). The ADC converts the continuous time, continuous amplitude analogue signal to a digital signal by quantizing the

amplitude of the signal into one of 2^B levels at discrete time intervals (T_S), where B is the number of bits used to represent the amplitude of a sample of the ADC. In this case each sample amplitude was represented by a signed 16-bit integer value, one bit to represent polarity and 15 bits for representing an integer number between 0 and $2^{15} - 1$. The relationship between the ADC resolution B, the measured analogue voltage ($V_{Analogue}$), the internal ADC reference voltage (V_{Ref}), and the quantised ADC output is given by Equation 2.6.

$$\frac{ADC\ resolution}{V_{Ref}} = \frac{ADC\ output}{V_{Analogue}} \tag{2.6}$$

To determine the internal ADC reference voltage, the peak amplitude of an analogue sinusoidal waveform was measured using an oscilloscope and the laptop ADC respectively. The analogue waveforms were generated using a Xiaomi Redmi 3s smart phone and the Simple Tone Generator mobile application. The corresponding measurements are shown in Table 2.6.

Table 2.6: Table showing measured maximum amplitude analogue and digitized values of a analogue sinusoidal waveform of various amplitudes as well as the calculated internal reference voltage of the signed 16-bit laptop ADC.

Oscilloscope (mV)	ADC output	Calculated Vref (mV)
600	10440	1883
432	7407	1911
300	5264	1867
228	3924	1904
160	2784	1883
114	1978	1888
84	1474	1867

Using Equation 2.6 with the ADC resolution B = 15, the mean reference voltage between the measurements was $V_{ref} = 1.878$ V $\pm$ 1.14 mV. The experimental setup is depicted in Figure 2.10.

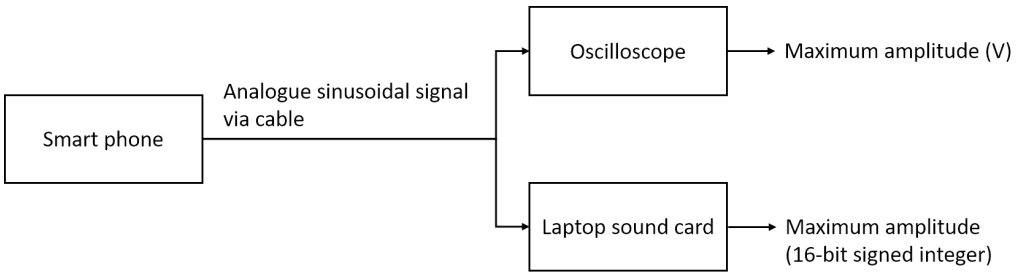

Figure 2.10: Experimental setup for acquiring the data in Table 2.6 used to calculate the internal reference voltage of the laptop sound card ADC.

2.2.2 Raw audio processing

The raw audio was recorded in Python and exported to Audacity 2.1.3 which was used to crop the appropriate sections of the total noise recording, taking only the

noise produced from the appropriate acquisition and to apply an A-weighting filter to the raw audio.

Given the microphone sensitivity of 15.5 mV at 1 Pa, the sound pressure waveform measured by the microphone is determined from the recorded, A-weighted audio voltage signal V_{Audio} using Equation 2.7. Figure 2.11 shows a flow diagram illustrating the entire audio recording process.

$$p(t) = \frac{V_{Audio}(t)}{Microphone\ sensitivity} \tag{2.7}$$

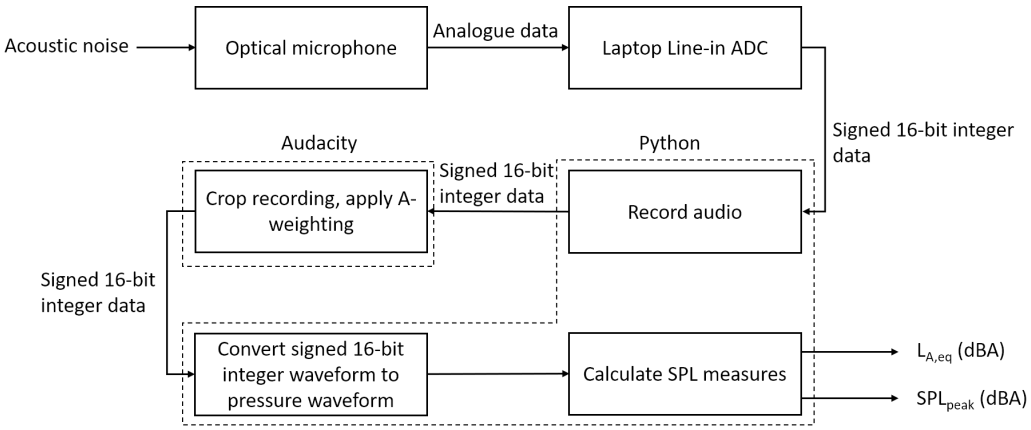

Figure 2.11: Flow diagram depicting the audio recording process.

2.3 Data analysis

2.3.1 Audio

Using the A-weighted pressure signal $p(t)$ in Pa, the peak sound pressure level (SPL), measured in dBA, was determined using

$$L_A = 20log_{10}(p) + 94, \tag{2.8}$$

in which p is the maximum absolute value of $p(t)$.

However, because of the fluctuating, impulse type nature of the noise, the peak SPL can be misleading in terms of the total noise heard by patients. To account for this the A-weighted equivalent continuous sound level $L_{A,eq}$ was used as a better measure of loudness; the steady state $L_{A,eq}$ value over a given time period has the same energy as the original audio signal $p(t)$ (Berenice Goelzer, Colin H. Hansen, 2001). $L_{A,eq}$ was computed using Equation 2.9 where p_0 is the threshold of human hearing at $20\,\mu Pa$, T is the time period (taken as the sequence acquisition time) and $p(t)$ is the measured time-varying noise pressure signal.

$$L_{Aeq,T} = 10log_{10}\left(\frac{1}{T}\int_0^T \left(\frac{p(t)}{p_0}\right)^2 dt\right) \tag{2.9}$$

2.3.2 Image Analysis

We compared image quality of the modified DTI sequences to that of the standard sequence by evaluating SNR and FA measurements in the corpus callosum (CC). Image analysis was performed using an open source Python based software project developed for computational neuroanatomy with a main focus on dMRI called DIPY (Diffusion Imaging In Python) (Garyfallidis et al., 2014). Using DIPY, the FA and coloured FA (CFA) maps as well as SNR measurements were calculated from acquired diffusion weighted images.

Six diffusion datasets were acquired from a healthy volunteer – two standard acquisitions, two acquisitions with filtered EPI readouts and regridding kernel window sizes of 2 and 4, respectively, and two acquisitions with sinusoidal EPI readouts and regridding kernel window sizes of 2 and 4, respectively. The first standard DTI acquisition dataset was used as the reference to which all subsequent acquisitions were registered before any analysis was performed. Registration was done using mri_robust_register (Reuter, Rosas, and Fischl, 2010).

Signal-to-noise ratio (SNR)

SNR was defined as the ratio of the signal mean in an ROI divided by the standard deviation of the background noise as shown in Equation 2.10.

$$SNR = \frac{mean\ signal\ in\ ROI}{Standard\ deviation\ of\ background\ noise} \tag{2.10}$$

For each acquisition, a mask covering the CC was generated by isolating the red voxels in the CFA maps, indicating high diffusivity in the left-right direction. This was implemented by setting thresholds for red, blue, and green as (0.6,1), (0,0.3), and (0,0.3), respectively, where red, blue, and green are indicative of diffusivity in the left-right, superior-inferior, and anterior-posterior directions, respectively.

Corpus callosum SNR was then calculated for each of the 30 diffusion-weighted images, and the b0 image, and the lowest, highest, and b0 values compared for the different acquisitions. The lowest SNR (worst-case) corresponds to diffusion weighting along (or nearly along) the CC tracts, and the highest (best-case) to diffusion weighting perpendicular (or nearly perpendicular) to the tracts. Processing was done using DIPY (Garyfallidis et al., 2014).

Fractional anisotropy (FA)

For each acquisition, a binary mask was generated of the corpus callosum in the sagittal view by isolating the red voxels in the coloured FA maps using a threshold of 0.6, similar to the mask generated for the SNR comparisons. The individual binary masks of all acquisitions were then multiplied with each other to produce a final mask containing only corresponding pixels that were above the threshold in all acquisitions. This final mask was applied to the FA maps from each acquisition to extract voxelwise FA values across the corpus callosum.

Paired two-tailed t-tests were performed to determine whether mean FA across the corpus callosum differed between the modified and standard acquisitions. The t-tests were implemented using Python's SciPy.stats package.

Chapter 3

Results

3.1 Noise reduction

The acoustic noise levels of the different acquisitions are shown in Table 3.1 both without and with the foam enclosure. Since fat saturation influenced the total recorded noise levels, we report results for both strong fat suppression and no fat suppression.

While simply increasing the EPI echo spacing of the standard DTI sequence reduces the peak SPL by a maximal 3.9 dBA for strong fat suppression without the foam enclosure, the lowest A-weighted equivalent continuous sound levels ($L_{A,eq}$) are recorded for the filtered and sinusoidal acquisitions, with the sinusoidal sequence performing similarly or marginally (no more than 0.5%) worse than the filtered on all outcomes. Without the foam enclosure, the filtered acquisition with filtered crushers reduces peak SPL by 3.4 and 4 dBA for strong and no fat suppression, respectively, and $L_{A,eq}$ by 2.5 and 2.8 dBA, respectively. Adding the foam enclosure increases the amount by which peak SPL is reduced to 4.8 and 7 dBA, respectively.

Table 3.1: Acoustic noise measurements for standard and modified DTI acquisitions with strong / no fat saturation, respectively. (Trap = Trapezoidal; Filt = Filtered; Sin = Sinusoidal; Const = Constant)

Without foam enclosure						
Sequence	**Frequency (Hz)**[1]	**Read-out**	**Phase**	**Crushers**	**Peak SPL (dBA)**	**$L_{A,eq}$ (dBA)**
					Fat suppression Strong / No	**Fat suppression Strong / No**
Standard	735	Trap	Trap blip	Trap	114.7 / 114.4	97.9 / 97.5
Standard	684	Trap	Trap blip	Trap	110.8 / 110.6	96.9 / 96.4
Filtered	684	Filt	Wider blip	Trap	111.4 / 110.6	95.5 / 94.8
Filtered	684	Filt	Wider blip	Filt	111.3 / 110.4	95.4 / 94.7
Sinusoidal	684	Sin	Const	Filt	111.8 / 111.0	95.5 / 94.7
With foam enclosure						
Standard	735	Trap	Trap blip	Trap	111.6 / 111.2	95.2 / 95.0
Filtered	684	Filt	Wider blip	Filt	109.9 / 107.4	92.7 / 92.0
Sinusoidal	684	Sin	Const	Filt	110.5 / 108.0	92.7 / 92.0

[1]EPI fundamental frequency given by $\frac{1}{2*echo\ spacing}$; 735 Hz corresponds to an echo spacing of 0.68 ms; 684 Hz corresponds to an echo spacing of 0.73 ms.

Figure 3.1 shows the effect of the echo spacing increase on the noise spectrum of the standard acquisition, and Figure 3.2 the effect of sequence modifications, both in the absence of the foam enclosure. Notably, the filtered and sinusoidal sequences demonstrate a significant reduction in EPI harmonics (Figure 3.2) compared to the standard sequence, but very little difference between each other. Figure 3.3 shows the combined effect of the filtered EPI readout and the foam enclosure on the noise spectrum, as well as the effect of additionally eliminating fat suppression.

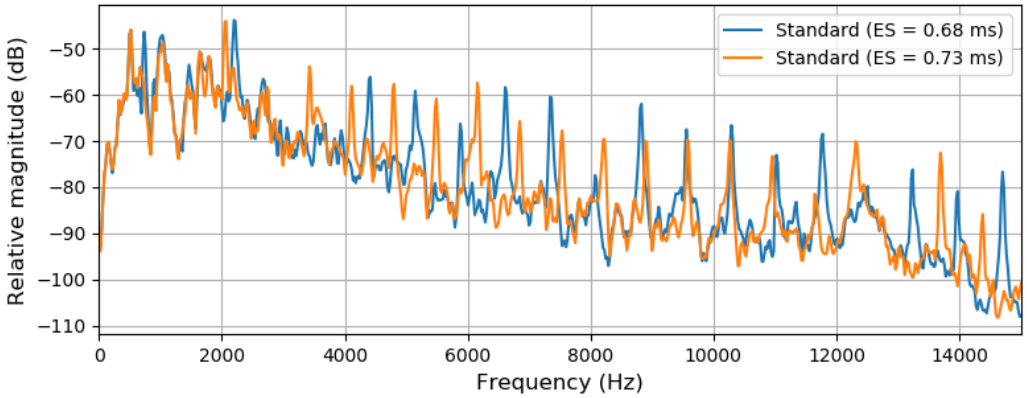

Figure 3.1: Noise spectra from standard DTI acquisitions with EPI echo spacings of 0.68 ms (blue) and 0.73 ms (orange).

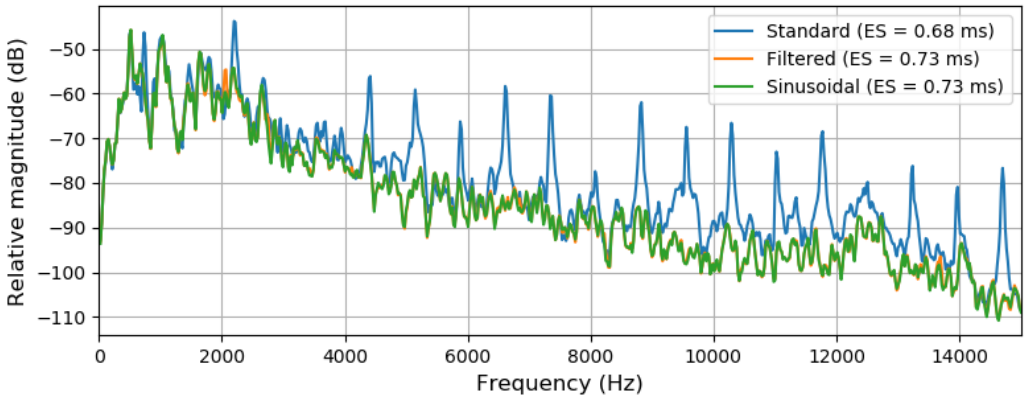

Figure 3.2: Comparison of noise spectra from the standard DTI sequence with echo spacing 0.68 ms (blue) to the filtered (orange) and sinusoidal (green) sequences, both with echo spacing 0.73 ms.

Figure 3.4 demonstrates how introducing pink noise into the room reduces changes in sound pressure level during the scanning session, which could potentially reduce a baby's startle response both at the start of and during an acquisition.

Figure 3.5 shows a near 40% reduction in acoustic noise amplitude compared to the standard DTI sequence for the filtered sequence with echo spacing 0.73 ms and foam enclosure, both with and without fat suppression.

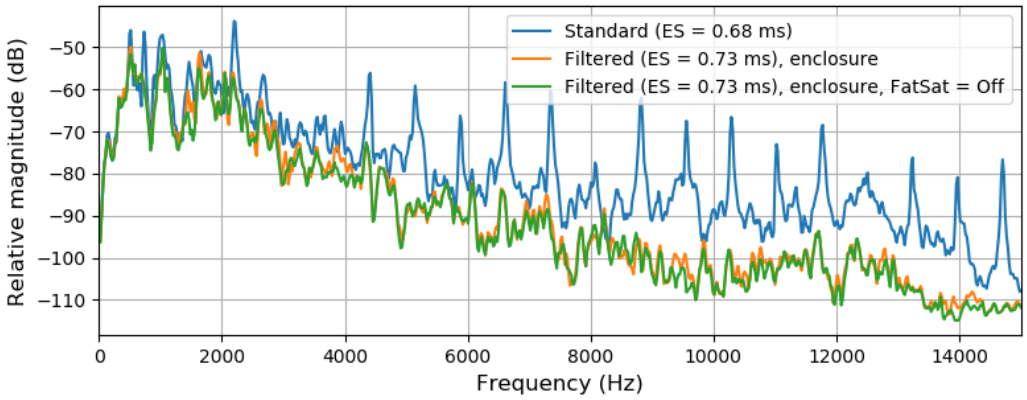

Figure 3.3: Comparison of the noise spectrum from a standard DTI acquisition with EPI echo spacing 0.68 ms, strong fat suppression and no foam enclosure (blue) to the noise spectra from the filtered sequence with longer echo spacing, a foam enclosure, and either strong (orange) or no (green) fat suppression.

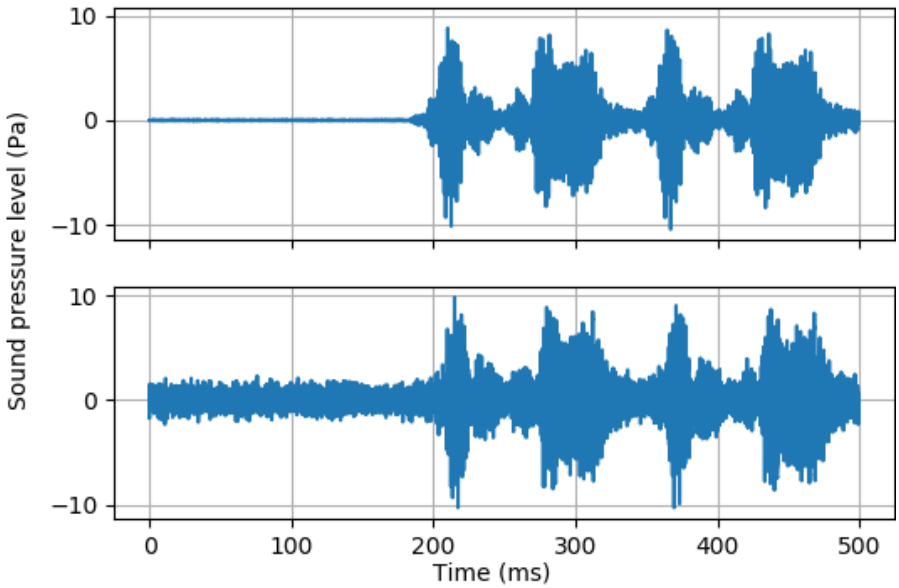

Figure 3.4: Recorded noise pressure levels without (top) and with (bottom) pink noise played through the scanner speaker.

3.2 Image analysis

Single slices from the non-diffusion weighted (b0) volumes of 6 *in vivo* acquisitions are shown in Figure 3.6. Subfigures A and D show images acquired using the standard sequence, B and E using a sinusoidal EPI readout with constant phase blip, and C and F using the filtered sequence. For B and C, regridding was performed using a kernel width (W) of 2, and for E and F, a kernel width of 4 was used.

Comparing SNR in the corpus callosum (CC) for three volumes – b0, the diffusion weighted (DW) volume with the lowest CC SNR, and the DW volume with the highest CC SNR – yielded values 9 to 44% higher for the filtered acquisition compared to the standard acquisition (Table 3.2), while SNR values for the sinusoidal acquisition

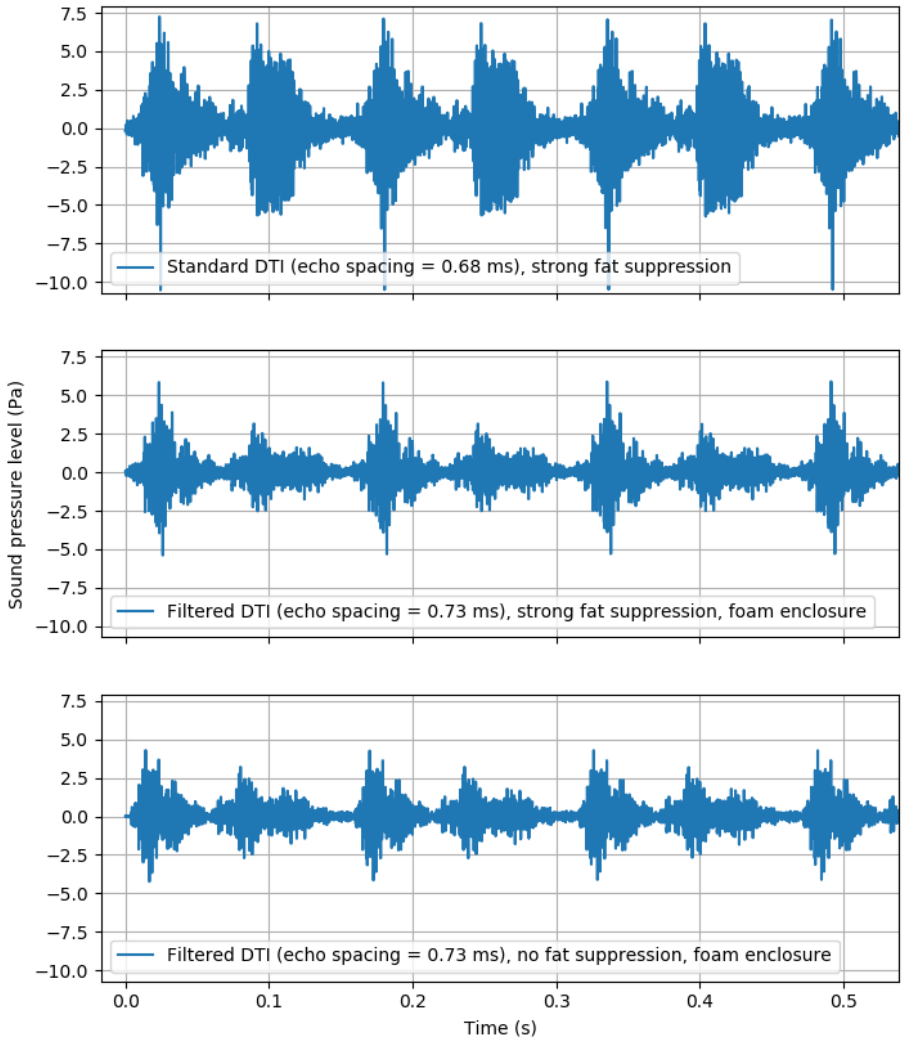

Figure 3.5: (Top) Noise recorded during a standard DTI acquisition (EPI echo spacing 0.68 ms) with strong fat suppression. (Middle and Bottom) Noise recorded during filtered acquisitions (EPI echo spacing 0.73 ms) in the foam enclosure, with (middle) and without (bottom) fat suppression, respectively.

were similar or lower than the standard. Notably, the observed SNR increase of the filtered acquisition was not attributable to lower standard deviation (STD) of the background noise, as this was greater for both modified sequences compared to the standard sequence. The STD of the background noise increased also with larger regridding kernel windows.

Figure 3.7 shows fractional anisotropy (FA) maps constructed from the standard, sinusoidal, and filtered acquisitions, respectively, as well as zoomed in views of the boxed region. Figure 3.8 shows the % difference in FA values of the sinusoidal and filtered acquisitions, respectively, relative to the standard. The absence of features, except for the skull, demonstrates similar performance of the modified sequences across all tissue types and different anatomical regions.

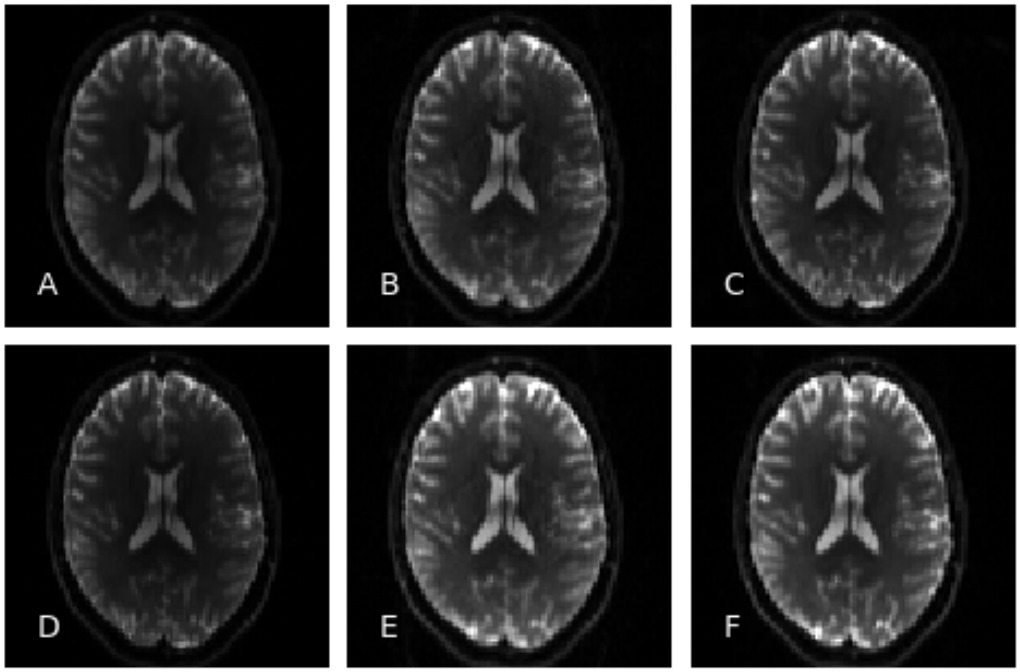

Figure 3.6: Non-diffusion weighted (b0) images from a healthy adult volunteer acquired using standard (A and D), sinusoidal (B and E), and filtered (C and F) DTI sequences, respectively. The regridding kernel window width was 2 for B and C, and 4 for E and F.

Table 3.2: Comparison of signal-to-noise ratio (SNR) in the corpus callosum (CC) for three volumes (b0, greatest CC signal attenuation, least CC signal attenuation) from each acquisition.

EPI readout	Window size (W)	SNR (b0)	SNR (min)	SNR (max)	Background noise STD
Standard 1	-	18.0	4.4	12.5	6.3
Standard 2	-	19.3	4.6	13.2	5.9
Filtered	2	21.4	5.2	14.4	7.5
Filtered	4	25.9	6.3	17.6	9.0
Sinusoidal	2	17.7	4.5	12.1	9.4
Sinusoidal	4	17.9	4.4	12.4	12.9

STD = Standard deviation

Figure 3.9 shows the mid-sagittal slice of the FA map for each of the acquisitions, and Figure 3.10 the corpus callosum mask within which voxelwise FA values were extracted for comparison between the different acquisitions (Table 3.3). FA values in the corpus callosum of the sinusoidal and filtered acquisitions did not differ from those of the standard acquisitions (pairwise student's t-test, all p's >0.2).

The increases in image reconstruction time required for the modified acquisitions using regridding kernel widths of W = 2 and 4 were 45 seconds and 7:45 minutes, respectively, for the 16 channel infant head coil.

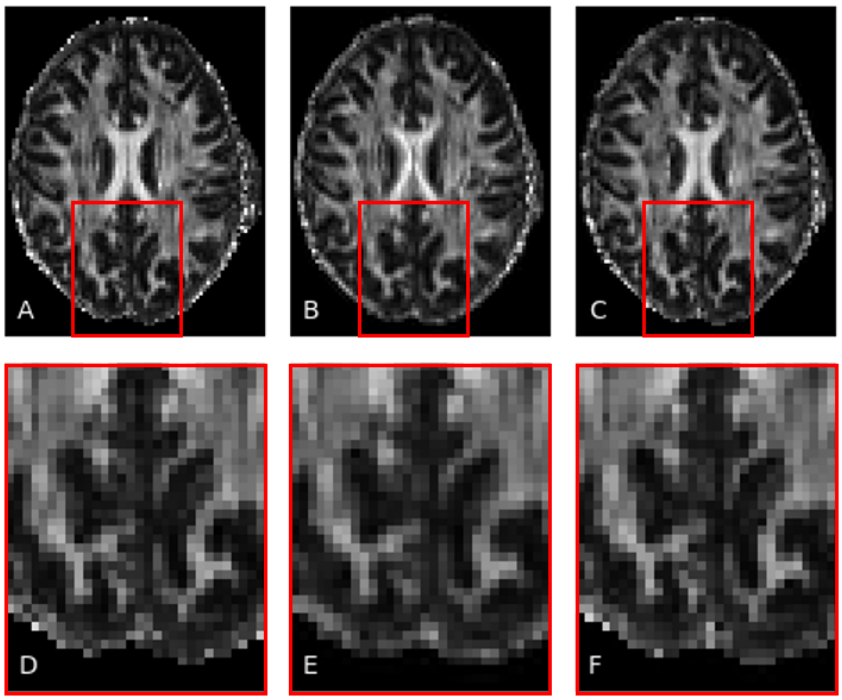

Figure 3.7: FA maps from standard (A), sinusoidal (B), and filtered (C) acquisitions, and zoomed views of the highlighted regions for each.

Table 3.3: Comparison of fractional anisotropies averaged across the corpus callosum mask for each of the different DTI acquisitions.

Acquisition	Window size (W)	Fractional anisotropy (mean ± STD)
Standard 1	-	0.80 ± 0.07
Standard 2	-	0.80 ± 0.07
Filtered EPI	2	0.80 ± 0.06
	4	0.80 ± 0.07
Sinusoidal EPI	2	0.79 ± 0.08
	4	0.79 ± 0.07

STD = Standard deviation

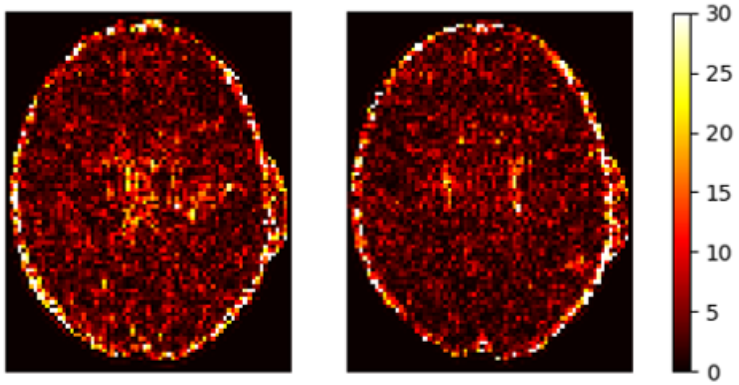

Figure 3.8: Voxelwise percentage difference in FA values of sinusoidal (left) and filtered (right) acquisitions relative to the standard acquisition.

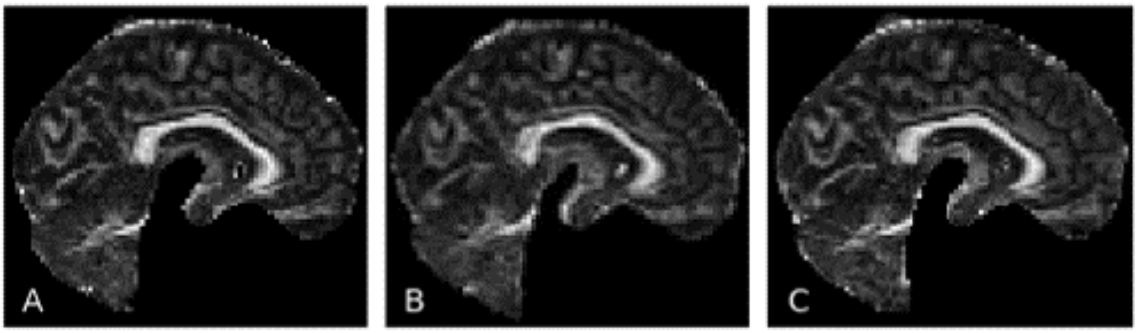

Figure 3.9: Mid-sagittal slices of the Fractional Anisotropy maps constructed from standard (A), sinusoidal (B), and filtered (C) acquisitions in a healthy adult volunteer.

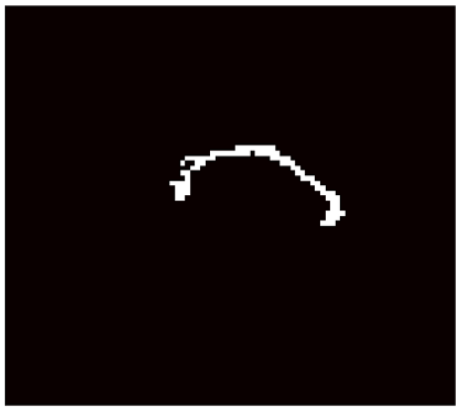

Figure 3.10: Binary corpus callosum mask within which voxelwise fractional anisotropy values were extracted for comparison between the different acquisitions.

Chapter 4

Discussion

4.1 Noise reduction

In this study we reduced acoustic noise generated during DTI acquisitions by reducing the high frequency content in the EPI readout gradient waveforms using digital filters, adjusting the EPI echo spacing to minimize noise from the fundamental frequency, and adding an acoustic foam enclosure over the sleeping infants for passive noise absorption.

Filtered EPI readout waveforms showed similar reductions in acoustic noise to the sinusoidal implementation with a constant phase encoding blip which suggests that the harmonics of the gradient waveforms need not be removed entirely but only require attenuation. This allows for pulse shapes that more closely resemble the trapezoidal standard, resulting in less deviation from the standard k-space trajectory, possibly reducing error inherent in the regridding process. Another advantage of filtered pulses is that they require lower peak amplitudes than sinusoidal waveforms to produce equivalent pulse moments, with the difference in amplitude depending on the filter cut-off frequency (Fc) and EPI echo spacing.

The acoustic noise spectra of the DTI sequences with filtered and sinusoidal EPI readouts demonstrate a significant reduction in EPI harmonics (Figure 3.2) compared to the standard sequence, but very little difference between each other. This not only leads to lower peak SPL and $L_{A,eq}$, but also lowers the tone of the remaining noise, possibly reducing patient annoyance and anxiety related to the high pitched dissonant tone of the standard EPI sequence.

While peak SPL values with and without fat suppression for the standard acquisition differed by only 0.2 - 0.4 dBA (<0.4%), removing fat suppression in the modified sequences reduced peak SPL values by 0.8 - 0.9 dBA (0.7 - 0.9%) compared to their fat-suppressed counterparts. This effect was exaggerated in the foam enclosure where peak SPL values were lower by 2.4 - 2.5 dBA (2.3%) when fat suppression was turned off.

The noise reduction achieved in this work (2.9 dBA peak SPL; 2.4 dBA $L_{A,eq}$) for a sinusoidal EPI readout with constant phase blip is significantly lower than the 22.5 dBA average SPL reported by Schmitter et al. (2008). It should be noted, however,

that the results were reported for different sequences and although both DTI and fMRI sequences contain an EPI readout, the sequences are notably different, making noise reduction results due to the EPI readout modifications difficult to compare. Furthermore, in addition to the sinusoidal EPI readout and constant phase blip, Schmitter et al. (2008) replaced the readout pre-winder with a sinusoidal waveform of linearly increasing amplitude and total moment equal to that of the original pre-winder and added a sinusoidal wave with linearly decreasing amplitude to the end of the EPI readout to prevent the abrupt ending of the waveform. These modifications could have contributed to increased noise reduction.

The 5.8 dBA SPL noise reduction achieved by Hutter et al. (2018) for a DTI sequence with sinusoidal EPI readout, a constant phase encoding blip, and sinusoidal CAIPIRINHA blips is more comparable to the results reported here (2.9 dBA peak SPL and 2.4 dBA $L_{A,eq}$). Differences may be due to the use of differing metrics, since Hutter et al. (2018) reported noise reduction in dBA SPL, but did not indicate whether values referred to peak, average, root-mean-squared (RMS), or $L_{A,eq}$ SPL. Notably, the present work also did not include CAIPIRINHA blips.

Both Schmitter et al. (2008) and Hutter et al. (2018) also used lower fundamental frequencies for the EPI readout (543 Hz and 507 Hz, respectively) than the present work. This would affect noise reduction as the A-weighting scale attenuates frequencies on the lower end of the spectrum more and since the scanner transfer functions reported by Schmitter et al. (2008), Hutter et al. (2018)) as well as the current work, show lower amplification in the 500 Hz range. In this work we chose to use a fundamental frequency of 684 Hz to keep the scan time the same as for the standard EPI fundamental frequency of 735 Hz.

The position of the microphone between the experiments conducted by Schmitter et al. (2008), Hutter et al. (2018), and this work also varied. In this work, the microphone was placed on a water phantom inside the 16 channel infant head coil at the isocenter along the Z-axis, close to where the infant's ear would be situated, and facing to the side of the scanner. In contrast, it was placed in the isocenter in an empty bore by Hutter et al. (2018) and on top of a water phantom in a 12 channel head coil by Schmitter et al. (2008). The orientation of the microphone was not specified in these works.

A major factor that impedes proper comparison between results reported in publications is the lack of standardisation in the type of measurement used to quantify the level of acoustic noise. Often the only noise reduction measure reported is SPL which, assuming this refers to peak SPL, on its own does not provide a clear picture of total noise reduction, especially for the impulse-like noise from an MRI acquisition.

This is illustrated in Figure 4.1 for two noise pressure waveforms that differ in peak SPL by 2.3 dBA and $L_{A,eq}$ by 5.1 dBA. The single repetitive peak in the second waveform, which is substantially higher than the rest of the signal, causes a smaller total reduction in peak SPL, leading to an underestimate of noise reduction compared to the RMS-type noise reduction measure $L_{A,eq}$, indicative of the noise energy over time, which is substantially reduced.

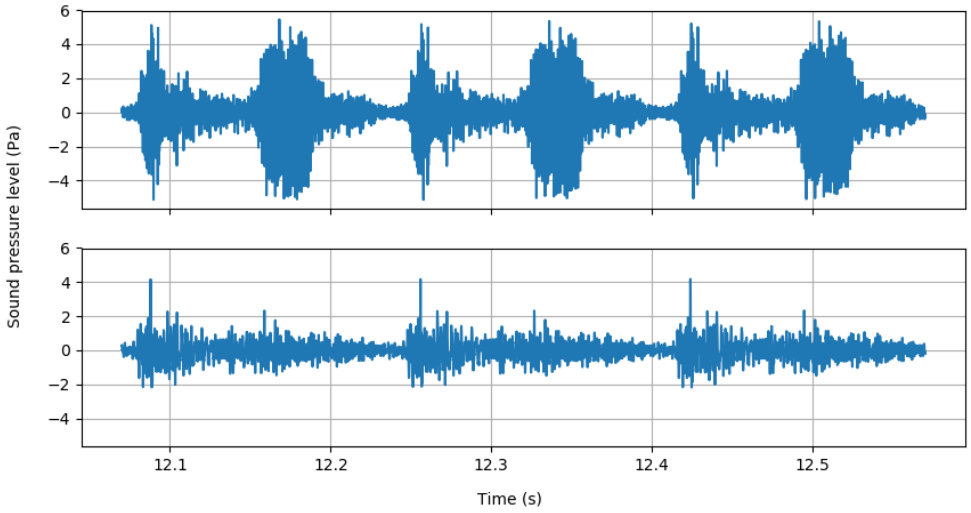

Figure 4.1: Two noise pressure waveforms that differ in peak SPL by only 2.3 dBA, but in $L_{A,eq}$ by 5.1 dBA.

Supplementing peak SPL measurements with continuous equivalent noise $L_{A,eq}$ measurements will provide a more comprehensive picture of the actual noise reduction.

The placement of the microphone during data acquisition could potentially be another major source of error in comparing noise reduction achieved between studies. Wu et al. (2014) showed large differences in the location and amplitudes of resonant peaks in the measured scanner transfer functions when the microphone was moved along the Z-axis in 5 cm increments close to a volunteer's left ear. A second experiment showed large differences in measured transfer functions when the microphone was at a fixed location during a scan of an empty scanner bore and thereafter with two different volunteers. Wu et al. (2014) showed that the transfer functions are highly dependent on microphone position along the scanner bore and is greatly influenced by the presence of objects / subjects inside the scanner bore. Considering this, the different microphone placements between this work and that of Schmitter et al. (2008) and Hutter et al. (2018) could contribute to the differences in measured noise reductions.

Pink noise

The pink noise introduced into the room via the scanner speaker did not completely mask the scanner noise but was able to break the silence between acquisitions and reduced the differences (jump) in noise amplitude thoughout the scanning session (Figure 3.4). This could potentially reduce the startling of patients, especially sleeping infants, caused by these abrupt changes in noise.

Foam enclosure

The addition of the foam enclosure proved an effective low-cost passive noise reduction method, although this solution is limited to infants due to size constraints. The enclosure is easily modifiable which allows for low cost improvements in the future.

Compared to results achieved by Nordell et al. (2009) who initially used a 38 mm polyurethane foam insert to achieve a 6 dB reduction in noise, our 50 mm foam only managed to reduce noise by 3.1 dBA for the standard acquisition. Compared to the 38 mm foam insert cut from a single large piece of foam (Nordell et al., 2009), our 50 mm foam enclosure was constructed by bending a flat foam plate, causing small density changes to the top and bottom surface of the foam at the bend which could have altered the sound absorbing properties, accounting for the lower reduction in noise.

4.2 Image analysis

The large difference in SNR between the DWIs acquired using the sinusoidal and filtered EPI readouts (Table 3.2) could be attributed to the substantial deviation from the standard EPI k-space trajectory in the phase encoding axis due to the constant phase blip of the sinusoidal readout implementation compared the widened blip of the filtered waveforms. Since additional regridding was only performed in the frequency axis the phase error persists. The error induced by the differences in sampling density between the filtered and sinusoidal readout in the frequency axis was mitigated by the density correction of the regridding algorithm in Equation 2.5.

The reason for the increase in SNR from the standard to filtered EPI readout seen in Table 3.2 along with the increased contrast seen in Figure 3.6 is not immediately apparent but could potentially be due to the lower maximum slew rate of the filtered sequence (see Table 2.1), which would lead to reduced eddy current induction and subsequently reduced eddy current related distortions. Induced eddy currents generate magnetic fields that oppose the direction of the intended gradient field, altering the net gradient field. This causes deviation in the expected k-space trajectory which leads to image distortion resulting in reduced SNR. Despite the sinusoidal waveform requiring the lowest slew rate, the filtered waveform deviates less from the standard k-space trajectory leading to greater SNR than the standard waveform. Based on the work of Pipe and Duerk (1995) who showed that the lowest estimator variance is achieved for a constant gradient amplitude, one would expect the SNR resulting from the filtered sequence to be somewhere between the trapezoidal and sinusoidal sequences. This might suggest that SNR could be optimized by finding the optimum ratio between reduced slew rate and deviation from standard k-space trajectory. Digital filters could easily be adjusted to find this ratio via simple filter adjustments.

Since the STD of the background noise was higher for both the filtered and sinusoidal readouts compared to the standard acquisition (Table 3.2) and increased with regridding kernel window size, we can conclude that the increase in SNR was not due to image smoothing.

The regridding kernel size greatly influenced the regridding time, with an additional 45 seconds or 7 minutes 45 seconds being required for image processing for acquisitions with window sizes of 2 and 4, respectively. This time is heavily dependent on the number of channels in the coil used for acquisitions as more channels increase the data throughput and thus regridding time. The additional regridding times required for a 32 channel adult head coil were 12:54 and 20:02 minutes for W = 2 and

4, respectively. This time could be reduced by optimising the reqridding algorithm and implementation.

4.3 Diffusion

The SNR differences in the corpus callosum did not translate to significant differences in the mean fractional anisotropy. Pairwise t-tests showed no significant differences in mean FA across the corpus callosum between the two standard acquisitions and any of the modified acquisitions. Since regridding kernel window widths of 2 and 4 produced similar results in terms of fractional anisotropy, there appears to be no advantage to using regridding window sizes larger than W = 2 as these greatly increase image reconstruction times.

Chapter 5

Conclusions and Future Work

The present work shows that digitally filtering the gradient waveforms effectively reduces the acoustic noise produced during MRI acquisition with no time penalty or reduction in image quality. The proposed method has the potential to be generalized to most gradient waveforms across a variety of sequences. With the addition of the passive noise reduction enclosure the combined noise reduction could greatly reduce infant anxiety during acquisitions, possibly increasing successful acquisition protocol completion rates. These methods could be particularly valuable when scanning infants in whom other noise reduction methods (such as active noise cancellation earphones) are not feasible.

Future work could include the filtering of all waveforms in a particular sequence. The proposed method is simple to implement and requires only the filtering of existing waveforms before being implemented on the scanner. This method can be realised with a fixed filter type and cut-off frequency or these could be varied as required. Since this method is not an analytical approach, it significantly reduces the complexity of designing new waveforms to reduce acoustic noise output.

The adjustment of the EPI echo spacing from 0.68 ms to 0.73 ms for the standard DTI had a large impact on peak acoustic noise SPL, with the maximum recorded difference being 3.9 dBA. By reducing the EPI harmonics, the effect of moving the EPI fundamental frequency to a trough in the FRF will be improved. This process is, however, highly dependent on the accuracy of the FRF. Wu et al. (2014) showed that MRI noise prediction using the independent linear model (Hedeen and Edelstein, 1997) with an error of less than 4% is possible, although prediction error increased to more than 30% for a subject mismatch where a FRF developed for one subject was used to predict noise during an acquisition with another subject in the scanner. In future, a patient specific FRF could be acquired before the commencement of the intended acquisition protocol to allow for quick adjustment of EPI echo spacing.

Bibliography

Berenice Goelzer, Colin H. Hansen, Gustav A. Sehrndt (2001). *Occupational exposure to noise : evaluation, prevention and control*. Geneva: World Health Organization, pp. 30–36.

Edelstein, William A. et al. (2002). "Making MRI quieter". In: *Magnetic Resonance Imaging* 20.2, pp. 155–163.

Edelstein, William A. et al. (2005). "Active-passive gradient shielding for MRI acoustic noise reduction". In: *Magnetic Resonance in Medicine* 53.5, pp. 1013–1017.

Forbes, Lawrence K. et al. (2007). "An analytical approach to the design of quiet cylindrical asymmetric gradient coils in MRI". In: *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* 31.4, pp. 218–236.

Foster, John R. et al. (2000). "Sound-Level Measurements and Calculations of Safe Noise Dosage During EPI at 3 T". In: *Journal of Magnetic Resonance Imaging* 12.1, pp. 157–163.

Garyfallidis, Eleftherios et al. (2014). "Dipy, a library for the analysis of diffusion MRI data". In: *Frontiers in Neuroinformatics* 8, p. 8.

Hedeen, Robert A. and William A. Edelstein (1997). "Characterization and prediction of gradient acoustic noise in MR imagers". In: *Magnetic Resonance in Medicine* 37.1, pp. 7–10.

Hennel, F., F. Girard, and T. Loenneker (1999). "'Silent' MRI with soft gradient pulses". In: *Magnetic Resonance in Medicine* 42.1, pp. 6–10.

Hennel, Franciszek (2001). "Fast spin echo and fast gradient echo MRI with low acoustic noise". In: *Journal of Magnetic Resonance Imaging* 13.6, pp. 960–966.

Hutter, Jana et al. (2018). "Quiet echo planar imaging for functional and diffusion MRI". In: *Magnetic Resonance in Medicine* 79.3, pp. 1447–1459.

Jackson, James M. et al. (2011). "Could image processing enable reduced acoustic noise during MRI scans?" In: *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* 39 B.4, pp. 191–205.

Jackson, John I. et al. (1991). "Selection of a Convolution Function for Fourier Inversion Using Gridding". In: *IEEE Transactions on Medical Imaging* 10.3, pp. 473–478.

Kannan, Govind et al. (2011). "An efficient feedback active noise control algorithm based on reduced-order linear predictive modeling of fMRI acoustic noise". In: *IEEE Transactions on Biomedical Engineering* 58.12 PART 1, pp. 3303–3309.

Moelker, Adriaan, Mika W. Vogel, and Peter M. T. Pattynama (2003). "Efficacy of passive acoustic screening: Implications for the design of imager and MR-suite". In: *Journal of Magnetic Resonance Imaging* 17.2, pp. 270–275.

Nordell, Anders et al. (2009). "The acoustic hood: a patient-independent device improving acoustic noise protection during neonatal magnetic resonance imaging". In: *Acta Paediatrica* 98.8, pp. 1278–1283.

Pipe, James and Jeffrey L. Duerk (1995). "Analytical resolution and noise characteristics of linearly reconstructed magnetic resonance data with arbitraryk-space sampling". In: *Magnetic Resonance in Medicine*.

Ravicz, Michael E and Jennifer R Melcher (2001). "Isolating the auditory system from acoustic noise during functional magnetic resonance imaging: Examination of noise conduction through the ear canal, head, and body a)". In: *The Journal of the Acoustical Society of America* 45.1, pp. 216–231.

Reuter, Martin, H Diana Rosas, and Bruce Fischl (2010). "Highly Accurate Inverse Consistent Registration: A Robust Approach". In: *NeuroImage* 53.4, pp. 1181–1196.

Salahat, Ramy et al. (2013). "A simple and efficient approximation to the modified Bessel functions and its applications to Rician fading". In: *2013 7th IEEE GCC Conference and Exhibition, GCC 2013*, pp. 351–354.

Schmitter, S. et al. (2008). "Silent echo-planar imaging for auditory FMRI". In: *Magnetic Resonance Materials in Physics, Biology and Medicine* 21.5, pp. 317–325.

Smink, J, GJ Plattel, and PR Harvey (2007). "General methods for acoustic noise reduction by avoiding resonance peaks". In: *Proceedings of the 16th*.

Tomasi, D et al. (2005). "fMRI-acoustic noise alters brain activation during working memory tasks." In: *NeuroImage* 27.2, pp. 377–86.

Wu, Ziyue et al. (2014). "Evaluation of an independent linear model for acoustic noise on a conventional MRI scanner and implications for acoustic noise reduction". In: *Magnetic Resonance in Medicine* 71.4, pp. 1613–1620.

Zhang, Nanyin, Xiao Hong Zhu, and Wei Chen (2005). "Influence of gradient acoustic noise on fMRI response in the human visual cortex". In: *Magnetic Resonance in Medicine* 54.2, pp. 258–263.