

Super Resolution of HARDI images Using Compressed Sensing Techniques

by

Eric Lum

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2015

© Eric Lum 2015

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Effective techniques of inferring the condition of neural tracts in the brain is invaluable for clinicians and researchers towards investigation of neurological disorders in patients. It was not until the advent of diffusion Magnetic Resonance Imaging (dMRI), a noninvasive imaging method used to detect the diffusion of water molecules, that scientists have been able to assess the characteristics of cerebral diffusion *in vivo*. Among different dMRI methods, High Angular Resolution Diffusion Imaging (HARDI) is well known for striking a balance between ability to distinguish crossing neural fibre tracts while requiring a modest number of diffusion measurements (which is directly related to acquisition time).

HARDI data provides insight into the directional properties of water diffusion in cerebral matter as a function of spatial coordinates. Ideally, one would be interested in having this information available at fine spatial resolution while minimizing the probing along different spatial orientations (so as to minimize the acquisition time). Unfortunately, availability of such datasets in reasonable acquisition times are hindered by limitations in current hardware and scanner protocols. On the other hand, post processing techniques prove promising in increasing the effective spatial resolution, allowing more detailed depictions of cerebral matter, while keeping the number of diffusion measurements within a feasible range.

In light of the preceding developments, the main purpose of this research is to look into super resolution of HARDI data, using the modern theory of compressed sensing. The method proposed in this thesis allows an accurate approximation of HARDI signals at a higher spatial resolution compared to data obtained with a typical scanner. At the same time, ideas for reducing the number of diffusion measurements in the angular domain to improve the acquisition time are explored. Accordingly, the novel method of applying two distinct compressed sensing approaches in both spatial and angular domain, and combining them into a single framework for performing super resolution forms the main contribution provided by this thesis.

Acknowledgements

I would like to begin by thanking my supervisor Oleg Michailovich for teaching me countless signal and image processing techniques during my later undergraduate years, and inspiring me in the field of diffusion MRI during my time as a MASC student. His remarkable patience and tolerance are distinguishing characteristics among the many teachers I have had throughout my times as a student.

My colleagues Iván Camilo Salgado Patarroyo and Quan Zhou were great friends that have provided me with invaluable feedback during our relatively short but cherished times together. Their support during difficult times was highly appreciated, and it was truly an honour to have the opportunity to work beside these great minds during my time as a MASC student.

Table of Contents

List of Tables	vii
List of Figures	viii
1 Introduction	1
1.1 Medical Imaging	1
1.2 Brain Imaging	2
1.3 Contributions of this thesis	3
1.4 Organization of this thesis	3
2 Principles of MRI	5
2.1 Physics of MRI	5
2.2 MRI Image Acquisition	8
2.3 Image Encoding	10
2.4 Echo Planar Imaging	11
2.5 Spatial Resolution and FOV relationships in MRI	12
2.6 Diffusion MRI	14
2.7 Diffusion parameters and the diffusion equation	17
3 Fundamental dMRI models	19
3.1 Diffusion Tensor Imaging	20
3.2 Diffusion Spectrum Imaging	23
3.3 High Angular Resolution Diffusion Imaging	24

4	Working with HARDI data	26
4.1	Gaussian Mixture Model	26
4.2	Diffusion Basis Functions	28
4.3	Q ball Imaging	29
4.4	QBI with Spherical Harmonics	31
4.5	Spherical Ridgelets	35
5	Resolution Enhancement in Spatial and Angular Domains	40
5.1	Spatial Methods	41
5.2	Angular Methods	42
5.3	Super resolution via Compressed Sensing: DBF Decomposition	43
6	Main Contribution	47
6.1	Organization of HARDI datasets	47
6.2	Framework of GCSR	49
6.3	Solution via ADMM	54
7	Results	59
7.1	Qualitative assessment with ODFs	60
7.2	Quantitative metrics for HARDI data	65
7.3	Fiber Tractography	68
7.3.1	Streamline Tractography	69
7.3.2	Quantitative metrics in Streamline Tractography	71
7.3.3	Experimental results from Fiber Tractography	72
7.3.4	Comparison with denoised measurements	74
8	Summary and Future Works	76
	References	78

List of Tables

7.1	Comparison of tractography metrics between GCSR and GMM Reconstructed data	74
7.2	Comparison of tractography metrics between GCSR and NLM filtered version of ground truth data	75

List of Figures

2.1	Traversal of k-space using EPI. Acquisition occurs during frequency encoding, a subset of which are depicted by black squares.	12
2.2	Pulse sequences corresponding to a) Spin-Warp and b) EPI. The white lines in figures b) and d) show their respective trajectory in the k-space. Note that application of the π pulses have been omitted from the diagram to avoid cluttering. Credit [39]	13
2.3	Phase alignment of water molecules in the absence of diffusion maximizes the measured signal. Credit [51]	15
2.4	Example of how scan-dependent parameters are incorporated into the PGSE sequence. Credit [39].	17
3.1	Figure illustrating how the q-space is sampled for (a) DSI (b) DTI (c) HARDI (d) Multi-shell Diffusion Imaging. The q-space origin is highlighted in red. Credit [4].	25
4.1	Representation of the Real and Symmetric Spherical Harmonic basis functions up to $l = 6$. Note that they functions have been normalized such that they lie in the range [-1,1]. Credit [4].	33
4.2	Illustration of Spherical Ridgelets created using $j = 3$ resolution levels, and bandwidth $\rho = 2$	38
5.1	Different Sparse Representation of HARDI data. Credit [46].	44
6.1	Matrix representation of a $K \times \Omega_k $ HARDI dataset s , where $K = 4$ and $ \Omega_k = 9$. Diffusion encodings of the center pixel is represented as s_5 , whereas the DWI corresponding to diffusion encoding direction $[1 \ 1]/\sqrt{2}$ is depicted by s^1	48

7.1	Projection of each individual fiber via ODFs onto a 2D plane. The final computer data is the superposition of all three fibers onto the same FOV, along with grey matter simulation in locations without fiber presence. . . .	61
7.2	ODFs captured from computer simulated data at a FOV of 32×32 , where correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction.	62
7.3	ODFs captured from computer simulated data contaminated with Rician noise at $\text{SNR} = 5$, where correspondence is defined as: (Left) Ground Truth without noise, (Middle) GCSR, (Right) GMM Reconstruction	63
7.4	ODFs captured from HCP <i>in vivo</i> data near the left and right ventricles, where correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction	64
7.5	ODFs captured from HCP <i>in vivo</i> data in regions exhibiting bending and curving geometries. Correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction	65
7.6	NMSE (top row) and AAE metrics (bottom row) for reconstructions of computer simulated data at various SNR levels and b-values.	67
7.7	NMSE (left) and AAE metrics (right) for reconstructions of HCP <i>in vivo</i> data in various regions. Fiber geometries of smooth, bending and curving correspond to Figures 7.4, 7.5 (top) and 7.5 (bottom) respectively.	68
7.8	Various bundles of fiber tracts generated via streamline tractography using the Camino software.	69
7.9	Fiber tractography using HCP <i>in vivo</i> data (see Figure 7.5), where correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction	70
7.10	Selective fibers in the designated ROIs are chosen for error metrics based on their distinctive geometries.	73

Chapter 1

Introduction

1.1 Medical Imaging

The first instance of body anatomy imaging was documented in 1895, when the German physics professor Wilhelm Röntgen used X-rays to form an image illustrating the bone structures within a human hand [35]. What made this discovery significant was its ability to capture contrasts delineating between bone and tissue from *within* the body without the need for surgery on living subjects or dissection in post-mortem analysis. As a result of his revolutionary discovery, he was issued the first Nobel Prize in history [35]. Notably, this technique proves to be extremely useful in the diagnosis stage where abnormalities could be identified as part of preliminary analysis. From then on, the term medical imaging was used to describe methods of obtaining information about a patient's internal body structure for clinical analysis. This is of a notably different philosophy from images obtained through photography, which is clearly impractical in this context owing to the impossible placement of the imaging device.

In the last century, advances in science have uncovered different methodologies for medical imaging, each relying on varying properties of physics allowing each to retrieve unique information in different parts of the body. For example, ultrasound deploys sound waves above the capability of human hearing and can be used to provide real time information on blood flow, while Magnetic Resonance Imaging (MRI) uses strong magnetic fields to capture the spin density of water molecules which can be later used to infer anatomical structures and pathologies. Furthermore, medical imaging is a continuously evolving field in the sense that new breakthroughs will be reflected in improved image quality, leading to more accurate diagnosis. For this reason, research is always ongoing for ways to

improve both the acquisition procedure and post-processing methods of improving image reconstruction. In this thesis, the focus will be on a specific variant of MRI referred to as diffusion MRI (dMRI) and its applications towards brain imaging.

1.2 Brain Imaging

The brain is considered one of the most mysterious parts of the human body due to its complicated micro-structure. While its understanding is still far from being complete, considerable efforts over the past years have yielded much information that has been applied towards the identification of different anatomical structures and their correspondence to motor and cognitive functions, as well as the identification of mental disorders such as depression or schizophrenia [35]. Examples of different modalities of medical imaging well equipped to brain imaging include and are not limited to, X-ray computed tomography (CT), positron emitted tomography (PET) and MRI. The presence of multiple modalities suggests that each provides a different set of information. Indeed, as an example CT is used to probe structures of the brain but cannot provide insight on the change in signal due to functional changes. PET can overcome this shortcoming by capturing functional behaviour in different parts of the brain, but requires the injection or inhaling of radioactive agents to facilitate the imaging process.

With that said, MRI is an invaluable tool for brain imaging due to its ability to generate high levels of detail, while being a non-invasive imaging modality (that is, avoiding the use of radioactive tracers that have a low probability of inducing cancer in the subject). Another advantage of MRI is its versatility; small changes in the imaging protocol defines what information can be acquired from the scan. For example, functional MRI (fMRI) is often used to interpret the change in signal arising as a consequence of neural activity, while structural MRI is more traditional in the sense that it provides information on the structure of underlying white matter, grey matter and cerebrospinal fluid (CSF). On the other hand, diffusion MRI has seen a rapid growth over the previous two decades owing to increased scanner hardware performance and its ability to sense the diffusivity of water molecules within the subject. This leads to interesting applications such as the ability to reconstruct the trajectories of neural fiber tracts in the human brain. Unfortunately, one significant downfall of dMRI is its prohibitively long acquisition times. This is particularly problematic for patients afflicted with dementia or children, since high levels of subject motion arising from these patients causes significant imaging artifacts in the resulting scan. This problem has been recognized and has been the subject of much research in the dMRI community, leading to a variety of techniques and models proposed to tackle

this problem. Further compounding this problem is the need for higher spatial resolution (e.g., a smaller voxel size) and angular sampling (e.g., number of diffusion orientations measured) in dMRI data, the addressing of which is a main goal in this thesis.

1.3 Contributions of this thesis

To expand on the aforementioned low resolution issues, the limitations regarding spatial and angular resolution will be elaborated. Owing to the inherent low spatial resolution of the dMRI data, the fiber structure at a given spatial location is subject to partial voluming effects; that is, the presence of different types of tissue within a voxel. The implications stemming from partial voluming effect can be seen in the following example; instead of identifying a single fiber within a $1\text{ mm} \times 1\text{ mm} \times 1\text{ mm}$ voxel, in a $2\text{ mm} \times 2\text{ mm} \times 2\text{ mm}$ voxel we must account for the presence of multiple differently oriented fibers. On the other hand, the diffusion profile is constructed by scanning the entire Field of View (FOV) K times, where K corresponds to the number of diffusion measurements varying in both magnitude and direction. The ability to detect the presence of multiple oriented fibers is highly dependent on sufficient probing in different diffusion directions. Sufficiently accurate representation of the diffusion profile would require very dense sampling, which is out of the question due to prohibitively long acquisition times. Fortunately, there are sampling schemes in the angular domain such as High Angular Resolution Diffusion Imaging (HARDI) that keep K at feasible levels while retaining reasonably accurate modelling of the underlying diffusion process. That being said, the goal of the thesis is to post-process the dMRI data to provide high angular and spatial resolution reconstructions that can alleviate the partial voluming effect and long acquisition times. Through the introduction of regularization constraints, the method will be posed as a convex optimization problem. A convenient framework for solving this problem is by using an iterative variable splitting approach. Through this method, a complicated optimization problem in multiple variables can be reduced to a sequence of sub-problems of a single variable, each admitting simpler closed form solutions. Moreover, the idea of variable splitting alludes to the possibility of parallel processing, which can make it remarkably computational efficient in practice.

1.4 Organization of this thesis

The remainder of the thesis is organized as follows. Chapter 2 will talk about the basics of MRI as well as an introduction to diffusion MRI, the primary modality of interest in

this thesis. Chapter 3 builds on dMRI by talking about different formalisms regarding how dMRI data can be best used to model the underlying diffusion processes. Chapter 4 will focus on the HARDI philosophy and how the techniques discussed in the previous chapter can be adapted to work with HARDI data. The primary goal of Chapter 5 is to address the problem of inherent low resolution of HARDI data and study different approaches in the literature taken to address this problem. In Chapter 6, the main contribution of this thesis will be revealed, which builds on some of the ideas developed by other dMRI researchers to propose a framework for applying super resolution to available HARDI data. Chapter 7 tests this proposed methodology in a series of experiments ranging from computer simulated data to *in vivo* data and fiber tractography. Finally, Chapter 8 will summarize the key ideas covered in this thesis as well as providing a conclusion and ideas that could be used in future works.

Chapter 2

Principles of MRI

The branch of medical imaging known as MRI was made possible by the phenomenon of Nuclear Magnetic Resonance (NMR) along with the contributions from many scientists throughout the last several decades. Its success stems from many factors, including high signal to noise ratio (SNR) compared to other medical imaging techniques such as ultrasound, and avoiding the need to introduce radiochemical tracers into the subject which, despite the extremely low probabilities, has seen cases where cancer was induced in the subject. Last but not least, it is also non-invasive and does not depend on the use of ionizing radiation. This goal of this section is to present an overview of the physics involved in MRI as well as studying imaging techniques which can drastically reduce the acquisition time. The information will prove invaluable in providing insights to the need of such methods as well as downfalls associated.

2.1 Physics of MRI

Ever since MRI has been discovered about half a century ago, acquisition times have significantly improved through advances not only in technology, but equally importantly by optimizing the acquisition process. This section will first explain the physics behind MRI, followed by special acquisition techniques which will give insight to how spatial resolution can be improved.

According to quantum mechanics, atoms with an odd number of protons are designated a spin of $\pm\frac{1}{2}$, through which they intrinsically carry an associated magnetic moment and angular momentum. Under the presence of a strong magnetic field B_0 , the magnetic

moments of the atoms will tend to align along the B_0 field. The strength of the B_0 field is typically measured in terms of Teslas [T], with higher strengths desirable due to their ability to align a larger proportion of protons in the sample. However, it is important to note that having too strong of a magnetic field can potentially cause safety hazards. Common field strengths of MRI scanners for human brain imaging are 1.5 T or 3 T, with 3 T scanners becoming more popular in recent times due to increased hardware capability. In addition to alignment of magnetic moments, the atoms also undergo a process known as *precession* about the induced magnetic field. Precession refers to rotation about an axis that is constantly changing; the effect can be visualized as the spinning of a top. Furthermore, the angular frequency ω of the spins and the net magnetic field B experienced by the atom are proportional to each other as governed by the Larmor equation

$$\omega = \gamma B \tag{2.1}$$

where γ is the gyromagnetic constant and varies between different nuclei. For hydrogen atoms (protons), $\gamma = 2.675 \times 10^6$ rad/s · T. The application of the B_0 field can be thought of as an initial alignment of the magnetic moments in the sample along a direction, which without loss of generality is commonly referred to as the *longitudinal axis*.

Note that the MRI scanner is unable to detect different tissue contrasts with the B_0 field alone. Rather, it detects the rate at which protons return to their equilibrium state (e.g., precession along the longitudinal axis) after an external stimuli is applied and shortly removed thereafter. To facilitate this, an electromagnetic RF pulse is applied to the sample for a short duration and subsequently turned off. This RF pulse is commonly known as an *excitation pulse* due to its effect of providing a stimuli to the protons, and also as a $\frac{\pi}{2}$ pulse due to the orientation of the spins being altered by an angle of $\frac{\pi}{2}$. This excitation pulse effectively realigns the spins into the plane perpendicular to the longitudinal axis, which goes by the name of the *transverse plane*. Upon removal of the RF pulse, the protons will try to recover their original orientation along the longitudinal axis. This process is known as relaxation and comes in two different flavours; longitudinal and transverse relaxation. As their name implies, the different types of relaxation describe recovery towards equilibrium along the longitudinal axis and the transverse plane respectively.

Owing to the $\frac{\pi}{2}$ pulse pushing the spin orientation entirely into the transverse plane, during longitudinal relaxation the magnitude of the magnetic moments along the longitudinal axis is steadily increasing as the protons settle towards the equilibrium state (with magnetization at equilibrium given by M_0). The differential equation governing longitudinal relaxation is given by

$$\frac{dM_z}{dt} = \frac{M_0 - M_z}{T_1} \quad (2.2)$$

where M_z describes the projection of the magnetization vector along the longitudinal axis, and T_1 is a constant used to characterise the longitudinal relaxation times of the sample and is typically around one second. It can be seen that (2.2) is a separable differential equation with solution $M_z = M_0(1 - e^{-\frac{t}{T_1}})$. The solution agrees with the observation that the longitudinal magnetization recovers as time goes on, eventually settling back to M_0 at $t = \infty$. Owing to the ability of longitudinal relaxation to illustrate the differences between certain tissues effectively, in MRI literature this is commonly referred to as T_1 relaxation.

On the other hand, transverse relaxation refers to the recovery of spin orientation towards equilibrium state, but this time being measured along the transversal plane. Since the initial magnetization was entirely along the longitudinal axis and thus exhibits no transversal component, one would expect that transversal relaxation would result in a decrease of magnitude. Indeed, given M_{xy} being the projection of the magnetization vector along the transversal plane and T_2 being a constant used to characterise the transversal relaxation time of the sample, the differential equation for transversal relaxation is given by

$$\frac{dM_{xy}}{dt} = -\frac{M_{xy}}{T_2} \quad (2.3)$$

from which we can see the solution is $M_{xy} = M_{xy}(0)e^{-\frac{t}{T_2}}$, an equation describing exponential decay which agrees with our observation. However, transversal relaxation is complicated by magnetic field inhomogeneity, which causes excited protons at different locations to precess at slightly different frequencies. This translates into loss of phase coherence as times passes. This effect is known as dephasing, and causes the net magnetization to be attenuated in addition to transversal relaxation. Since this phase incoherence does not have a useful physical interpretation, it is desirable to have it removed. To facilitate this, another RF pulse is applied after a specified duration $TE/2$ with the goal of reversing the phases. This RF pulse is known as the π pulse owing to phase reversal requiring the spin orientations to be modulated by an angle of π . In MRI nomenclature, it is also referred to as the refocusing pulse due to its objective of recovering phase coherency which in turn strengthens the measured signal. After the π pulse is applied, the reversal of the phases causes dephasing to occur in the opposite direction. For the phases to resume coherency, a symmetric duration of $TE/2$ must occur so that the initial dephasing can be completely undone. At this point, dephasing has been removed from the transverse magnetization and is ready to be captured by the MRI scanner. Altogether, this process of removing

phase incoherence is known as the spin echo protocol. The quantity TE refers to Time Echo which is an MRI acquisition parameter describing the time between application of the initial RF excitation and the time at which a spin echo occurs. For experiments which measure multiple instances of spin echoes, TE is usually measured against the spin echo with the largest magnitude.

2.2 MRI Image Acquisition

We have discussed how protons are excited by an RF pulse and subsequently undergo relaxation along both the longitudinal axis and the transversal plane in a typical MRI experiment. However, the mechanism behind the MRI scanner’s ability to detect the changing magnetization during relaxation has not been discussed. Recall that the Larmor equation (2.1) states that the angular frequency of the spins is proportional to the magnetization exerted on the proton. By virtue of Faraday’s Law, the receiver coils of the MRI scanner are able to pick up the current induced by the magnetic field of the proton’s precessing spins. The contrasts of the tissue captured by the MRI scanner depends on the strength of the current induced, which will depend both on the coherence of the phases as well as the rate of relaxation as discussed earlier.

Now that it is clear what signals are being captured by the MRI scanner, we will move on to how the subject is captured by the scanner. The MRI scanner reads data in a 2D manner and uses slice selection techniques to move between different slices of the subject. Since a 3D volume can be thought of as a sequence of 2D slices stacked together, this technique allows a 3D volume to be captured by the MRI scanner by scanning each slice and consolidating the results into a single volume. The key to selecting parts of the sample to be imaged lies in magnetic field gradients, which are used as temporary measures of changing the net magnetic field. In the following developments, we use \mathbb{R}^3 to represent the standard 3D Cartesian grid with unit vectors $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}\}$ along each orthogonal axis. Denoting the homogeneous magnetic field by B_0 and the magnetic field gradient vector by $\mathbf{g} = [G_x \ G_y \ G_z] \in \mathbb{R}^3$, the net magnetic field B at spatial location $\mathbf{r} = [x \ y \ z]^T \in \mathbb{R}^3$ is described by the equation

$$B = B_0 + \mathbf{g} \cdot \mathbf{r} \tag{2.4}$$

Without loss of generality, let the unit vector $\hat{\mathbf{z}}$ be parallel to the B_0 field (e.g., pointing along the longitudinal axis). With this formulation, the z-component G_z is responsible for slice selection while G_y and G_x are used to image each 2D slice. In light of this development,

slice selection and how the MRI scanner captures each individual slice will be discussed separately.

The slice select gradient causes the frequencies in the sample to change linearly along the \hat{z} direction. By doing so, a distinct frequency can be assigned to each transversal plane slice of the sample. Recall that an excitation RF pulse rotates the spin orientation by $\frac{\pi}{2}$. Aside from that purpose, slice selection can be achieved by coinciding the central frequency of the excitation pulse with the Larmor frequency of precessing spins within the desired slice (due to the slice selection gradient, each slice experiences a different Larmor frequency).

Next, we will illustrate how the MRI scanner reads each transversal slice. The magnetic field gradient components along the \hat{x} and \hat{y} direction together changes the effective frequency in each spatial location along the selected transversal slice. After a duration t , the accumulation of the phase θ at each spatial location (x, y) can be expressed as

$$\theta = 2\pi(k_x(t)x + k_y(t)y) \quad (2.5)$$

where the quantities $k_x(t)$ and $k_y(t)$ are related to the integral of the corresponding component of the magnetic field gradients and are given by

$$\begin{aligned} k_x(t) &= \gamma \int_0^t G_x(\tau) d\tau \\ k_y(t) &= \gamma \int_0^t G_y(\tau) d\tau \end{aligned} \quad (2.6)$$

In the following discussion, the dependence on time for $k_x(t)$ and $k_y(t)$ is understood and usually omitted for brevity in notation. Since the MRI scanner's receiver coils senses the current induced by the precessing spins in accordance to Faraday's Law, the net current is equivalent to the summed contributions of the spin density $f(x, y)$ at each spatial location (x, y) on the transversal plane. Accounting for the distinct phase accumulation at each spatial location, the current $s(t)$ detected by the receiver coils is thus

$$s(t) = \int \int f(x, y) e^{2\pi i(k_x x + k_y y)} dx dy \quad (2.7)$$

where $i = \sqrt{-1}$, which indicates a Fourier relationship between the spin density and the induced current. With equation (2.7), from the scanner data it is possible to recover $f(x, y)$ via inverse Fourier transform. In the MRI community, the k-space $(k_x, k_y) \in \Omega_k$ is typically

referred to as scanner space since it refers directly to the space in which data is captured, whereas $(x, y) \in I$ is denoted as the image space since it represents the spatial locations within the FOV I .

2.3 Image Encoding

In the previous section we have discussed how the spins (and by extension, the magnetization due to the Larmor equation) can be recovered from the scanner data. However, the given procedure of applying gradients in the transversal plane after RF excitation and slice selection is only able to capture one point in the k-space. By equation (2.7), it follows that the image of the spins can be recovered only after the k-space is fully sampled. Clearly, repeating this whole process for each location in the k-space is an extremely inefficient endeavour, thus the motivation of this section is to present time efficient acquisition of the k-space.

Let us first revisit the k-space. It consists of two orthogonal axes k_x and k_y , which in MRI nomenclature are used to represent directions corresponding to frequency and phase encoding respectively. The phase encoding gradient is responsible for altering the phases of the magnetization in the protons. Through phase encoding different rows of the k-space become out of phase with respect to another while all protons of the same row all share the same phase perturbation. This technique allows the MRI scanner to distinguish between different different rows in the k-space. On the other hand, the frequency encoding gradient works in the same way but alters the precession frequency of the protons along each column of the k-space. In this manner, we can see frequency and phase encoding in conjunction are able to isolate a specific k-space coordinate.

The development of MRI acquisition schemes relies on the application of RF excitation pulses along with the phase / frequency encoding gradients in an orderly manner. This is collectively known as a pulse sequence which dictates the trajectory travelled in the k-space. The Spin-Warp pulse sequence is a basic yet practical MRI imaging method, and is the foundation for more complicated sequences such as Echo Planar Imaging (described later). After the initial excitation RF pulse, a phase encoding gradient is applied for a short duration to act as a row selector. It is subsequently turned off and the frequency encoding gradient is then activated. Simultaneously, the MRI scanner's receiver coils are turned on during the duration of the frequency encoding gradient, and an entire row of the k-space is read per RF excitation in this manner. The time between successive RF excitations on a fixed slice is referred to as Repetition Time (TR). The number of samples made per TR is chosen to minimize artifacts within the given FOV [13]; typically it is between 128 and 512

samples. Due to the frequency encoding step occurring on the order of several milliseconds whereas phase encoding entails application of a new RF excitation pulse, it is clear that a standard MRI acquisition resolution in the phase encoding direction (k_y) is limited when compared to frequency encoding (k_x).

2.4 Echo Planar Imaging

Previously we have seen how an entire line of k-space can be captured during one RF excitation. In this section Echo Planar Imaging (EPI) will be introduced, a scanner protocol capable of capturing the *entire* k-space in one RF excitation, as shown in Figure 2.1. A key feature is its single continuous k-space trajectory, which maximizes the uptime of the readout gradient. It can also be seen that the frequency encoding follows a zigzag pattern. This can be achieved by switching the sign of the frequency encoding gradient. Consequentially, EPI demands high performing gradients that are able to switch rapidly (recall that one row of frequency encoding is typically performed on the order of milliseconds). This is quantified by the gradient's slew rate which determines the trajectory's acceleration. However, any phase errors that may accumulate during scanner readout are not reset at each phase encoding. This was not a problem with the Spin-Warp pulse sequence because applying a new RF excitation for each phase encoding effectively meant each k-space line undergoes a phase reset. Last but not least, EPI is constrained by the relaxation parameters of the subject for imaging of T2 contrasts. Recall that T2 relaxation is governed by an exponential decaying function of the transverse relaxation starting from application of the RF excitation pulse. Due to the long k-space trajectory of EPI, sampling of the entire k-space must be done before this decay reaches undetectable levels. Fortunately, this can be overcome by improved scanner hardware in the form of high magnitude gradients that will increase the speed of the k-space trajectory.

Despite the above downfalls of EPI, the primary motivation behind using this protocol comes from subject motion during the scanning of a slice. With slower imaging schemes, motion artifacts can still be observed in the resultant image. On the other hand, the ability of EPI to scan an entire slice in 40 ms [73] means subject motion will not be significant during this small time frame. The use of EPI is practically restricted to cardiac imaging as well as dMRI and fMRI. These applications share the need of having multiple acquisitions of the subject in quick succession in order to minimize the misalignment introduced by motion.

Recall that the k-space can be interpreted as the frequency representation of the magnetization which resembles an image of the sample, and as such it shares some characteristics




Figure 2.1: Traversal of k-space using EPI. Acquisition occurs during frequency encoding, a subset of which are depicted by black squares.

of natural images. Namely, the magnitude of the frequency components close to the origin of the k-space are in general more prevalent than the high frequency components. Therefore, the definition of TE is applied to the spin echo corresponding to the measurement at the origin of k-space. To solidify the concepts of how the frequency and phase encoding gradients work in conjunction in each of the aforementioned pulse sequences, it is helpful to analyse a timing diagram describing the state of each gradient at any given time. This is shown in Figure 2.2. Recently, more sophisticated k-space sampling techniques such as spiral sampling [39] have shifted to denser sampling near the k-space origin in order to improve image reconstruction quality. Implementing the spiral trajectory requires more complicated waveforms for the phase and frequency encoding, and the k-space samples cannot be arranged in a Cartesian grid. Unless the measurements from spiral sampling are properly extrapolated to a Cartesian grid, the standard Fast Fourier Transform cannot be used to reconstruct the image via (2.7).

2.5 Spatial Resolution and FOV relationships in MRI

From equation (2.7), a Fourier relationship was discovered between the current picked up by the MRI scanner and the resultant image consisting of the spin densities. Based on this, we can formulate the Nyquist conditions determining the FOV encompassed by the resultant image as follows

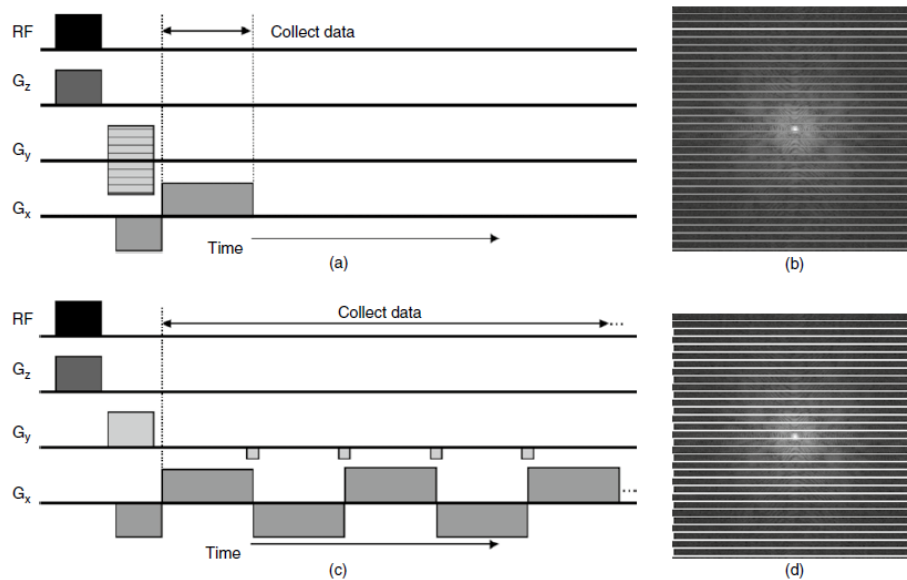


Figure 2.2: Pulse sequences corresponding to a) Spin-Warp and b) EPI. The white lines in figures b) and d) show their respective trajectory in the k-space. Note that application of the π pulses have been omitted from the diagram to avoid cluttering. Credit [39]

$$FOV = \frac{1}{\Delta k} \quad (2.8)$$

where Δk denotes the step size between adjacent samples in the k-space. This equation tells us that a densely sampled k-space yields a larger area covered in a MRI scan, and is of particular importance to MRI technologists for determining an optimal sampling density based on the desired rendering of anatomical structures. Violation of this criterion, that is insufficient sampling density in the k-space results in image aliasing artifacts.

Of more importance to this thesis however is the second Nyquist condition describing limitations in spatial resolution, which refers to the minimum distance for which different image intensities can be detected. This relation is given by

$$\text{Spatial Resolution} \propto |k_{max}| \quad (2.9)$$

where the quantity $|k_{max}|$ represents the maximum sampled frequency in the k-space. Intuitively, an analogy can be made to a natural image whereby higher frequency components reveal more subtle details over close proximities. Unfortunately, high frequencies are typi-

cally also associated with noise contamination which degrades the image quality, resulting in a lower Signal to Noise Ratio (SNR). Thus, it can be seen that improvement of spatial resolution comes at the cost of lower SNR, which is particularly problematic for dMRI where the desired signal is already attenuated from the diffusion of water molecules. Finally, it should be noted that an improved spatial resolution in a fixed FOV results in an increased number of voxels, which will increase the acquisition time of the experiment. Therefore, the ability of EPI to acquire the whole transversal plane in a single TR is offset by the loss in SNR and spatial resolution. Because dMRI studies typically use EPI to obtain the full set of DWIs in a timely manner, this is why a typical spatial resolution in dMRI is about $2 \text{ mm} \times 2 \text{ mm} \times 2 \text{ mm}$.

2.6 Diffusion MRI

While the MRI pulse sequences described in the previous section allows one to view the interior structure of the brain, it is unable to map the axonal tracts within the brain which are responsible for the transfer of information. Because the diameter of these tracts are on the order of $20 \mu\text{m}$ [26, 28], it is very difficult to identify them even during post-mortem analysis on the subject. To overcome this issue, researchers have formulated the myelin hypothesis [43] which states that within these axonal tracts, diffusion of water molecules perpendicular to the axonal length will be severely hindered by the myelin sheathes enclosing the axon. On the other hand, water diffusion along the tracts will be relatively unobstructed. The purpose of diffusion MRI (dMRI) is to measure and classify water diffusion behaviour.

Before proceeding to how dMRI works, it is instructive to review the concept of self diffusion or Brownian motion. Generally speaking, diffusion is characterised by a mass transport process in nature, resulting in the mixing mechanism of molecules or particles [39]. While it is well known that diffusion can describe the net flux of molecules from a high concentration region spreading to low concentration region as given by Fick's first law, perpetual molecular movement can be observed even after thermodynamic equilibrium is established in the sample. This phenomenon is referred to as self diffusion, and is unrelated to molecular transport induced by bulk motion (e.g., convection, dispersion), where the molecules move in a coherent manner. Instead, the self diffusion (henceforth abbreviated as simply diffusion) of any given particle over a given time can be described as a random walk [39]. Given this revelation, it is clearly not feasible to draw conclusions on water diffusion in the brain by observing the random walk of each water molecule within the brain. In 1905, Albert Einstein proposed that the three dimensional random walks exhibited by molecules

in unrestricted environments have a mean square displacement $X \in \mathbb{R}^3$ characterized by

$$E[X^2] = 6Dt \tag{2.10}$$

where D is the diffusion coefficient of the molecule in question and t corresponds to the time duration allotted for molecules to diffuse. This equation allows one to study the diffusion behaviour by analysing the ensemble average displacement of these random walks as opposed to each individual molecule. With this in mind, let us see how MRI can be adapted to measure the extent of diffusion.

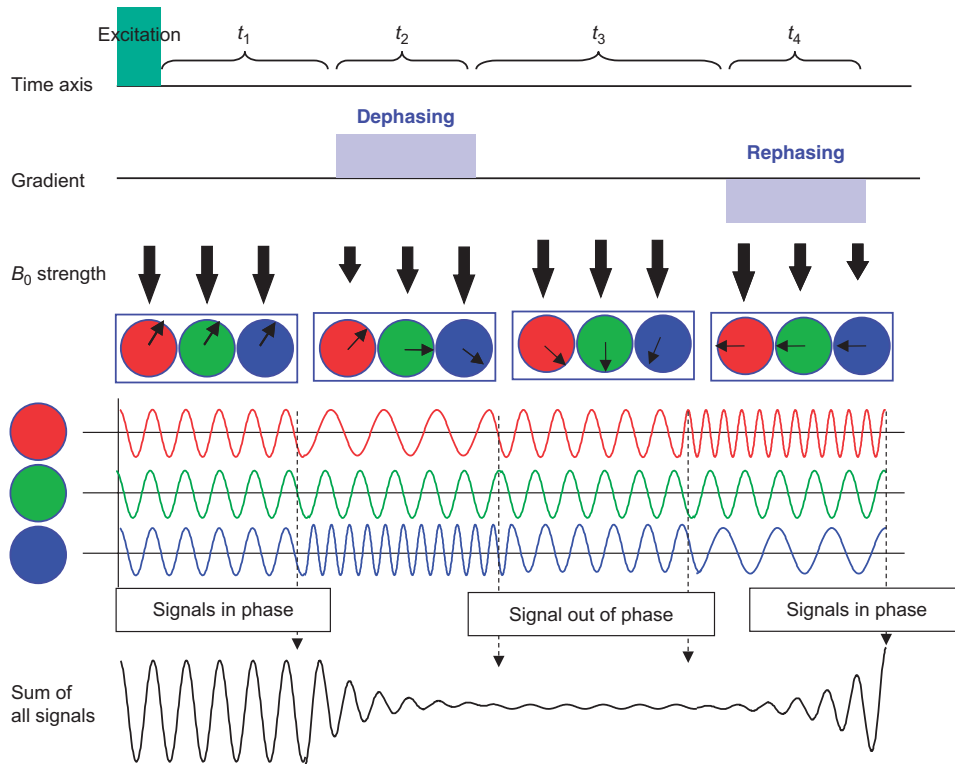


Figure 2.3: Phase alignment of water molecules in the absence of diffusion maximizes the measured signal. Credit [51]

It turns out that measurement of diffusion can be integrated into the MRI infrastructure fairly easily by introducing small modifications to the pulse sequence. The most popular adaptation is known as the Pulse Gradient Spin Echo (PGSE), introduced by Stejskal and Tanner in 1965 [65]. This can be done by introducing a pair of dephasing and rephasing gradient pulses that has the effect of desynchronizing and resynchronizing the spins of

water molecules respectively. To prevent confusion with the imaging gradients previously described to encode spatial information in conventional MRI, the dephasing/rephasing gradient pairs are sometimes collectively referred as diffusion encoded gradients, or *diffusion gradients* for short. Figure 2.3 shows the effect of the dephasing and rephasing gradients on the water molecules when no diffusion is present, as well as a simplified version of how the diffusion gradient pulses are integrated into the pulse sequence. Prior to application of the dephasing gradient, all water molecules are in phase as shown by the common alignment of the arrows. This is no longer true after the dephasing gradient is applied, and the phase misalignment persists even after the dephasing gradient is turned off. Notice how there is no apparent activity during t_3 , where neither of the diffusion gradients are activated. Assuming this is true, application of the rephasing gradient will realign the spin of the water molecules and the measured MR signal is the same as prior to dephasing. However, in the presence of diffusion, water molecules will be experiencing displacement during t_3 , which has the effect of permuting the order of the water molecules in Figure 2.3. As a result, the rephasing gradient will be unable to re-align the spins of the water molecules into a common orientation. As the spins are not oriented coherently, the signal picked up by the MRI scanner will be attenuated, with the degree of attenuation depending on the degree of diffusion present along the direction specified by the diffusion gradients. This gives rise to a MR image, with noticeable areas of reduced intensity in the presence of diffusion.

Repetition of this experiment using diffusion gradients varying in strength and direction forms a set of images collectively referred to as Diffusion Weighted Images (DWI). Note that signal attenuation in DWIs are not the result of water molecule diffusion alone. To isolate the effects of diffusion, dMRI also requires a conventional T_2 scan without applying diffusion encoded gradients. When normalizing the DWIs by the T_2 scan, the signal attenuation due to diffusion of water molecules along the direction of the diffusion encoded gradients is isolated. Because the directionality of the dephasing gradient pulse characterizes diffusion along that particular direction only, we can see the necessity of obtaining diffusion profiles along different directions to obtain a good understanding of water molecule diffusion behaviour. The *q-space* Ω_q is commonly used to denote the space of all possible diffusion gradients. The main focus of dMRI is to sample the q-space in an optimal configuration such that a balance is struck between acquisition time and accurate recovery of the diffusion profile in the subject. This will be covered in more detail in the next chapter where different variants of dMRI are discussed.




Figure 2.4: Example of how scan-dependent parameters are incorporated into the PGSE sequence. Credit [39].

2.7 Diffusion parameters and the diffusion equation

We have seen how applying the dephasing/rephasing gradient pairs allows one to measure the extent of water diffusion in the subject. Recalling that diffusion is quantified by parameters introduced in (2.10), one would naturally wonder how the parameters are reflected in the PGSE pulse sequence. Unsurprisingly, it turns out that the diffusion encoded gradients play a large role in determining the parameters; namely the magnitude of the gradient G along with the duration δ for which it is applied. Recall how the imaging gradients of the scanner cause phase perturbation proportional to the integral of the gradient pulse (2.6). Diffusion encoded gradients behave in much the same way in the sense that increasing G or δ will cause a greater extent of initial dephasing. The remaining parameter Δ depends on the time interval between the application of the dephasing and rephasing gradients. Based on the definition of the mixing time Δ , we can infer that a short Δ causes only fast diffusion processes to occur, while having a longer Δ allows more slower diffusion processes to manifest themselves. Thus, it can be seen that increased signal attenuation will occur when 1) greater initial spin dephasing is applied, and 2) more spatial displacement of molecules happens due to an increased number of ongoing diffusion processes. A rough illustration showing the diffusion parameters applied to the PGSE pulse sequence can be seen in Figure 2.4. From the above, we see that the PGSE pulse sequence can be tuned to measure different diffusion regimes. Specifically, it was shown in [65] the signal attenuation due to diffusion can be expressed as

$$\frac{S}{S_0} = e^{-\gamma^2 \delta^2 G^2 (\Delta - \frac{\delta}{3}) D} \quad (2.11)$$

where S and S_0 represent the diffusion weighted signal and the signal acquired without applying diffusion gradients respectively. D here is known as the Apparent Diffusion Coefficient which is used to highlight the fact that diffusion measurements represent the *observed* displacement over the mixing time; how this value is evaluated will be explained in greater detail in the upcoming DTI section. For the sake of notational convenience, the scanner acquisition parameters are typically consolidated into a single parameter $b = \gamma^2 \delta^2 G^2 (\Delta - \frac{\delta}{3})$ [7]. This was widely adopted due to the simplified representation, and also emphasizes that these scanner based parameters typically have little bearing on subsequent analysis once the DWIs have been acquired. Thus, (2.11) simply becomes $\frac{S}{S_0} = e^{-bD}$.

Chapter 3

Fundamental dMRI models

Owing to each DWI being only capable of determining water diffusion of a given intensity along the direction specified by the diffusion gradient, it only yields a tiny bit of the diffusion profile. Thus, to probe different diffusion regimes occurring along various spatial directions, the PGSE experiment needs to be repeated for a large number of diffusion gradients. This cannot always be accomplished for practical reasons. To overcome this barrier, dMRI researchers have come up with different models and methods allowing one to get a good grasp on diffusion behaviour while keeping the acquisition time feasible. In this thesis, we will cover the more common models. Before that however, we will briefly discuss various studies conducted by dMRI researchers that have contributed to a greater understanding of the human brain.

Although the procedure for measuring diffusion of water molecules using existing MRI infrastructure has been shown in the previous chapter, at this point it is still unclear what conclusions can be drawn based on these measurements. Towards that end, there have been a multitude of studies showing how dMRI data can be used in classification of the morphology of dendrites and axons in the human brain. Early studies observed that water molecules in the brain have a strong propensity to diffuse along the length of an axon as opposed to the perpendicular direction, leading to the myelin hypothesis [43, 66], which states that the myelin sheath enclosing neural axonal tracts is extremely effective at hindering water diffusion across the sheath. Empirical evidence through fiber tractography experiments has suggested similar fiber tract modelling relative to post-mortem analysis. In 2005, Assaf et al. [3] proposed the Composite Hindered and Restricted Model of Diffusion (CHARMED) which states that the diffusion profile in white matter can be classified as a combination of hindered diffusion exhibiting Gaussian diffusion, and restricted diffusion for intra-axonal volumes for which non-Gaussian diffusion behaviour is observed. While this

model has been positively noted by accounting for the presence of restricted and hindered diffusion behaviour through the notion of volume fractions, its greatest limitation stems from inability to accurately represent bending and fanning distribution of axons [81].

A more recent study by Zhang et al. [81] attempts to present a more complete depiction of diffusion behaviour by extending the classification of brain tissue to grey matter, which is known to contain a high degree of axonal orientation dispersion. The proposed Neurite Orientation Dispersion and Density Imaging (NODDI) scheme argues that the observed diffusion signal can be expressed as a sum of individual contributions from intra-cellular, extra-cellular and CSF compartments. Specifically, the intra-cellular compartment refers to the area enclosed by dendrites and axons, which are modelled by a series of cylinders with zero radius to emphasize the observation of unhindered diffusion along an axonal length as opposed to highly restricted diffusion in the perpendicular direction [81]. On the other hand, extra-cellular compartments referred to the area surrounding dendrites and axons, which are represented by hindered yet unrestricted diffusion behaviour. Finally, the diffusion within CSF compartments were characterised by unhindered and unrestricted diffusion, modelled by a symmetric Gaussian function and generally referred to as *isotropic* diffusion.

Owing to the ability of dMRI to differentiate between healthy and diseased patients by noting discrepancies in the pattern of cerebral diffusion, it has become an invaluable tool for clinicians and researchers towards diagnosing diseases such as Multiple Sclerosis [33]. The continuing development of more sophisticated characterization of diffusion behaviour in brain matter has lead to an interest of different dMRI models to more effectively illustrate different aspects of diffusion behaviour, which is the primary focus of this chapter.

3.1 Diffusion Tensor Imaging

The preceding discussion on the characterisation of diffusion behaviour in the human brain encompassed the notion of hindered and restricted diffusion initially proposed by Assaf [3]. Whereas restricted diffusion described a non-Gaussian displacement distribution observed in neuronal tissues under the presence of strong diffusion gradients [53], hindered diffusion can be modelled by a Gaussian displacement pattern. Usage of the terms hindered and restricted suggests that the displacement probability is not equivalent in every direction. Indeed, physical obstruction can reduce the extent of diffusion in any given direction. When this occurs, the overall diffusion process is classified as being *anisotropic*. On the other hand, isotropic diffusion behaviour refers to the case where displacement is equally probable in all directions.

In the early days of dMRI where sophisticated frameworks for describing diffusion processes in the brain did not exist, it was critical that diffusion behaviour could be inferred from a limited number of DWIs. Towards that end, one of the first widespread techniques for working with dMRI data was known as Diffusion Tensor Imaging (DTI). It assumes that the diffusion displacement probability followed a Gaussian distribution, which today we know accounts for the unhindered diffusion but is incapable of properly describing restricted diffusion. Despite this shortcoming, DTI is still the most common protocol used in clinical studies. As its namesake implies, DTI uses a symmetric positive definite 2nd order tensor \mathbf{D} , which made it possible to characterize anisotropic diffusion. Note that the use of the tensor does not exclude the possibility of isotropic diffusion, in which \mathbf{D} becomes proportional to a 3x3 identity matrix.

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix} \quad (3.1)$$

Each of the diagonal terms within the tensor can be interpreted as the diffusivity along the directions aligned with standard Cartesian axes x, y, z . Given that, it is important not to jump to the conclusion that cross terms indicate the diffusivity along the direction containing equal weighting between two axes. For example, while D_{yy} represents the diffusivity along the y-axis, D_{xz} does not correspond to diffusivity along the unit vector $\frac{1}{\sqrt{2}}(1, 0, 1)$. Rather, the cross terms denote the correlation between the diffusivity along each Cartesian axes. Consequentially, we can see that the tensor D is in fact symmetric and has precisely 6 unknown coefficients $\{D_{xx}, D_{yy}, D_{zz}, D_{xz}, D_{yz}, D_{xy}\}$.

To account for differences in diffusion measurements between different directions in the anisotropic case, the ADC is a scalar value describing the diffusivity along a given direction and has units $[\text{s}/\text{mm}^2]$. Different dMRI models provide their own approach for evaluating the ADC. In the case of DTI, the ADC is evaluated using the aforementioned diffusion tensor \mathbf{D} , by observing the signal attenuation due to diffusion can be extended from (2.11) into the following equation

$$\frac{S}{S_0} = e^{-b\mathbf{u}^T\mathbf{D}\mathbf{u}} \quad (3.2)$$

where \mathbf{u} is a unit vector denoting the desired diffusion direction, and $b = \gamma^2\delta^2G^2(\Delta - \frac{\delta}{3})$ is the b-value containing the scanner parameters γ, δ, Δ as introduced prior. Since the diffusion tensor has a total of 6 unknowns, a minimum of 6 DWIs is required to fully solve for \mathbf{D} . Furthermore, an additional MRI scan in the absence of diffusion encoding

is required for normalization purposes. In practice DTI experiments use 20 to 30 evenly distributed directions over the unit sphere [4], after which a least squares fit could be used to minimize the effect of measurement noises. An extensive review of such methods can be found in [42].

Although the diffusion tensor is capable of characterizing the diffusion profile in the DTI model, by itself it does not provide an intuitive metric which would allow one to relate the observed ADCs to the physiological properties of the underlying cerebral matter. For this reason, researchers have instead looked at the three eigenvalues $\{\lambda_i\}_{i=1}^3$ of the diffusion tensor and developed several useful metrics. These metrics include Mean Diffusivity (MD), Fractional Anisotropy (FA) and Volume Ratio (VR) [51] which are defined as follows.

$$\begin{aligned} \text{MD} &= \bar{\lambda} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \\ \text{FA} &= \sqrt{\frac{3}{2}} \sqrt{\frac{(\lambda_1 - \bar{\lambda})^2 + (\lambda_2 - \bar{\lambda})^2 + (\lambda_3 - \bar{\lambda})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \\ \text{VR} &= \frac{\lambda_1 \lambda_2 \lambda_3}{\bar{\lambda}^3} \end{aligned}$$

Of particular interest to clinicians is the FA metric ($\text{FA} \in [0, 1]$) because it allows one to infer the degree of anisotropy at a given location. Under isotropic conditions, the eigenvalues are equivalent and the resulting FA approaches 0. Conversely, when it is strongly anisotropic the eigenvalue corresponding to that direction will be very large relative to the other two. In this case, the FA metric yields a value close to 1.

Unfortunately, it has been demonstrated that the DTI model is overly simplistic for modelling neural fiber tracts due to its assumption of unrestricted diffusion, which is only valid for locations with a single dominant fiber orientation [70, 2]. In the presence of multiple fibers with different orientations (e.g., crossing, kissing configurations), DTI is incapable of providing a satisfactory representation. For this reason, DTI is often used to view fiber tracts in the corpus callosum owing to the relatively aligned fiber bundles [9]. For other parts of the brain with complex fiber structures such as projections of callosal fibers to the cortex [26], we must resort to different methods of classifying diffusion behaviour.

3.2 Diffusion Spectrum Imaging

We have seen in the previous section that characterizing the diffusion behaviour with a Gaussian distribution is valid for areas with a single dominant fiber orientation, but is rendered inaccurate for more complicated geometries with multiple fiber orientations. One major reason for this is the assumption of an implicit model with the mindset of simplifying the underlying diffusion process. In this section, we will describe Diffusion Spectrum Imaging (DSI), another perspective proposed by Wedeen in 2005 [74] on recovering water diffusion behaviour without imposing a model on the diffusion profile. The philosophy behind DSI is best understood by revisiting the fundamental physics surrounding diffusion.

Central to DSI is the notion of diffusion propagator $P(\mathbf{x}_1, \mathbf{x}_2, \Delta)$ and the wave-vector \mathbf{q} . The diffusion propagator describes the probability a water molecule at initial location \mathbf{x}_1 undergoes a spatial displacement to \mathbf{x}_2 during the mixing time Δ . As mentioned previously it is more useful to use an aggregated measure of the diffusion behaviour as opposed to classifying each random walk separately. Towards that end, the ensemble average propagator (EAP) is classified as the total probability of spatial displacement to \mathbf{x}_2 from all possible initial spin densities $\rho(\mathbf{x}_1)$ during a given mixing time Δ . The corresponding formula is thus

$$P(\mathbf{x}, \Delta) = \int_{\mathbb{R}^3} \rho(\mathbf{x}_1) P(\mathbf{x}_1, \mathbf{x}_2, \Delta) d\mathbf{x}_1 \quad (3.3)$$

where $\mathbf{x} = \mathbf{x}_2 - \mathbf{x}_1$ represents the displacement vector. On the other hand the wave vector \mathbf{q} is, similar to the aforementioned b-value, an alternate interpretation of the scanner parameters G , δ and is given by

$$\mathbf{q} = \frac{\gamma \delta G \mathbf{u}}{2\pi} \quad (3.4)$$

where \mathbf{u} is a unit vector denoting the direction of diffusion encoding gradient. A quick warning from this simplified definition of the wave vector stems from the diffusion gradient pulse used in PGSE, which is applied for a very short duration compared to the mixing time (e.g., $\delta \ll \Delta$).

Using the formalism of the EAP and wave vector, it is possible to establish a connection between these and the normalized diffusion signal $E(\mathbf{q})$. Since a phase change is induced on the spin density during the transition from \mathbf{x}_1 to \mathbf{x}_2 , the resulting average spin density can be expressed as $P(\mathbf{x}, \Delta) e^{2\pi i \mathbf{q} \cdot \mathbf{x}}$. Integrating over all possible displacement vectors to consider all initial spatial locations, we arrive at the formula

$$E(\mathbf{q}) = \int_{\mathbb{R}^3} P(\mathbf{x}, \Delta) e^{2\pi i \mathbf{q} \cdot \mathbf{x}} d\mathbf{x} \quad (3.5)$$

from which we can deduce there exists a Fourier relationship between the normalized diffusion signal $E(\mathbf{q})$ and the EAP. Through the EAP, we would now be able to completely classify the diffusion behaviour. Note that DSI is completely independent of a model for the underlying diffusion process. However with this advantage comes a severe downfall for utilizing DSI. For DSI, the q-space must be sampled very densely in order for the EAP to be correctly reconstructed from the diffusion signal [4]. As noted in the MRI section, acquisition time is a very practical requirement owing to patient motion for *in vivo* analysis. A typical DSI experiment requires anywhere between 200 and 512 brain volume scans [75], which renders it infeasible with the existing methodology and technology.

3.3 High Angular Resolution Diffusion Imaging

So far we have seen how DTI and DSI represent opposite ends of the spectrum in terms of their ability to represent diffusion behaviour and acquisition times. Unfortunately this makes neither of them a suitable candidate for clinical use. To alleviate this, Tuch [72] proposed a new method of sampling the q-space named High Angular Resolution Diffusion Imaging (HARDI). The main innovation of HARDI stems from its goal of capturing angular information at the expense of magnitude. As such, the magnitude of the diffusion gradients is kept constant for all q-space samples. Collectively, the q-space samples in HARDI take on the appearance of a discretized sphere. To reduce the effects of correlation between data points, a uniform distribution of q-space samples is desired. Furthermore, since the diffusion process is assumed to be symmetric sometimes only the upper hemisphere is used.

Clearly, one would question whether capturing diffusivity within a narrow regime would present a good understanding of diffusion behaviour. It turns out that high angular resolution results in the ability to resolve the orientation of fibers, whereas radial resolution is more effective at capturing details at a very microscopic scale but is severely hindered by its small FOV [4]. Thus, while fine radial resolution is important for inferring characteristics such as axon diameter, it is less important for applications of dMRI such as understanding the trajectory of neural fiber tracts. Moreover, provided the b-value is chosen carefully such that a large variety of diffusion processes occur without sacrificing signal to noise ratio, the acquired data provides rich diffusion information. In some cases, several concentric spherical shells in q-space are deployed during data acquisition; this is known

as multi-shell diffusion imaging [78, 27] and provides some compromise to the problem of limited radial resolution.

For the above, a large portion of the ongoing research in fiber tractography uses HARDI data. But that is only the tip of the iceberg as HARDI simply represents the acquisition method of using diffusion gradients spread over a spherical shell. There have been many proposed methods of handling the data ever since Tuch introduced the concept of HARDI to the dMRI community, some of which borrow fundamental ideas from DTI and DSI. The next section will cover some of the more relevant approaches in greater detail.

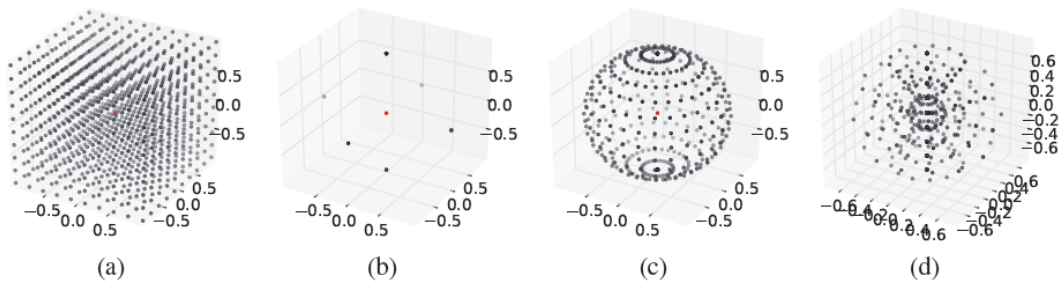


Figure 3.1: Figure illustrating how the q-space is sampled for (a) DSI (b) DTI (c) HARDI (d) Multi-shell Diffusion Imaging. The q-space origin is highlighted in red. Credit [4].

In conclusion, this chapter looked at different methods of modelling diffusion MRI data. From the simple tensor model of DTI to the EAP formalism used in DSI, the advantages and downfalls of each has been laid out in detail. Figure 3.1 provides an illustrative view of how the q-space is sampled; the dense rectangular lattice sampling of DSI, the 6 samples required of DTI, and the spherical shell sampling present in HARDI. The rightmost figure is more difficult to discern, but upon closer observation one can observe the multiple concentric spherical shells signifying a multi-shell model.

Chapter 4

Working with HARDI data

Over the years there have been different philosophies towards handling HARDI data. The first type is model based, namely by modelling each voxel by a small number of individual diffusion compartments. The other type strives to retain the EAP formalism presented by DSI while maintaining the modest acquisition time required of HARDI data.

4.1 Gaussian Mixture Model

One of the downfalls of DTI was its inability to represent areas with more than one dominant diffusion orientation. To combat this shortcoming while retaining the simplicity of a Gaussian diffusion profile, the Gaussian Mixture Model (GMM) was initially proposed by Tuch [70]. Its formalism relies on two key assumptions about the diffusion process which are fairly general and adhered to in practice [70]. Firstly, heterogeneity of the diffusion compartments within each voxel is restricted to several distinct homogeneous regions. Secondly, the aforementioned distinct homogeneous regions are separated by a large enough distance guaranteeing only slow exchange happening between regions. In conjunction, these criteria serve to minimize the dependencies between each individual diffusion compartment. Under these assumptions about the diffusion process, every spatial location can have its diffusion profile modelled by a linear combination of Gaussian functions, each following the tensor model introduced in DTI and is used to represent a single diffusion compartment. Recall that the diffusion tensor \mathbf{D} from DTI was used to characterize Gaussian diffusion through

$$\frac{S}{S_0} = e^{-\mathbf{b}\mathbf{u}^T\mathbf{D}\mathbf{u}} \quad (4.1)$$

where as before $b = \gamma^2 \delta^2 G^2(\Delta - \frac{\delta}{3})$ represents the diffusion regime of interest and \mathbf{u} defines the unit vector in the direction of the applied diffusion encoding gradient. Tuch's idea was to fit an individual tensor within each of the N diffusion compartments, giving rise to the set of tensors $\{\mathbf{D}_i\}_{i=1}^N$. To facilitate numerical estimation of the tensors $\{\mathbf{D}_i\}_{i=1}^N$, Tuch proposed to parametrise them using Euler angles. Furthermore, the eigenvalues of the tensors were predefined to avoid individual tensors from assuming oblate forms. In this manner, the i^{th} diffusion compartment was modelled as $\psi_i = e^{-b\mathbf{u}^T \mathbf{D}_i \mathbf{u}}$

Generally speaking, the contribution of each individual diffusion compartment will not be equivalent and must be noted accordingly. Towards that end, the term *volume fraction* $\{f_i\}_{i=1}^N$ was introduced to indicate the degree of presence for each compartment at the given spatial location. Noting that these volume fractions represent the proportion of the voxel occupied by each compartment, we have $f_i \in [0, 1]$, $i = 1, 2 \dots N$. Since the magnitude of the diffusion signal is tightly coupled with the strength of diffusion in the sample, the volume fractions are normalized such that $\sum f_i = 1$. Through the formulation of each distinct Gaussian function and the corresponding volume fractions, the theory of GMM states that the diffusion signal at each voxel is formed from a linear combination of N diffusion compartments in the following manner

$$E(\hat{\mathbf{q}}) = \sum_{i=1}^N f_i \psi_i = \sum_{i=1}^N f_i e^{-b\mathbf{u}^T \mathbf{D}_i \mathbf{u}} \quad (4.2)$$

It should be noted that in practice, the absence of *a priori* information on the number of diffusion compartments within any given voxel dictates that N must also be estimated. While in theory the choice of N can be any natural number, Tuch's study concluded that [70] values of $N > 2$ lead to stability issues during the model fitting. The procedure for estimating unknown parameters at each voxel required the minimizing of the error function χ between the estimated model and observed diffusion signal $E(\mathbf{q}_k)$ for each diffusion gradient

$$\chi = \sum_k \left(\sum_{i=1}^N f_i e^{-b\mathbf{u}_k^T \mathbf{D}_i \mathbf{u}_k} - E(\mathbf{q}_k) \right)^2 \quad (4.3)$$

which was subsequently solved using standard methods of numerical optimization. Unfortunately, it was reported that this minimization problem was plagued by the existence of local minima, which had to be accounted for by utilizing multiple restarts. Furthermore, it can be seen that the minimization of χ is non-linear in nature.

While the numerical issues as well as inability of modelling using a larger number of diffusion compartments poses a severe limitation, nevertheless GMM resulted in a paradigm shift in dMRI because it showed that DTI was ineffective at representing diffusion behaviour at the higher b-values ($b > 1000 \text{ s/mm}^2$) conducted in this study. This revelation led dMRI researchers to build upon the theory of GMM towards the development of more advanced models, one of which is discussed in the next section.

4.2 Diffusion Basis Functions

The main motivation of the Diffusion Basis Functions Decomposition technique proposed by Alonso et al. [57] was to overcome the numerical issues surrounding the original GMM approach. Towards that end, they proposed to use a linear combination of Gaussian functions illustrating different primary diffusion orientations to represent the observed diffusion signal. Similar to GMM, these Gaussian functions were represented using a diffusion tensor \mathbf{D} with pre-defined eigenvalues best describing the diffusivity of the sample in the presence of anisotropic diffusion. In their paper this set of Gaussian functions $\psi = \{\psi_i\}_{i=1}^N$ was referred to as the set of *diffusion basis functions* (DBFs).

To represent the principal orientation of each DBF ψ_i , the initial tensor \mathbf{D} must be adjusted accordingly to formulate \mathbf{D}_i . For the i^{th} DBF, this can be facilitated by applying a rotation matrix \mathbf{R}_i to the initial tensor to construct the modified tensor $\mathbf{D}_i = \mathbf{R}_i \mathbf{D} \mathbf{R}_i^T$. Replacing \mathbf{D} by \mathbf{D}_i , it follows that the function for the i^{th} DBF can be expressed as $\psi_i(b, \mathbf{u}) = e^{-b\mathbf{u}^T \mathbf{D}_i \mathbf{u}}$. For the sake of notational convenience, the dependence of the DBFs on b, \mathbf{u} is sometimes omitted but understood to exist implicitly. Using the notion of DBFs and the assumptions of the underlying diffusion process from GMM (slow exchange between compartments), it was stated that the observed diffusion signal could be represented as

$$\frac{S(b, \mathbf{u})}{S_0} = \sum_{i=1}^N f_i \psi_i(b, \mathbf{u}) \quad (4.4)$$

where S_0 denotes the MR signal acquired without diffusion encoding and used for normalization purposes. The primary innovation from DBF decomposition stems from using a pre-computed set of N tensors, as opposed to estimating the tensors in conjunction with the volume fractions and N as in GMM. Thus, only the volume fractions $\{f_i\}_{i=1}^N$ corresponding to each DBF requires estimation.

Before concluding the discussion, a note on the practical implementation of DBF decomposition is prudent. The first concern involves the choice of orientation for each of the

N DBFs. Under general circumstances, it is assumed that prior information about the underlying diffusion process is unavailable. In that case, the set of orientations should be as uniformly spaced out to provide optimal coverage. Different approaches such as modelling via the Thomson problem and recursive tessellation of platonic solids have been adopted by dMRI researchers [38, 29, 70].

4.3 Q ball Imaging

Recall the discussion of DSI in the previous chapter, which characterises self-diffusion using the EAP formalism without reliance on the Gaussian diffusion model. It turns out that this idea can be adapted to data acquired through HARDI. Since the primary motivation of HARDI is to represent the diffusion profile at high angular resolution, one way of applying this concept to the diffusion model is to look at how the diffusion profile is distributed along different directions on the unit sphere. For example, a uniform spherical distribution would indicate isotropic diffusion since diffusion among all directions is equally likely. On the other hand, anisotropic diffusion could be represented by a symmetric spherical distribution consisting of two identically shaped ellipsoids fused together at the origin. This method of modelling the diffusion profile as a spherical probability distribution function is known as an Orientation Distribution Function (ODF). To see how EAP can be linked to ODF, consider how it represents the probability of spatial displacement to any given location. In 2004, Tuch introduced the ODF $\psi(\mathbf{u})$ as a function of the EAP by integrating out the radial component, effectively creating a marginal distribution as follows [69]

$$\psi(\mathbf{u}) = \frac{1}{Z} \int_0^\infty P(r\mathbf{u})dr \quad (4.5)$$

where \mathbf{u} indicates an arbitrary unit vector indicating direction, and Z is simply a unitless normalization constant. Note this approach assumes that the EAP at different radii are weighted equivalently. However this has been pointed out by Aganj et al. [1] that it does not consider the fact that the differential spherical volume is greater at higher radii. Towards that end, it has been proposed that the ODF $\psi(\mathbf{u})$ is instead computed as

$$\psi(\mathbf{u}) = \int_0^\infty P(r\mathbf{u})r^2dr \quad (4.6)$$

where $\psi(\mathbf{u})$ is inherently normalized. In practice, it was noted by Tuch that performing radial projection on Cartesian sampled data from DSI would inevitably lead to Cartesian

artifacts during mapping from Cartesian coordinates to spherical coordinates [69]. Needless to say, this problem must be solved before attempting ODF construction from dMRI data. Tuch tackled this problem through two major contributions; firstly by introducing the idea of collecting dMRI data over a spherical shell, thereby avoiding the aforementioned Cartesian artifacts while at the same time being able to quantify the degree of angular resolution of the dMRI data. Tuch’s other major contribution was by showing that the diffusion ODF could be constructed from dMRI data collected in this manner by using the Funk Radon Transform (FRT) [71], a mathematical tool proposed by Paul Funk in 1915 [31]. It is an operator \mathcal{R} taking a spherical function as input and outputting a spherical function in return. Let \mathbf{v} be a unit vector representing the direction of an arbitrary diffusion gradient in a dMRI acquisition. Then, the FRT of a spherical function defined at \mathbf{v} , $s(\mathbf{v})$, is evaluated by taking the integral of s over all vectors orthogonal to \mathbf{v} . Mathematically this can be expressed as

$$\mathcal{R}[s(\mathbf{v})] = \int_{\mathbf{u} \perp \mathbf{v}, \mathbf{u} \in \mathbb{S}^2} s(\mathbf{u}) d\mathbf{u} = \int_{\mathbf{u} \in \mathbb{S}^2} s(\mathbf{u}) \delta(\mathbf{u}^T \mathbf{v}) d\mathbf{u} \quad (4.7)$$

where $\delta(\cdot)$ is the well known Dirac delta function. If we assume \mathbf{v} to be situated at the North Pole of the sphere, it turns out that the integral over \mathbb{S}^2 in (4.7) is defined over the set of vectors situated on the equator; this is also referred to as the great circle of \mathbf{v} . One key point regarding this implementation of the FRT is that it operates on functions defined in the \mathbb{S}^2 space. Since the diffusion gradients of q-space belongs to \mathbb{R}^3 as opposed to \mathbb{S}^2 , it is not directly applicable to HARDI. Fortunately, this problem can be resolved fairly easily by restricting the magnitude of the vectors such that a spherical shell of radius q' is formed in the q-space. With this in hand, the FRT can be extended to \mathbb{R}^3 without compromising the result in the following manner

$$\mathcal{R}[f(\mathbf{v})] = \int_{\mathbf{x} \in \mathbb{R}^3} f(\mathbf{x}) \delta(\mathbf{v}^T \mathbf{x}) \delta(\|\mathbf{x}\|_2 - q') d\mathbf{x} \quad (4.8)$$

Assuming that HARDI data over the diffusion gradients in q-space corresponded to a spherical function, Tuch’s proposal was to apply the FRT directly to the HARDI data $E(\mathbf{q})$ to estimate the ODF $\psi(\mathbf{u})$ for a given diffusion gradient *direction* \mathbf{u} , resulting in

$$\psi(\mathbf{u}) \approx \frac{1}{Z} \mathcal{R}[E(\mathbf{q})] \quad (4.9)$$

where \mathbf{q} is the wave vector as introduced in the DSI section. This method of constructing the ODF is referred to as *Q-ball Imaging* or QBI, owing to the notable geometry of a spherical shell present in q-space. There are several interesting points behind the deceptively

simple equation. Firstly, that the FRT does not inherently normalize the data, however for the ODF to be interpreted as a probability distribution function it is necessary to introduce the normalization constant Z . Another point of interest is that (4.9) represents a mere approximation of the true ODF. Last but not least, observe that the ODF expression on the left is only direction dependent, whereas the wave vector \mathbf{q} has magnitude information as well. This notation emphasizes that the ODF is formulated such that it is independent of magnitude. While results from HARDI data has shown that Tuch’s QBI proposal is successful at recovering the diffusion orientation, it has been noted that the approximation has some undesirable behaviour. Notably, Barnett [6] argues that usage of the FRT leads in the diffusion propagator being weighted by a zeroth order Bessel Function of the 1st kind J_0 and other terms as opposed to the desired Dirac delta function in the integration

$$\psi(\mathbf{u}) = 2\pi q' \int \int \int_{r=0}^{\infty} P(r, \theta, z) r J_0(2\pi q' r) dr d\theta dz \quad (4.10)$$

where $P(r, \theta, z)$ is the cylindrical coordinate representation of the EAP, and as before q' refers to the radius of the spherical shell corresponding to the HARDI data. The principal reason behind Barnett’s argument was that $xJ_0(x)$ grows in an envelope proportional to the square root of the parameter $x = 2\pi q' r$, which obviously did not resemble the behaviour of a Dirac Delta. Nevertheless, the success of the QBI approach for working with HARDI data has lead dMRI researchers to pursue more computationally effective methods in this area, which will be detailed in the next section.

4.4 QBI with Spherical Harmonics

The implementation of QBI using the FRT is subject to several problems in practice. Notably, the FRT involves integration of the HARDI signal alongside the great circle of the sphere. This is problematic from a computational efficiency perspective because the q-space sampling can never be dense enough to provide a continuum of data. In that case, spherical interpolation must be applied between the existing q-space points to accommodate estimation of the spherical function at points where data is unavailable. To alleviate the computational burden of computing the FRT, another approach allows an easier method of evaluating the spherical integral of FRT. This relies on the theory of Spherical Harmonics (SH), which will be introduced shortly.

A function $F(x, y, z)$ in \mathbb{R}^3 is defined to be harmonic if it satisfies the Laplace equation,

which is to say

$$\Delta F := \frac{\partial^2 F}{\partial x^2} + \frac{\partial^2 F}{\partial y^2} + \frac{\partial^2 F}{\partial z^2} = 0 \quad (4.11)$$

Since it is more convenient to deal with spherical coordinates when working with spherical functions such as the ODF, the Laplace equation can be expressed in spherical coordinates with the standard change of variables

$$\frac{1}{q^2} \left[\frac{\partial}{\partial q} \left(q^2 \frac{\partial}{\partial q} \right) + \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2}{\partial \phi^2} \right] F = 0, \quad (4.12)$$

$$q = \sqrt{x^2 + y^2 + z^2}$$

$$\theta = \cos^{-1} \left(\frac{z}{\sqrt{x^2 + y^2 + z^2}} \right)$$

$$\phi = \tan^{-1} \left(\frac{y}{x} \right)$$

where $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. Since HARDI is mainly interested in the angular portion of the data, the angular part of the solution to the above is of more practical use in this context. With that in mind, a slight abuse of notation leads to the term spherical harmonics referring only to the angular solution of (4.12). Letting $l \in \mathbb{N}$ and $m \in [-l, l], m \in \mathbb{Z}$, this definition of the spherical harmonics $Y_l^m(\theta, \phi)$ take the form

$$Y_l^m(\theta, \phi) = (-1)^m \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\phi} \quad (4.13)$$

where P_l^m are the associated Legendre polynomials and are given by

$$P_l^m(x) = \frac{(1-x^2)^{m/2}}{2^l l!} \frac{d^{l+m}}{dx^{l+m}} (x^2-1)^l \quad (4.14)$$

When the parameter of the associated Legendre polynomial P_l^m is a trigonometric function (e.g., $x = \cos \theta$ in our case), it turns out that the result is a polynomial in terms of the trigonometric function.

Since the number of q-space samples made in HARDI data directly determines the number of coefficients we can fit for the spherical harmonics basis functions, it is imperative that we make efficient use of them. Towards that end, several assumptions are made

about the HARDI signal [30]. Firstly, the corresponding ODFs are symmetric on the unit sphere, and therefore odd spherical harmonic coefficients would evaluate to zero. Secondly, collected HARDI data are real valued. These assumptions allow us to reformulate the spherical harmonics in HARDI y_l^m which uses only the real and symmetric subset of SH functions, in the following manner [32].

$$y_l^m = \begin{cases} \frac{1}{\sqrt{2}}(Y_l^m + (-1)^m Y_l^{-m}) & 0 < m \leq l \\ Y_l^0 & m = 0 \\ \frac{i}{\sqrt{2}}((-1)^m Y_l^m - Y_l^{-m}) & -l \leq m < 0 \end{cases} \quad (4.15)$$

The first few orders of the reformulated spherical harmonics are illustrated in Figure 4.1; from this it can be inferred that higher order spherical harmonics tend to exhibit many jagged features and are not suitable for modelling ODFs which, like real life phenomena, tend to be relatively smooth in nature. As an example, the spherical harmonic expansion of HARDI signals at $b = 3000\text{s/mm}^2$ truncated to $l = 4$ provides a reasonable balance between angular resolution and accuracy; a more detailed analysis of the tradeoffs for different values of l and b can be found in here [36].




Figure 4.1: Representation of the Real and Symmetric Spherical Harmonic basis functions up to $l = 6$. Note that they functions have been normalized such that they lie in the range $[-1,1]$. Credit [4].

As noted above, the set of SH functions satisfy orthogonality, which is to say

$$\int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} y_l^m(\theta, \phi) y_{l'}^{m'}(\theta, \phi) d\phi d\theta = \frac{4\pi}{2l+1} \delta(l-l') \delta(m-m') \quad (4.16)$$

Since any square-integrable spherical function may be expressed as a linear combination of the spherical harmonics, the HARDI signal $s(\theta, \phi)$ can be approximated with a truncated spherical harmonic expansion up to order L in the following manner

$$s(\theta, \phi) \approx \sum_{l=0, l \in 2\mathbb{N}}^L \sum_{m=-l}^l c_l^m y_l^m(\theta, \phi) \quad (4.17)$$

where c_l^m are the coefficients of each spherical harmonic in the expansion and $l \in 2\mathbb{N}$ signifies that only even orders of l are considered. Henceforth the notation $l \in 2\mathbb{N}$ will be dropped for brevity in notation. Note that the expansion also has an equivalent representation as a function of the unit vector \mathbf{v} as opposed to the angles $\{\theta, \phi\}$; this will be used interchangeably to simplify the notation whenever possible. We replace the definition of the HARDI signal in the calculation of the ODF using the FRT, which leads to

$$\begin{aligned} \psi(\mathbf{v}) &= \int_{\mathbf{u} \in \mathbb{S}^2} \delta(\mathbf{u}^T \mathbf{v}) s(\mathbf{u}) d\mathbf{u} \\ &= \int_{\mathbf{u} \in \mathbb{S}^2} \delta(\mathbf{u}^T \mathbf{v}) \sum_{l=0}^L \sum_{m=-l}^l c_l^m y_l^m(\mathbf{u}) d\mathbf{u} \\ &= \sum_{l=0}^L \sum_{m=-l}^l c_l^m \int_{\mathbf{u} \in \mathbb{S}^2} \delta(\mathbf{u}^T \mathbf{v}) y_l^m(\mathbf{u}) d\mathbf{u} \end{aligned} \quad (4.18)$$

Thus, the primary difficulty in evaluating the ODF is now reduced to evaluation of the spherical integral on the last line of (4.18). It turns out this spherical integral can be very elegantly simplified using the Funk-Hecke Theorem, which states a function $f(t)$ defined on the interval $t \in [-1, 1]$ satisfies

$$\int_{\mathbf{u} \in \mathbb{S}^2} f(\mathbf{u}^T \mathbf{v}) Y_l^m(\mathbf{u}) d\mathbf{u} = 2\pi Y_l^m(\mathbf{v}) \int_{-1}^1 f(t) P_l(t) dt \quad (4.19)$$

Applying the Funk-Hecke Theorem to the above results in a trivial integral that may be evaluated using the sifting property of the Dirac Delta Function. Thus, the resulting ODF simply becomes a linear combination of spherical harmonics scaled by $2\pi P_l(0)$,

$$\psi(\mathbf{v}) = \sum_{l=0}^L \sum_{m=-l}^l 2\pi P_l(0) c_l^m y_l^m(\mathbf{v}) \quad (4.20)$$

In practice, it is rather unwieldy to retain the double summation when evaluating the ODF. By observing that for a given order L , there are precisely $R = \frac{(L+1)(L+2)}{2}$ of the real and symmetric spherical harmonic basis functions, the resulting expression can be consolidated into a single summation which can be subsequently evaluated by a matrix vector multiplication (e.g., $\psi(\mathbf{v}) = \sum_{l=0}^R 2\pi P_{j_l}(0) c_{j_l}^m y_{j_l}^m(\mathbf{v})$, where $\{j_l\}_{l=0}^R = \{0, 2, 2, 2, 2, 2, 4, 4, \dots\}$ is introduced to preserve the mapping between different the different orders of SH functions). The usage of spherical harmonics towards evaluating the ODF via the Funk Radon Transform has become very prevalent due to the speed and ease of computation.

4.5 Spherical Ridgelets

The theory of spherical ridgelets was introduced by Michailovich et al. [50] to provide an alternative representation of HARDI signals. The primary motivation behind this development can be attained by examining the formation of a diffusion signal from GMM theory, which says the diffusion signal measured in any given voxel can be classified as the summed contribution from N compartments, where the diffusion behaviour within the i^{th} compartment consists of a primary diffusion orientation \mathbf{v}_i and is modelled by an anisotropic Gaussian function. Since diffusion results in attenuation of the measured MR signal, we would expect the net signal contribution of the i^{th} compartment (towards the overall signal) to be weak along \mathbf{v}_i . On the other hand, we would see strong signal contribution along directions perpendicular to \mathbf{v}_i due to low degrees of diffusivity along those directions. By summing the signals from the N departments, we can deduce the majority of the energy from the resultant diffusion signal would be concentrated along the N great circles. Unfortunately, the spherical harmonics proposed in the previous section do not share this property and thus it was argued they are not optimal for representing HARDI signals. To tackle this problem, it was suggested that the Gauss-Weierstrass Kernel $K_{\mathbf{v}}$ defined at orientation \mathbf{v} was used to generate a function of the form

$$K_{\mathbf{v}}(\mathbf{u}) = \sum_{n=0}^{\infty} \frac{2n+1}{4\pi} e^{-\rho n(n+1)} P_n(\mathbf{u} \cdot \mathbf{v}) \quad (4.21)$$

where once again P_n denotes the associated Legendre polynomial of degree n , ρ is a scalar controlling the bandwidth of the resulting kernel, and \mathbf{u} is a unit vector indicating the spherical coordinate at which $K_{\mathbf{v}}$ is evaluated. At this point it should be noted that the dependence of \mathbf{v} and \mathbf{u} in $K_{\mathbf{v}}(\mathbf{u})$ is restricted to the quantity $\mathbf{u} \cdot \mathbf{v}$, and thus $K_{\mathbf{v}}(\mathbf{u})$ can consequentially be classified as a zonal function. Crucial to the development of spherical

ridgelets is the addition theorem of the Legendre polynomials that states

$$P_n(\mathbf{u} \cdot \mathbf{v}) = \sum_{m=-l}^l Y_l^m(\mathbf{u})Y_l^m(\mathbf{v}) \quad (4.22)$$

Unfortunately, $K_{\mathbf{v}}(\mathbf{u})$ is ill-equipped for representation of HARDI signals owing to majority of its energy being concentrated along \mathbf{v} . However, it turned out that application of the FRT to $K_{\mathbf{v}}$ provided a remedy to this problem. Towards that end, the *ridgelet generating function* (RGF) $\Phi_{\mathbf{v}}$ was constructed by integrating $K_{\mathbf{v}}$ over the great circle $\sigma(\mathbf{v})$ situated about \mathbf{v} .

$$\Phi_{\mathbf{v}}(\mathbf{u}) = \frac{1}{2\pi} \int_{\mathbf{q} \in \sigma(\mathbf{v})} K_{\mathbf{q}}(\mathbf{u}) d\eta(\mathbf{q}) \quad (4.23)$$

where $\eta(\mathbf{q})$ is used to denote the standard rotational invariant measure on \mathbb{S}^2 . It was previously seen that the Funk-Hecke Theorem (4.19) associated with spherical harmonics resulted in considerable ease towards evaluation of the spherical integral. The presence of $P_n(\mathbf{u} \cdot \mathbf{v})$ in conjunction with the addition theorem allows us to use the Funk-Hecke Theorem, resulting in

$$\Phi_{\mathbf{v}}(\mathbf{u}) = \frac{1}{2\pi} \sum_{n=0}^{\infty} \frac{2n+1}{4\pi} \lambda_n e^{-\rho n(n+1)} P_n(\mathbf{u} \cdot \mathbf{v}) \quad (4.24)$$

where $\lambda_n = \begin{cases} 2\pi(-1)^{n/2} \frac{1 \cdot 3 \cdots (n-1)}{2 \cdot 4 \cdots n} & , n \text{ is even} \\ 0 & , n \text{ is odd} \end{cases}$

It was subsequently shown that with an appropriate choice of bandwidth ρ , RGFs provided a good approximation of the diffusion signal within a GMM compartment owing to the aforementioned concentration of energy around the great circles. Therefore, RGFs presented a natural choice of generating function towards HARDI signal decomposition. Furthermore, this was done with one single carefully tuned kernel $K_{\mathbf{v}}$ as opposed to the whole collection of Spherical Harmonic functions! Along that note however, while an ideal value of ρ could be selected to represent the k^{th} diffusion compartment, it did not necessarily imply the same ρ would perform equally well for other diffusion compartments. This observation motivated the development of applying multiresolution analysis towards the proposed RGFs, thereby allowing reconstruction of the HARDI signal via scaled versions of the RGF.

Observe that the definition of the Gauss Weierstrass kernel (4.21) included the function $e^{-\rho n(n+1)}$, a discrete analogy of the continuous function $\kappa_\rho(x) = e^{-\rho x(x+1)}$. Constructing dilated versions of κ_ρ using $j \in \mathbb{N}_0$ to define the scaling factor results in

$$\kappa_j(x) = \kappa_\rho(2^{-j}x), \quad \text{where } j = 0, 1, 2 \dots \quad (4.25)$$

Substituting $\kappa_j(n)$, the discretized version of $\kappa_j(x)$, into the original definition of the Gauss Weierstrass kernel allows one to formulate the set of scale-discrete scaling functions in the following manner

$$K_j(t) = \sum_{n=0}^{\infty} \frac{2n+1}{4\pi} \kappa_j(n) P_n(t), \quad t \in [-1, 1] \quad (4.26)$$

where the parameter t replaces $\mathbf{u} \cdot \mathbf{v}$ from the original definition in efforts to simplify the notation. A particular desirable characteristic of $K_j(t)$ was its compliance to the *discrete approximation identity*, which states that a function $F \in \mathbb{L}^2(\mathbb{S}^2)$ can be represented via convolution with $K_j(t)$ at sufficiently fine scale via

$$\lim_{j \rightarrow \infty} \|F - K_j * F\| = 0 \quad \forall F \in L^2(\mathbb{S}^2) \quad (4.27)$$

$$\text{where } (H * F)(\mathbf{v}) = \int_{\mathbb{S}^2} H(\mathbf{v} \cdot \mathbf{u}) F(\mathbf{u}) d\eta(\mathbf{u})$$

In accordance with standard wavelet analysis, the spherical wavelet $W_j(t) = K_{j+1}(t) - K_j(t)$ was proposed to represent the difference function between two successive resolutions. Denoting the operation of convolution by $*$, this allowed a function $F \in \mathbb{L}^2(\mathbb{S}^2)$ to be reconstructed via

$$F = K_0 * F + \sum_{j=0}^{\infty} W_j * F \quad (4.28)$$

So far, the discussion of multiresolution analysis was applied towards scaled discretized versions of the Gauss Weierstrass kernel $K_{\mathbf{v}}$. Recalling that RGFs were related to $K_{\mathbf{v}}$ by means of the Funk Radon Transform, is it possible to extend this analysis towards RGFs as well? It turns out that the answer is yes, but the rigorous proof of this is beyond the scope of this thesis and can instead be found in the original works [50]. Proceeding under this assertion, the base RGF $\Phi_{0,\mathbf{v}}$ corresponding to orientation \mathbf{v} was defined as

$$\Phi_{0,\mathbf{v}}(\mathbf{u}) = \frac{1}{2\pi} \sum_{n=0}^{\infty} \frac{2n+1}{4\pi} \lambda_n \kappa_0(n) P_n(\mathbf{u} \cdot \mathbf{v}) \quad (4.29)$$

The next step is to build the corresponding spherical ridgelet $\psi_{j,\mathbf{v}}$ as the difference function between successive resolutions of the RGF.

$$\begin{aligned} \psi_{j,\mathbf{v}}(\mathbf{u}) &= \Phi_{j+1,\mathbf{v}}(\mathbf{u}) - \Phi_{j,\mathbf{v}}(\mathbf{u}) \\ &= \frac{1}{2\pi} \sum_{n=0}^{\infty} \frac{2n+1}{4\pi} \lambda_n[\kappa_{j+1}(n) - \kappa_j(n)] P_n(\mathbf{u} \cdot \mathbf{v}) \end{aligned} \quad (4.30)$$

As a side note, spherical ridgelets were given this name owing to their property of having majority of its energy concentrated along the great circles with respect to their defined orientation \mathbf{v} . This is as opposed to the aforementioned spherical wavelet functions, which instead have energy concentration around \mathbf{v} . To illustrate these properties, the spherical ridgelets at the first three resolution levels are presented in Figure 4.2.




Figure 4.2: Illustration of Spherical Ridgelets created using $j = 3$ resolution levels, and bandwidth $\rho = 2$.

It should be noted that the set of spherical ridgelets are not linearly independent, in the sense that they form an overcomplete frame as opposed to a basis. Thus, there is no unique representation of the HARDI signal using spherical ridgelets. Although at first glance this may seem to be a detrimental characteristic, on the contrary it allows one to specify regularizers based on assumptions about the behaviour of a typical HARDI signal. For example, the notion of sparsity in spherical ridgelet coefficients was explored in the original work [50], where it was proposed to recover the set of coefficients using the orthogonal matching pursuit algorithm [68].

In this chapter a variety of methods of working with HARDI data have been introduced. As we can see, GMM and DBF decomposition methods relies on the fundamental assumption that the proposed model is appropriate for modelling the underlying diffusion

behaviour. At the same time, they are advantageous in the sense that they offer robustness to noise over model-free approaches due to their avoidance of physically improbable solutions [57]. On the other hand, the representation of a HARDI signal using spherical ridgelets hinges on the criterion that they provide a sparse representation of HARDI signals. This representation provided a useful method of inferring primary diffusion orientations through the notion of an ODF.

Chapter 5

Resolution Enhancement in Spatial and Angular Domains

Now that different approaches to dMRI have been presented, we will look at a practical application of how they can be used to model the location and shape of neural fiber tracts (henceforth abbreviated as fibers) in the brain. This field is also known as fiber tractography, and has been performed in different studies across different dMRI techniques. For example, DTI tractography was employed to explore morphology on the corpus callosum at the mid-sagittal level [37]. Another tractography study using DSI data was conducted by Wedeen et al. [75] in efforts to provide accurate anatomical modelling of fibers in the centrum semiovale among others regions renowned for exhibiting complex fiber geometry. This thesis will focus on working with HARDI data since it provides a very appealing balance between acquisition time and angular resolution, both of which are very critical factors in fiber tractography.

Before proceeding with the discussion on available tools and methods for resolving low resolution dMRI data, let us first re-emphasize reasons for this shortcoming. Namely, why are long acquisition times particularly prevalent in dMRI as opposed to other variations of MRI such as structural MRI? Since dMRI data consists of entire brain volumes scanned K times, where K corresponds to the number of deployed diffusion gradients, the scanning time scales linearly with respect to K . In HARDI, these diffusion gradients are chosen to represent diffusion behaviour along different orientations. This is as opposed to structural MRI in which one acquisition provides the necessary information required to identify different anatomical structures as well as the presence of pathologies within a patient.

On the other hand, spatial resolution is governed by the acquisition protocol of the MRI

scanner. We have previously discussed how EPI acquires the entire k-space in the time frame of one TR. For dMRI, EPI dominates the scene as the prevalent acquisition method due to the drastically reduced scan times. At the same time, dMRI suffers from inherent limitations of EPI, namely the reduced spatial resolution. Unfortunately, spatial resolution is critical towards effective fiber tractography due to inevitable partial volume effects [2] appearing at low spatial resolution. For example, if the voxel size was sufficiently large that it encompasses tissues from both grey and white matter, resolving the fiber orientations at that spatial location would be much more difficult than a more favourable case where different tissues can be isolated into different voxels at sufficiently high spatial resolution. Thus, the primary goal of this chapter will be looking towards methods that can help alleviate the difficulties of working with low resolution dMRI data, from both spatial and angular perspectives.

5.1 Spatial Methods

There have been many proposed methods in the literature to improve the spatial resolution of dMRI images, and they can be loosely classified as being either post-processing or scanner-based techniques. Scanner based techniques typically increase spatial resolution by increasing the magnetic field, improving the performance of the gradient-switching process or sacrificing the Field-of-View (FOV). While it has been successful in several applications [34], improving the hardware is not always feasible as MRI scanner upgrades tend to be very expensive. The technique involving FOV reduction has seen success in specific applications such as imaging of the spinal cord [77], but is not suitable for fiber tractography since it gives rise to the possibility of a fiber not being captured in its entirety.

Post-processing techniques of improving spatial resolution can be further classified as being either single-frame or multi-frame. In multi-frame super resolution, a series of low resolution dMRI images representing the same DWI are acquired. The images are acquired such that each is offset by a sub-voxel with respect to one another. In theory, unique information can be collected from each low resolution image in this manner. Therefore, the low resolution images can be fused together to reconstruct a single image at higher spatial resolution. One such study conducted by Peled et al. can be found in [55]. Unfortunately, multi-frame super resolution is incapable of reducing the overall acquisition time owing to the need of performing multiple acquisitions towards the reconstruction of one high resolution image. On the other hand, single-frame super resolution hinges on the exploration of spatial dependence between any given pair of adjacent voxels to obtain a robust estimate of the underlying dMRI signal via sophisticated interpolation schemes. Examples of this

include patch-based super resolution proposed by Coupe et al. [23], track-density imaging proposed by Calamante et al. [14] among many others. A recent study conducted by Yap et al. [79] exploring the idea of fiber-driven spatial resolution enhancement. Specifically, an advanced spatial interpolation framework was used to estimate the diffusion signal in regions where a high probability of fiber connection was detected between adjacent voxels. Compared to scanner based techniques, post-processing proved to be more versatile overall due to independence of MRI scanner hardware requirements nor pre-specified methods of obtaining the data.

Yet another proposal deserving mention is the multi-slice technique, which is interesting in the sense that both scanner based and post-processing philosophies are present. Multi-slice techniques work by acquiring multiple acquisitions of the same brain volume at low resolutions, where each acquisition is performed orthogonally with respect to the other. From this series of volumes, a MAP probabilistic framework [61] is used to estimate the volume at high spatial resolutions. While not requiring expensive hardware upgrades, it does require data being collected in a specific method that has not been widely adopted by the dMRI community. Furthermore, similar to multi-frame techniques there is little to no reduction of scan time due to the need of multiple low resolution acquisitions.

5.2 Angular Methods

Unlike the spatial domain where data can be brought into a higher resolution with relative ease, it is more difficult to directly increase the angular resolution since an effective form of spherical interpolation would be required. While it is known that data acquired through closely oriented diffusion gradients share some correlation, it is important to remember that dMRI measures apparent diffusion; the extent to which diffusion is hindered with respect to the strength and direction given by the applied diffusion gradient. Therefore, numerical spherical interpolation is susceptible to incorrect estimation in the presence of altered tissue morphology.

In lieu of such information, it is difficult for angular resolution problems to be overcome with hardware techniques; that is to say, having a MRI scanner acquire raw dMRI data at increased angular resolution without any processing whatsoever. Thus, angular resolution enhancement is typically carried out through post-processing techniques. Many such methods make use of Compressed Sensing (CS) introduced in 2006 by Donoho and Lustig [45], and independently by Candes [17]. This technique allows incoherent sub-critically sampled data (with respect to the Nyquist Criterion) to be suitably reconstructed at higher resolutions by taking advantage of the data having a sparse representation in another domain.

Initially, compressed sensing was applied to standard MRI data measured in the k-space, which would correspond only to the spatial domain. Recently however, researchers have been exploring CS applied towards dMRI data which sees acquisition in both k-space and q-space. For dMRI, these studies are typically carried out for DSI or HARDI data. Usage of compressed sensing requires the data to have a sparse representation with a suitable set of basis functions. For HARDI, examples of such studies include Spherical Ridgelets [50], Diffusion basis functions [46], and a rotationally invariant set of atoms depicting common diffusion configurations [5]. Examples of CS applied to DSI data, on the other hand, can be found in [48, 49, 10].

Performing sub-critical sampling with standard linear decoders is prone to noise and aliasing artefacts due to violation of the Nyquist criterion. One way of mitigating this effect would be to use prior information known about the image. In a similar fashion, regularization can be incorporated into the compressed sensing framework to improve results. Examples of these approaches rely on controlling the extent of sparsity, or desired traits of piece-wise constant images which lead to minimization of the Total Variation semi-norm being an effective constraint for the solution. These techniques will be discussed in greater detail later.

5.3 Super resolution via Compressed Sensing: DBF Decomposition

Taking note that dMRI data is characterized by both angular and spatial sampling, recent studies have approached the idea of increasing resolution in both spatial and angular domains. These studies note that the k-space and q-space formalisms are independent, which allows one to propose separate sparse representations in each domain and combine both techniques into a single compressed sensing framework. These methods differ from each other by their choice of bases for sparse representation, regularization and consequentially the cost function to be minimized. In this thesis, the super resolution approach introduced by Mani et al. [46] has been chosen as a reference benchmark, owing to proven effectiveness in their choice of basis functions and usage of regularizers commonly depicted in the literature. Specifically, Mani et al. [46] proposes to use the well known Diffusion Basis Functions decomposition as a foundation for building their super resolution framework in both spatial and angular domains. Their strategy can be roughly split into two stages; the acquisition process and their framework for recovering the data at increased spatial resolution via the theory of compressed sensing.

Previously, we have discussed how MRI acquisition in the k-space is usually performed using EPI due to the significant reduction in acquisition time through scanning a slice in one TR. Instead of using EPI data, they have chosen a variable density spiral sampling scheme [44] which can be achieved by having the phase encoding follow a sinusoidal trajectory with increasing amplitude. The main motivation behind their decision to use variable spirally sampled data is owing to the philosophy of CS, which dictates that correlation between data points is minimized via an incoherent sampling scheme. From statistical signal processing, minimizing the correlation between sampled data points is key to reducing the redundancy in the sampled set, which in turn maximizes the presence of unique data. In addition to the spiral acquisition scheme, they seek to *jointly* undersample the k and q-space by arbitrarily retaining samples. Their method of jointly undersampling the combined spaces ensures that the k-space between any pair of DWIs are sampled differently.

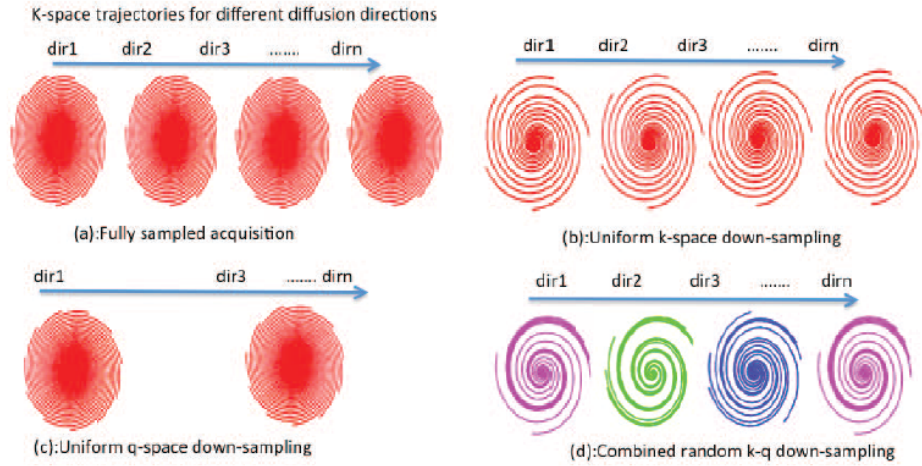


Figure 5.1: Different Sparse Representation of HARDI data. Credit [46].

Moving on to the actual compressed sensing framework, DBF decomposition states that at each voxel, under certain assumptions on the underlying diffusion process such as slow exchange (refer to GMM section for a detailed explanation) the diffusion profile can be modelled by a linear combination of N Gaussian diffusion functions $\{\psi_i\}_{i=1}^N$ aligned towards primary diffusion orientation (also referred to as diffusion compartments), having a mathematical equivalent of

$$\frac{S(b, \mathbf{g})}{S_0} = \sum_{i=1}^N f_i \psi_i(b, \mathbf{g}) \quad (5.1)$$

where the goal is to recover the volume fractions $\{f_i\}_{i=1}^N$ corresponding to the proportion of the voxel covered by each of the N diffusion compartments. The set of these N orientations are typically chosen to represent uniform coverage on the unit sphere. Denoting the number of diffusion gradients chosen by K ($K < N$), this results in K linear equations in N unknowns in the following manner. Formally, at every voxel location a matrix $\mathbf{A} \in \mathbb{R}^{K \times N}$ can be constructed which contains the N basis functions for each of the K diffusion gradients,

$$\mathbf{A} = \begin{pmatrix} \psi_1(b, \mathbf{g}_1) & \cdots & \psi_N(b, \mathbf{g}_1) \\ \vdots & \vdots & \vdots \\ \psi_1(b, \mathbf{g}_K) & \cdots & \psi_N(b, \mathbf{g}_K) \end{pmatrix}. \quad (5.2)$$

Before proceeding, recall that data measured by the MRI scanner in k-space and the diffusion signal are related via the Fourier transform \mathcal{F} . Therefore, the normalized diffusion signal can be expressed in terms of the k-space measurements through $\mathbf{y} = \mathcal{F}\{S_k/S_0\}$. Substituting this new definition into (5.1), the relation between measurements and desired volume fractions can be rewritten as $\mathbf{y} = \mathbf{A}\mathbf{f} + \epsilon$, where ϵ accounts for the presence of measurement noises and \mathbf{f} denotes concatenation of the volume fractions $\{f_i\}_{i=1}^N$ into a \mathbb{R}^N column vector.

Noting that this is an underdetermined linear system admitting an infinite number of solutions, the next step is to establish regularizers which constrain the set of solutions to satisfy additional requirements. In accordance with the theory of compressed sensing, sparsity of the volume fractions $\{f_i\}_{i=1}^N$ is enforced, which has a strong physical interpretation of limiting the number of fibers present in any given voxel. This gives rise to the minimization of the \mathbb{L}_1 norm as a regularization term $\|\mathbf{f}\|_1$. In addition to sparsity, the authors also proposed to minimize the Total Variation semi-norm across each reconstructed DWI, which effectively imposes piecewise continuity on the resulting normalized diffusion signal. Historically, this method has been shown to be extremely effective in image denoising algorithms. Aggregation of the above terms leads to the compressed sensing problem being formulated as

$$\mathbf{f} = \arg \min_{\mathbf{f}} \{ \|\mathbf{A}\mathbf{f} - \mathbf{y}\|_2^2 + \lambda_1 \|\mathbf{A}\mathbf{f}\|_{TV} + \lambda_2 \|\mathbf{f}\|_1 \}, \quad (5.3)$$

with $\lambda_1, \lambda_2 \in \mathbb{R}^+$ being regularization constants controlling the weighting between the terms and which are typically experimentally determined. An important observation stems from noticing this is a convex optimization problem. In their original work they elected to apply the Conjugate Gradient algorithm with the \mathbb{L}_1 substituted by an iterative reweighted \mathbb{L}_2

norm minimization (since the gradient of the \mathbb{L}_1 norm is not defined at every location). Following the recovery of the volume fractions $\{f_i\}_{i=1}^N$, their resulting diffusion ODF was reconstructed using the well known analytical expression [69],

$$\Phi(b, \mathbf{u}_k) = \sum_{i=1}^N \frac{f_i}{Z} \sqrt{\frac{\pi b}{\mathbf{u}_k^T \mathbf{D}_i^{-1} \mathbf{u}_k}}, \quad (5.4)$$

where \mathbf{D}_i refers to the rotated diffusion tensor for the i^{th} diffusion compartment denoting primary diffusion orientation and Z is an appropriate normalization constant. In their experiments it was shown that even under 6-fold reduction of sampling, reconstruction of the diffusion ODF was reasonably accurate under these conditions [46].

Chapter 6

Main Contribution

In the previous chapter, it was seen how compressed sensing can be applied towards HARDI data; namely how the diffusion profile at any given spatial location has a sparse representation in terms of diffusion basis functions. In this chapter, the main contribution of this thesis involving another compressed sensing framework will be discussed. This approach, henceforth referred to as Generalized Compressed Sensing Reconstruction (GCSR), builds upon the use of Spherical Ridgelets for representing HARDI data proposed by Michailovich et al. [50].

6.1 Organization of HARDI datasets

Before moving on to the compressed sensing framework, an overview of the notations and brief methodologies used to represent HARDI data is in order. Recall that single shell HARDI data consists of a set of diffusion weighted images, each representing the attenuation in MR signal in the presence of diffusion encoding along a given direction and fixed strength.

We have seen earlier how the k-space Ω_k and the q-space Ω_q has been used in HARDI to represent the domains of the scanner acquisition and diffusion gradients respectively. This approach allows a better classification of large HARDI data. Typically each DWI is presented as a two dimensional image of size $N_x \times N_y$, where $N_x, N_y \in \mathbb{N}$. For purposes of data storage it is desirable to aggregate this into a single vector of length $|\Omega_k| = N_x \times N_y$ through lexicographical ordering, which can be physically viewed as separating an image by its columns, followed by stacking individual columns to form a single column vector.

Consequentially, a HARDI signal S can be said to exist in the space $\Omega_q \times \Omega_k$, giving rise to two different depictions of the data. Discretization of the q-space is necessary in practice, corresponding to K points indicating diffusion gradient encoding during the acquisition process.

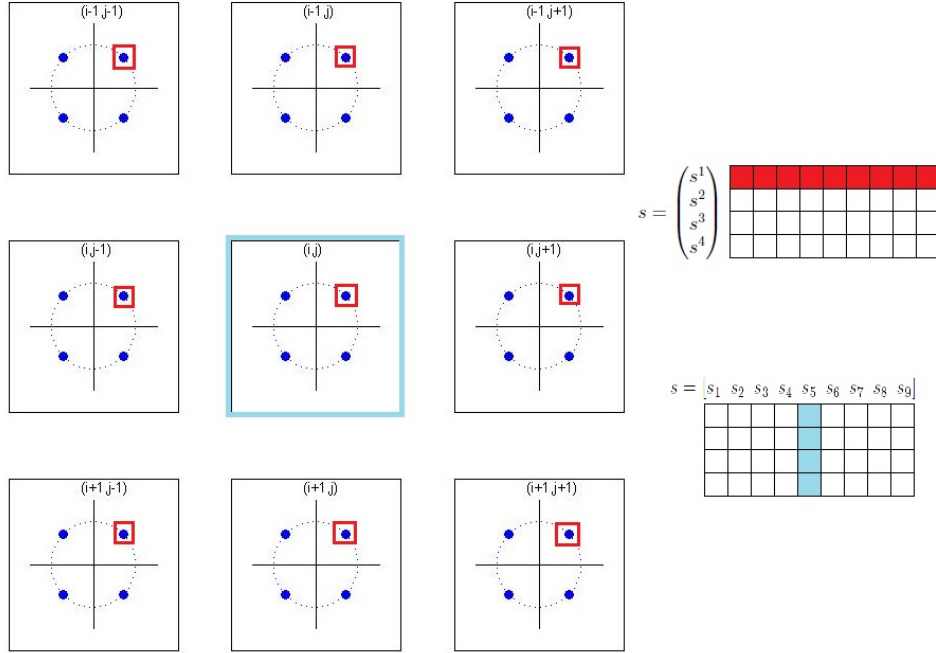


Figure 6.1: Matrix representation of a $K \times |\Omega_k|$ HARDI dataset s , where $K = 4$ and $|\Omega_k| = 9$. Diffusion encodings of the center pixel is represented as s_5 , whereas the DWI corresponding to diffusion encoding direction $[1 \ 1]/\sqrt{2}$ is depicted by s^1 .

Up until now, HARDI data has been exclusively referenced as a set of K images, each under the influence of diffusion encoding and having $|\Omega_k|$ voxels. This perspective provides an intuitive way of explaining the image acquisition process of the MRI scanner. However, another way of expressing HARDI data would be as a collection of $|\Omega_k|$ vectors, with each vector containing K voxels depicting one single spatial location but encoded with a different diffusion gradient. This interpretation of HARDI data is useful for interpretations of ODFs, allowing a thorough understanding of the underlying diffusion behaviour at any given spatial location. To provide the convenience of switching between any of the two representations, a HARDI dataset s can be compactly stored into a single matrix such that $s \in \mathbb{R}^{K \times |\Omega_k|}$. Using this formulation, each row of s denotes a separate DWI and can collectively be represented as $\{s^i\}_{i=1}^K$, where K is the number of diffusion gradients in the

HARDI data. Similarly, each column of s can be treated as a voxel sampled by the set of K diffusion gradients and are collectively represented as $\{s_k\}_{k=1}^{|\Omega_k|}$. In addition to notational convenience, the aforementioned lexicographically ordering allows dimensional reduction in the HARDI dataset. As opposed to a three dimensional array, a matrix representation suffices which in practice substantially simplifies computations in numerical software such as MATLAB.

To consolidate the preceding developments, consider a hypothetical 3×3 segment of Ω_k depicting the $(i, j)^{th}$ spatial location and its 8 immediate neighbours in two dimensional space, sampled over $K = 4$ diffusion gradients uniformly distributed on the unit circle. A visual aid presenting each perspective of HARDI data in matrix form is shown in Figure 6.1. It should be noted that after lexicographical ordering, the $(i, j)^{th}$ pixel corresponds to the fifth column of the matrix.

6.2 Framework of GCSR

Similar to how the compressed sensing approach of Mani [46] relied on sparsity of DBFs at any given spatial location, the success of GCSR relies heavily on the measured signal being sparsely represented by Spherical Ridgelets. As seen previously, in practice HARDI signals has the majority of its energy being concentrated along the great circles. Spherical ridgelets were formulated to satisfy this criterion, making it a natural candidate for the problem at hand.

Consider a set of K unit vectors $\{\mathbf{u}_i\}_{i=1}^K$ denoting the orientation of the diffusion encoded gradients over which a HARDI signal is sampled. For each \mathbf{u}_i , let the frame of J spherical ridgelets $\{\psi_j(\mathbf{u}_i)\}_{j=1}^J$ be constructed. It was previously shown that spherical ridgelets form a dense set in the space of HARDI signals. Therefore, the HARDI signal evaluated at orientation \mathbf{u}_i given a fixed spatial location \mathbf{r} can be expanded as

$$s'(\mathbf{u}_i|\mathbf{r}) = \sum_{j=1}^J \psi_j(\mathbf{u}_i)c_j(\mathbf{r}) \quad (6.1)$$

where $c_j(\mathbf{r})$ denotes the corresponding coefficient for each spherical ridgelet (at spatial location \mathbf{r}) and is known to exist but has yet to be determined. Using the previously defined notation s_k to represent the HARDI signal for the k^{th} spatial location leads to

$$\begin{pmatrix} s_k(\mathbf{u}_1) \\ s_k(\mathbf{u}_2) \\ \vdots \\ s_k(\mathbf{u}_K) \end{pmatrix} = \begin{pmatrix} \psi_1(\mathbf{u}_1) & \psi_2(\mathbf{u}_1) & \cdots & \cdots & \psi_J(\mathbf{u}_1) \\ \psi_1(\mathbf{u}_2) & \psi_2(\mathbf{u}_2) & \cdots & \cdots & \psi_J(\mathbf{u}_2) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \psi_1(\mathbf{u}_K) & \psi_2(\mathbf{u}_K) & \cdots & \cdots & \psi_J(\mathbf{u}_K) \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_J \end{pmatrix} \quad (6.2)$$

or equivalently $s_k = \Phi c$

where the $K \times J$ matrix Φ was introduced as an aggregation of the spherical ridgelets for all K diffusion gradients. Extending upon (6.2) once more to include the full HARDI signal $s' = [s_1 \ s_2 \ \cdots \ s_{|\Omega_k|}] \in \mathbb{R}^{K \times |\Omega_k|}$ leads to $s' = \Phi c'$, where c' is a $J \times \Omega_k$ matrix containing the full set of ridgelet coefficients for all voxels in the HARDI dataset.

The next step in this development is to recover c' , the unknown set of spherical ridgelet coefficients. Although the presence of a matrix equation suggests the naive approach of matrix inversion, this should be avoided. This is because in practice, $K < J$ which means the system in (6.2) is underdetermined. Here is where the bold motivation of GCSR is introduced. Namely, given HARDI data measurements $s' \in \mathbb{R}^{K \times |\Omega_k|}$ we attempt to recover the set of SR coefficients $c \in \mathbb{R}^{J \times |\Omega_{k'}|}$ used to represent the HARDI data at increased spatial resolution $\Omega_{k'}$. Subsequently, c is used to reconstruct $s = \Phi c$, where $s \in \mathbb{R}^{K \times |\Omega_{k'}|}$. In this thesis, a preliminary study of doubling the spatial resolution was explored (e.g., assuming Ω_k has dimensions $N_x \times N_y$ where $N_x, N_y \in \mathbb{N}$ as introduced above, $\Omega_{k'}$ would have dimensions $2N_x \times 2N_y$). This proposal is made possible by the fact that spherical ridgelets provide a sparse representation of HARDI data, with the underlying hypothesis being HARDI data at low spatial resolution providing sufficient information for reconstruction via spherical ridgelets at higher spatial resolution. This methodology suggests that additional information is necessary to constrain the solution for c . Using previously explored ideas, one possible method of recovering the optimal set of SR coefficients c for representing s is by solving a convex optimization problem of the form

$$c^* = \arg \min_c \left\{ \frac{1}{2} \|D\Phi c - s'\|_F^2 + \lambda \|c\|_1 + i_{\geq}(\Phi c) \right\} \quad (6.3)$$

$$\text{where } i_{\geq}(\Phi c) = \begin{cases} 0 & , \Phi c \in \mathbb{R}_+^{K \times \Omega_{k'}} \\ \infty & , \text{otherwise} \end{cases}$$

where the notation $\mathbb{R}_+^{K \times \Omega_{k'}}$ restricts all elements of a $K \times |\Omega_{k'}|$ matrix to be non-negative real numbers, and $\|\cdot\|_F$ represents the Frobenius matrix norm, and D is a spatial downsampling

operator. In this fashion, the first term controls data fidelity between measurement s' and downsampled reconstruction $D\Phi c$, the second term necessitates sparse representation of HARDI data by means of spherical ridgelets through minimizing the \mathbb{L}_1 norm of its coefficients, and the third term regulates that the reconstruction must yield non-negative image intensities, as expected of HARDI data. Finally, $\lambda \in \mathbb{R}^+$ is a regularization constant controlling the weighting between the first and second term.

At this point, it should be noted that the spherical ridgelet frame consists of J ridgelet functions which in practice is much larger than K , the number of diffusion measurements in the HARDI data. Under typical circumstances, estimating ridgelet coefficients c given K measurements in the q-space is impossible. However, the theory of compressed sensing renders this task tractable since the HARDI signal is known to have sparse representation in spherical ridgelets; thus sparsity of its coefficients becomes critical to the success of the optimization framework.

While (6.3) provides a basic reconstruction framework, it can be refined by taking advantage of additional regularization constraints in the spatial domain. Towards this end, two independent regularizers were considered. The first method explores the original idea of enforcing ridgelet coefficient continuity across spatially adjacent voxels. The goal behind this development comes from observing that tissue structures should not exhibit large changes within a local neighbourhood, thereby implicating small deviation in the ridgelet coefficients. For a given spherical ridgelet defined along \mathbf{v} , this can be achieved by minimizing the directional gradient of ridgelet coefficients along \mathbf{v} .

Suppose J spherical ridgelets are constructed by specifying l resolution levels and a bandwidth of ρ . In this manner, let $c_{j,i(j)}$ denote the spherical ridgelet coefficients at the j^{th} resolution level ($j = 1, 2, \dots, l$) and $i(j)$ be an index representing the orientation of the ridgelet ($i = 1, 2, \dots, g(j, \rho)$), g is a function of j and ρ denoting the number of ridgelets at resolution level j . Without loss of generality, we shall assume that $j = 1$ corresponds to the coarsest level of resolution associated with the ridgelet generating functions. Collecting $c_{j,i(j)}(\mathbf{r})$ at each voxel \mathbf{r} in the image space Ω , an image of ridgelet coefficients $I_{j,i(j)} = \{c_{j,i(j)}(\mathbf{r})\}_{\mathbf{r} \in \Omega}$ can be formed in this manner. The corresponding partial differences $\nabla_d : \mathbb{R}^I \rightarrow \mathbb{R}^I$ ($d \in \{1, 2, 3\}$) in the spatial domain along the direction parallel to the Cartesian axes $\{x, y, z\}$ can be computed for each $I_{j,i(j)}$. Applying the aforementioned differencing operators ∇ to each ridgelet coefficient image $I_{j,i(j)}$, the result can be aggregated into the matrix $DI_{j,i(j)} = [(\nabla_1 I_{j,i(j)})^T (\nabla_2 I_{j,i(j)})^T (\nabla_3 I_{j,i(j)})^T]^T$, where $DI_{j,i(j)} \in \mathbb{R}^{3 \times |\Omega|}$. Subsequently, the partial differencing along each Cartesian direction can be extended to partial differencing along any given direction $\mathbf{v}_{i(j)} \in \mathbb{R}^3$ by simply computing $\mathbf{v}_{i(j)}^T DI_{j,i(j)}$. This leads to a convenient expression for assessing the spatial continuity of ridgelet coefficient image $I_{j,i(j)}$, which can be treated equivalently as a minimization of its energy

$$\|\mathbf{v}_{\mathbf{i}(j)}^{\mathbf{T}} DI_{j,i(j)}\|_2^2.$$

Before wrapping up the concept of ridgelet coefficient continuity, it deserves mention that this method was inspired by the notion of fiber continuity [58], which argues that the directional derivative of any given fiber along \mathbf{v} will be minimized. While effective, this method has been shown to produce oversmoothing of the data. On the other hand, imposing continuity on ridgelet coefficients is a less heavy-handed tactic and is more likely to preserve a higher degree of data fidelity. Along that line of thinking, the concept of ridgelet coefficient continuity was abstained in the coarsest resolution of the spherical ridgelets ($j = 1$). Thus, the final expression of preserving continuity on the entire set of ridgelet coefficients c is

$$\|c\|_a^2 = \sum_{j=2}^l \sum_{i(j)} \|\mathbf{v}_{\mathbf{i}(j)}^{\mathbf{T}} DI_{j,i(j)}\|_2^2 \quad (6.4)$$

Another approach to spatial domain regularity explores the minimization of the Total Variation (TV) semi-norm on each DWI in the HARDI dataset. This method has been used in many different applications of image processing with success, by exploring the image in question exhibiting piecewise constant behaviour. Since each DWI corresponds to a brain image with smooth regions excluding the edges separating white matter, grey matter, and CSF, it is a reasonable assumption to classify it as a function of bounded variation. Using this assumption we can define the TV semi-norm on each of the K DWIs in the SR reconstruction $\|\Phi c\|_{TV}$ in the following manner. Adopting the partial differencing operators from ridgelet coefficient continuity, the TV semi-norm can be mathematically expressed as

$$\|\Phi c\|_{TV} = \sum_{k=1}^K \sum_{i=1}^{|\Omega|} \sqrt{\sum_{d=1}^3 |\nabla_d(\Phi c)[k, i]|^2} \quad (6.5)$$

where $(\Phi c)[k, i]$ refers to the k^{th} reconstructed DWI evaluated at the i^{th} spatial location. As a side note, $\|\cdot\|_{TV}$ was classified as a semi-norm owing to an isotropic image (e.g., all intensity values equivalent to $a \in \mathbb{R}^+$, $a \neq 0$) yielding a value of zero while having non-zero intensities. Incorporating both spatial regularizations into the optimization framework, the SR coefficients c can be recovered via solving the minimization problem

$$\min_c \left\{ \frac{1}{2} \|D\Phi c - s'\|^2 + \lambda \|c\|_1 + \mu \|c\|_a^2 + \nu \|\Phi c\|_{TV} + i_{\geq}(\Phi c) \right\}. \quad (6.6)$$

Due to the presence of terms such as $\|c\|_1$, (6.6) is a non-smooth convex optimization problem, which unfortunately rules out the use of gradient based methods barring an approximation to those terms. An alternate approach was to explore the concept of proximal gradient methods, whereby the subdifferential [59] was used in place of the gradient for non-smooth functions. Furthermore, this optimization problem consists of the sum of various functions and difficult to solve directly. Towards this end, we applied the theory of proximal algorithms, which in turn builds upon the notion of proximity operators [54] defined for function f as

$$\text{prox}_f(y) = \arg \min_x \left\{ f(x) + \frac{1}{2} \|x - y\|_2^2 \right\} \quad (6.7)$$

While the rigorous discussion of how proximity algorithms work is beyond the scope of this thesis, this technique was adopted in this thesis due to its proven effectiveness for solving complicated optimization problems such as (6.6). In particular, Alternating Direction Method of Multipliers (ADMM) is a specific proximity algorithm which incorporates variable splitting and iterative variable updates. In this manner, proximity operators need only be evaluated on a single function (for which a closed form solution is typically available) as opposed to the complicated cost function depicted in (6.6).

In what follows, we will see how variable splitting followed by the introduction of augmented Lagrange Multipliers allows one to formulate individual subproblems bearing resemblance to the application of proximal operators. Proceeding with variable splitting, let us first recast this convex optimization problem in the form of a constrained optimization problem by introduction of auxiliary variables ξ , η and ζ as follows

$$\begin{aligned} \min_{c, \eta, \xi, \zeta} \left\{ \frac{1}{2} \|D\Phi c - s'\|^2 + \lambda \|\eta\|_1 + \mu \|\zeta\|_a^2 + \nu \|\xi\|_{TV} + i_{\geq}(\xi) \right\}, \\ \text{s.t. } \Phi c = \xi, \\ c = \eta, \\ c = \zeta \end{aligned} \quad (6.8)$$

Note that the assignment of the auxiliary variables is by no means unique; in fact it is important to choose the assignment wisely such that each individual optimization problem can be solved easily. In general, introducing more auxiliary variables leads to more problems to solve and possibly slow down the process; however the individual problems will likely be easier to solve. A standard procedure for handling these equality constrained problems is

to introduce the augmented Lagrange multipliers for each auxiliary variable. Introducing the Lagrange multipliers $\{p_\xi, p_\eta, p_\zeta\}$ for each respective auxiliary variable, the cost function \mathcal{C} can be revised into the following form

$$\begin{aligned} \mathcal{C}(c, \eta, \xi, \zeta | p_\xi, p_\eta, p_\zeta) = & \frac{1}{2} \|D\Phi c - s'\|^2 + \lambda \|\eta\|_1 + \mu \|\zeta\|_a^2 + \nu \|\xi\|_{TV} + i_{\geq}(\xi) + \\ & \frac{\delta_\xi}{2} \|\Phi c - \xi + p_\xi\|^2 + \frac{\delta_\eta}{2} \|c - \eta + p_\eta\|^2 + \frac{\delta_\zeta}{2} \|c - \zeta + p_\zeta\|^2 \end{aligned} \quad (6.9)$$

where $\{\delta_\xi, \delta_\eta, \delta_\zeta\}$ are positive scalars controlling the difference between the auxiliary variable and its desired value, with the algorithm having guaranteed convergence regardless of their choice. One way to approach minimization of the above cost function, is to do so iteratively via

$$\begin{aligned} (c^{t+1}, \eta^{t+1}, \xi^{t+1}, \zeta^{t+1}) = & \arg \min_{c, \eta, \xi, \zeta} \mathcal{C}(c, \eta, \xi, \zeta | p_\xi, p_\eta, p_\zeta) \\ p_\eta^{t+1} = & p_\eta^t + c^{t+1} - \eta^{t+1}, \\ p_\zeta^{t+1} = & p_\zeta^t + c^{t+1} - \zeta^{t+1}, \\ p_\xi^{t+1} = & p_\xi^t + \Phi c^{t+1} - \xi^{t+1} \end{aligned} \quad (6.10)$$

subject to initial conditions $\eta = \zeta = \xi = 0$, and initial values for the Lagrange Multipliers $p_\zeta^0 = p_\xi^0 = p_\eta^0 = 0$. Through the minimization of \mathcal{C} , the set of SR coefficients c are recovered at increased spatial resolution and subsequently used to construct the estimated HARDI signal $s = \Phi c$. Although this concludes the framework for GCSR, minimization of the cost function \mathcal{C} is very difficult to tackle head on. In what follows, we outline the procedure of how ADMM decomposes the problem into smaller subproblems, each of which admits closed form solutions.

6.3 Solution via ADMM

The strength of ADMM lies in its ability to resolve the optimization of multiple variables into separate optimization problems of one variable each, which makes the solution much more tractable. Using the cost function \mathcal{C} introduced in the GCSR framework, minimization on the function of four variables is reformulated into four sequential minimizations of functions over a single variable

$$\begin{aligned}
c^{t+1} &= \arg \min_c \mathcal{C}(c, \eta^t, \zeta^t, \xi^t) \\
\eta^{t+1} &= \arg \min_\eta \mathcal{C}(c^{t+1}, \eta, \zeta^t, \xi^t) \\
\zeta^{t+1} &= \arg \min_\zeta \mathcal{C}(c^{t+1}, \eta^{t+1}, \zeta, \xi^t) \\
\xi^{t+1} &= \arg \min_\xi \mathcal{C}(c^{t+1}, \eta^{t+1}, \zeta^{t+1}, \xi)
\end{aligned}$$

where the notation $\mathcal{C}(c, \eta^t, \zeta^t, \xi^t)$ should be interpreted as the function of c evaluated at the values from the previous iteration η^t, ζ^t, ξ^t . One has to pay particular attention to how the later minimizations use the most recent values from the other variables as opposed to strictly from the previous iteration (e.g., minimization of ζ uses ξ^t and c^{t+1} due to the later having been updated in the same iteration). While this may seem to discourage parallel processing at a first glance, it will be shown that some of the sub-problems are indeed independent. Let us begin by examining the solution for c^{t+1} . As they do not affect the minimization of c , discarding all the terms independent of c leads to the following optimization problem

$$c^{t+1} = \arg \min_c \left\{ \begin{aligned} &\frac{1}{2} \|D\Phi c - s'\|^2 + \frac{\delta_\xi}{2} \|\Phi c - (\xi^t - p_\xi^t)\|^2 \\ &+ \frac{\delta_\eta}{2} \|c - (\eta^t - p_\eta^t)\|^2 + \frac{\delta_\zeta}{2} \|c - (\zeta^t - p_\zeta^t)\|^2 \end{aligned} \right\} \quad (6.11)$$

Although the presented optimization problem appears at first glance to be a standard least squares problem, it is complicated by the presence of the downsampling operator D . It is effectively alluding to the fact that measurements s' are only available at the original spatial resolution, whereas the reconstruction Φc exist at increased spatial resolution. To account for this in our solution, the data fidelity term $\frac{1}{2} \|D\Phi c - s'\|^2$ was omitted for locations where measurements did not exist. In what follows, the indicator function $1_{s'}(\mathbf{r})$ is used to represent whether measurements s' were present at spatial location \mathbf{r} in the following manner

$$1_{s'}(\mathbf{r}) = \begin{cases} 1, & s'(\mathbf{r}) \text{ exists} \\ 0, & \text{otherwise} \end{cases} \quad (6.12)$$

To simplify notations, the dependence of the indicator function on spatial location is understood and will be omitted. Thus, the solution for c can be succinctly represented as

$$c = [(1_{s'} + \delta_\xi)\Phi^T\Phi + (\delta_\eta + \delta_\zeta)I]^{-1}[1_{s'}\Phi^T s' + \delta_\xi\Phi^T(\xi^t - p_\xi^t) + \delta_\eta(\eta^t - p_\eta^t) + \delta_\zeta(\zeta^t - p_\zeta^t)] \quad (6.13)$$

where I is the identity matrix of size $J \times J$. A few notes regarding the above solution are in order. Firstly, the solution involves the inversion of two matrices $[(1_{s'} + \delta_\xi)\Phi^T\Phi + (\delta_\eta + \delta_\zeta)I]$. Fortunately, this does not pose a problem owing to its modest size $J \times J$, where J denotes the number of SR functions and depends on the number of resolution levels j as well as the bandwidth ρ . In our implementation, we used $j = 3$ and $\rho = 2$ leading to $J = 119$. Furthermore, since the matrices are independent of the variables in the ADMM scheme, in practice it can (and should) be pre-computed to avoid recalculation in each iteration. Secondly, this solution of c accounts for the entire set of SR coefficients required to reconstruct the HARDI data.

The optimization problem for η is also commonly known as a Least Absolute Shrinkage and Selection Operator (LASSO) problem in the literature. It is well known that LASSO problems can be solved using the soft thresholding operator

$$\mathcal{S}_\tau(x) = \max(|x| - \tau, 0) \cdot \text{sign}(x) \quad (6.14)$$

for a given threshold constant $\tau \in \mathbb{R}$. To present the problem such that it resembles the proximity operator form, we choose a threshold value of $\tau = \frac{\lambda}{\delta_\eta}$. Similar to the solution for c , only one single application during each iteration is necessary for recovery of all SR coefficients.

$$\begin{aligned} \eta^{t+1} &= \arg \min_{\eta} \left\{ \frac{\delta_\eta}{2} \|\eta - (c^{t+1} + p_\eta^t)\|^2 + \lambda \|\eta\|_1 \right\}, \\ &= \max \left(|c^{t+1} + p_\eta^t| - \frac{\lambda}{\delta_\eta}, 0 \right) \cdot \text{sign}(c^{t+1} + p_\eta^t) \end{aligned} \quad (6.15)$$

The solution for ζ involves the SR continuity term and takes the form of

$$\zeta^{t+1} = \arg \min_{\zeta} \left\{ \mu \|\zeta\|_a^2 + \frac{\delta_\zeta}{2} \|\zeta - (c^{t+1} + p_\zeta^t)\|^2 \right\} \quad (6.16)$$

One approach involves expanding the SR continuity term. However, we note that each of the SRs are independent, which hints at the possibility of solving each sub-problem and reassembling the results to form iterative update for ζ . Using this approach, each of the sub-problems is a least square problem which can be solved using the concept of a directional

low pass filter in the spectral domain (the implementation for this technique is detailed in the appendix of [82]). Recalling how the RGFs are to be excluded from this process, we define $g(1, \rho)$ as the number of RGFs. Thus, for $i(j) = g(1, \rho) + 1, g(1, \rho) + 2, \dots, J$ each problem becomes

$$\zeta_{i(j)}^{t+1} = \arg \min_{\zeta_{i(j)}} \left\{ \mu \|v_{i(j)}^T DI_{i(j)}\|_2^2 + \frac{\delta_\zeta}{2} \|\zeta_{i(j)} - (c_{i(j)}^{t+1} + p_{\zeta_{i(j)}}^t)\|^2 \right\} \quad (6.17)$$

where informally the term $p_{\zeta_{i(j)}}^t$ refers to the extraction of the $i(j)^{th}$ component on the Lagrange Multiplier at the previous iteration p_ζ^t . Finally, the problem for ξ becomes

$$\xi^{t+1} = \arg \min_{\xi} \left\{ \nu \|\xi\|_{TV} + \frac{\delta_\xi}{2} \|\xi - (\Phi c^{t+1} + p_\xi^t)\|^2 + i_{\geq}(\xi) \right\} \quad (6.18)$$

which is identical to the classic TV denoising problem from image processing, except for the positivity constraint. Unfortunately, this additional term complicates the resulting problem. To handle this, we follow a similar approach of applying variable splitting and proximity algorithms towards this sub-problem. As mentioned previously, the solution requires decomposing the above to account for each of the K DWIs within the HARDI dataset. While there are many such implementations in the literature to solve this problem, the semi-implicit fixed-point solution proposed in [20] was adopted due to its proven numerical stability and fast convergence. Last but not least, the augmented Lagrange multipliers must be updated during every iteration of the ADMM algorithm.

$$p_\eta^{t+1} = p_\eta^t + c^{t+1} - \eta^{t+1}, \quad (6.19)$$

$$p_\zeta^{t+1} = p_\zeta^t + c^{t+1} - \zeta^{t+1}, \quad (6.20)$$

$$p_\xi^{t+1} = p_\xi^t + \Phi c^{t+1} - \xi^{t+1} \quad (6.21)$$

In light of the above developments, the GCSR framework can be summarized as an algorithm over T iterations, with the output being reconstruction of the HARDI data at increased spatial resolution. In practice, T is usually determined by convergence of the cost function. For the GCSR algorithm, convergence was attained within $T = 60$ iterations. Pseudo code for the entire algorithm is provided as a summary of the above developments.

Algorithm 1 Generalized Compressed Sensing Reconstruction

```

1: procedure GCSR( $j, \rho, \delta_\eta, \delta_\xi, \delta_\zeta, \lambda, \mu, \nu, \{\mathbf{v}\}_{j=1}^J, s'$ )
2:  $\Phi \leftarrow \text{SphericalRidgelets}(j, \rho)$ 
3:  $R_1 \leftarrow [(1_{s'} + \delta_\xi)\Phi^T\Phi + (\delta_\eta + \delta_\zeta)I]^{-1}$ 
4: Initial Conditions
    $\xi^0 \leftarrow 0$ 
    $\eta^0 \leftarrow 0$ 
    $\zeta^0 \leftarrow 0$ 
5: Initialize Lagrange Multipliers
    $p_\xi^0 \leftarrow 0$ 
    $p_\eta^0 \leftarrow 0$ 
    $p_\zeta^0 \leftarrow 0$ 
6: for  $t = 0$  to  $T - 1$  do
7:   begin
8:      $c^{t+1} \leftarrow R_1[1_{s'}\Phi^T s' + \delta_\xi\Phi^T(\xi^t - p_\xi^t) + \delta_\eta(\eta^t - p_\eta^t) + \delta_\zeta(\zeta^t - p_\zeta^t)]$ 
9:      $\eta^{t+1} \leftarrow \max\left(|c^{t+1} + p_\eta^t| - \frac{\lambda}{\delta_\eta}, 0\right) \cdot \text{sign}(c^{t+1} + p_\eta^t)$ 
10:    for  $i(j) = g(1, \rho) + 1$  to  $J$  do
11:      begin
12:         $\zeta_{i(j)}^{t+1} \leftarrow \arg \min_{\zeta_{i(j)}} \left\{ \mu \|\mathbf{v}_{i(j)}^T D I_{i(j)}\|_2^2 + \frac{\delta_\zeta}{2} \|\zeta_{i(j)} - (c_{i(j)}^{t+1} + p_{\zeta(i)}^t)\|^2 \right\}$ 
13:      end
14:    for  $k = 1$  to  $K$  do
15:      begin
16:         $\xi_k^{t+1} \leftarrow \arg \min_{\xi_k} \left\{ \nu \sum_{i=1}^{|\Omega|} \sqrt{\sum_{d=1}^3 |\nabla_d(\xi_k)[i]|^2} + \frac{\delta_\xi}{2} \|\xi_k - (\Phi c_k^{t+1} + p_{\xi_k}^t)\|^2 + i_{\geq}(\xi_k) \right\}$ 
17:      end
18:       $p_\eta^{t+1} \leftarrow p_\eta^t + c^{t+1} - \eta^{t+1}$ 
19:       $p_\zeta^{t+1} \leftarrow p_\zeta^t + c^{t+1} - \zeta^{t+1}$ 
20:       $p_\xi^{t+1} \leftarrow p_\xi^t + \Phi c^{t+1} - \xi^{t+1}$ 
21:    end
22:   $s \leftarrow \Phi c$ 

```

Chapter 7

Results

The main objective of this chapter is to evaluate the feasibility of the proposed GCSR algorithm through a series of experiments. This can be objectively measured by determining the degree to which GCSR is able to recover the integrity of fiber orientations for a given HARDI dataset. Towards this end, the testing framework used in this thesis will first be discussed.

Given a HARDI dataset at sufficiently dense spatial and angular dataset which will be henceforth referred to as the ground truth, we will subject it to both spatial and angular downsampling by a factor of d_a and d_s respectively. The resulting dataset will be referred to as measured data. As discussed in the previous chapter, applying the GCSR algorithm to the measured data will result in estimated HARDI data at increased resolution, which will be referred below as reconstructed data. This dataset will have the same spatial resolution as the ground truth but remain at reduced angular resolution. Although the ground truth and the reconstructed data are at different angular resolutions, the underlying diffusion behaviour can be compared by calculating the ODF at each spatial location. In this manner, the ODF can be inferred as consolidation of the angular profile at each voxel into a single function. For our tests, $d_a = 3$ and $d_s = 2$ was chosen to balance the tradeoff between reconstruction integrity and the advantage gained via reduced sampling burden.

Finally, a note on the method of ODF calculation is in order. One candidate was to deploy QBI using spherical harmonics to evaluate the Funk Radon Transform, which was proven successful through different works in the literature. Recall however that there were two different proposals for projection of the EAP onto the unit sphere, namely

$$\begin{aligned}\psi_0(\mathbf{u}) &= \frac{1}{Z} \int_0^\infty P(\mathbf{u}r) dr \\ \psi_2(\mathbf{u}) &= \int_0^\infty P(\mathbf{u}r) r^2 dr\end{aligned}\tag{7.1}$$

where Z is a normalization constant. Both ψ_0 (original formulation proposed by Tuch [69]) as well as ψ_2 (solid angle formulation accounting for increased differential volume at higher radii [1]) were considered during preliminary testing. It was observed that ODFs with sharper maxima were generated by using ψ_2 . Naturally, this came at a cost of noise amplification which subsequently led to a hefty time investment through regularization tuning. On the other hand, ψ_0 produced smoother spherical functions which required substantially less post-processing, and for this reason was adopted for our experiments.

Aside from our proposed GCSR algorithm, the approach to super resolution via GMM proposed by Mani [46] was also implemented (henceforth referred to as GMM Reconstruction) to provide a reference. It should be noted at this point that their original proposal was to minimize the associated cost function through a Conjugate Gradient approach, with the \mathbb{L}_1 norm being approximated by an iterative reweighted \mathbb{L}_2 norm. By noticing how their formulation of the compressed sensing problem bears remarkable resemblance to GCSR in terms of chosen regularizers, it was decided that solving it using proximal algorithms (namely ADMM) had two major advantages. Specifically, using ADMM eliminated the need of approximating the gradient of the \mathbb{L}_1 term. Furthermore, it allowed us to re-use much of our existing framework towards minimizing their cost function.

7.1 Qualitative assessment with ODFs

The standard procedure for data testing involves computer simulated data for preliminary analysis of the performance in a controlled environment, followed by real life (*in vivo*) data to evaluate its feasibility in a clinical setting where more possible unexpected sources of error are likely to surface. Towards that end, let us first look at the computer simulated results. The data was generated using the standard test case of a three fiber GMM model. Two straight cylindrical fibers were constructed to intersect at a 90 degree angle forming a cross-shaped pattern. The remaining fiber resembled a torus and was positioned to surround the aforementioned cross. In accordance with the empirical observation of water diffusivity in highly anisotropic white matter structures [56], the diffusion tensor was set to $\mathbf{D} = 10^{-6} \text{diag}[1700; 300; 300] \text{mm}^2/\text{s}$. This method left a small amount of spatial locations

not covered by any of the three fibers. At these locations, we elected to fit a single quasi-isotropic tensor along a random direction with mean diffusivity of $1.052 \times 10^{-3} \text{ mm}^2/\text{s}$ and fractional anisotropy of 0.22, in accordance to prior dMRI experiments conducted on human grey matter [19].

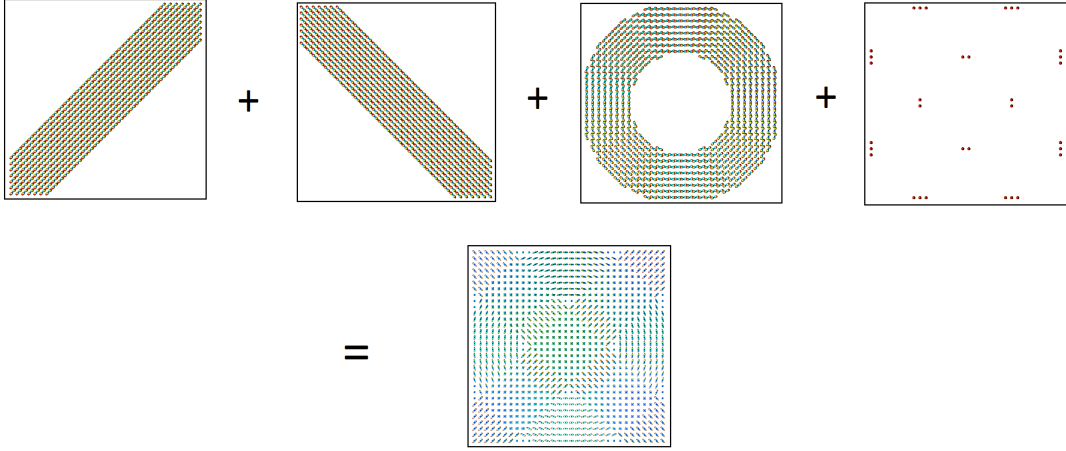


Figure 7.1: Projection of each individual fiber via ODFs onto a 2D plane. The final computer data is the superposition of all three fibers onto the same FOV, along with grey matter simulation in locations without fiber presence.

The simulated data was constructed at a $32 \times 32 \times 10$ field of view. In accordance to HARDI data which typically adopts a b-value of $b \in [1000, 3000] \text{ s/mm}^2$ [4], a b-value of $b = 2000 \text{ s/mm}^2$ was chosen. The method of generalized spiral sampling [60] was adopted for q-space sampling, yielding a total of 162 points providing approximate uniform coverage of the unit sphere. Since the HARDI signal is typically assumed to be symmetric, only the points on the upper hemisphere were retained resulting in $K = 81$ points. Thus, the ground truth data has a dimension of $32 \times 32 \times 10 \times 81$. Downsampling in the angular domain by a factor of $d_a = 3$, the resulting dimension of the measured data was $16 \times 16 \times 5 \times 27$.

The experimental results in the form of ODFs can be seen in Figure 7.2, where a central horizontal slice of dimension 32×32 was chosen for the convenience of visual interpretation. At a first glance the ODF reconstruction from both GCSR and GMM provide very accurate results in areas where only a single fiber is present. In regions with fiber crossings however, it is seen that the ODFs are rather smoothed out when compared to the ground truth. This effect is exemplified in several spatial locations (one of which is enclosed in a grey box), whereby the ground truth contained only one single fiber. However, both reconstructions




Figure 7.2: ODFs captured from computer simulated data at a FOV of 32×32 , where correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction.

failed to account for this by modelling it with the prevalent fiber crossing model among its neighbours. This corner case aside, the main fiber orientations can still be inferred from the reconstructions which is of critical importance in practical applications.

The previous experiment did not account for the presence of noise, an ubiquitous element seen in measured data. Thus, another computer simulation was conducted whereby after fitting the tensors, the data was subsequently contaminated with Rician noise giving rise to SNR of 5; while the Gaussian noise model works relatively well at low b -values ($b < 1000 \text{ s/mm}^2$), it has been shown that the Rician noise model is more appropriate for MRI due to the use of in-phase and quadrature receiver antennae to capture magnitude data [76]. As opposed to choosing a central horizontal slice in the previous experiment, an outer slice further from the middle was chosen to increase the presence of simulated grey matter. By examining the ODFs in Figure 7.3, one can see that the quality of both reconstructions suffered in this experiment. This effect is particularly noticeable in the grey matter areas, where the quasi-isotropic diffusion behaviour cannot be readily recognized in reconstructed data. On the other hand, areas with presence of fibers were represented clearly even in the presence of noise. From these experiments, it can be seen that qualitative results from the reconstructions could not provide a perfect representation of the ground truth. However, they provided an accurate depiction of the underlying fiber structure, which led to the pursuit of testing with *in vivo* data.

The requirements for the *in vivo* data was much stricter for our testing purposes. This is because HARDI data is typically acquired with a spatial resolution of $2 \text{ mm} \times 2 \text{ mm} \times$




Figure 7.3: ODFs captured from computer simulated data contaminated with Rician noise at $\text{SNR} = 5$, where correspondence is defined as: (Left) Ground Truth without noise, (Middle) GCSR, (Right) GMM Reconstruction

2 mm, which unfortunately leads to significant partial volume effects within voxels and is not suitable for accurate fiber tractography [26, 28]. For our purposes, the optimal dataset would contain data sampled at $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$ spatial resolution. Measured data obtained through our downsampling routine would effectively have $2 \text{ mm} \times 2 \text{ mm} \times 2 \text{ mm}$ spatial resolution, meaning our results here would closely resemble that of typical HARDI data. With current MRI scanners it is very difficult to achieve $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$ resolution across an entire brain volume, since even top of the line gradient coils would overheat and be unable to complete imaging. Fortunately, the problem of inherent low spatial resolutions for HARDI data has been recognized in the dMRI community. After much collaboration between dMRI researchers, the Human Connectome Project (HCP) consortium ¹ have successfully acquired HARDI data at $1.25 \text{ mm} \times 1.25 \text{ mm} \times 1.25 \text{ mm}$ spatial resolution.

The HCP *in vivo* data was obtained using a SS-EPI protocol on a modified Siemens

¹*In vivo* data used in the preparation of this work were obtained from the MGH-USC Human Connectome Project (HCP) database (<https://ida.loni.usc.edu/login.jsp>). The HCP project (Principal Investigators : Bruce Rosen, M.D., Ph.D., Martinos Center at Massachusetts General Hospital; Arthur W. Toga, Ph.D., University of California, Los Angeles, Van J. Weeden, MD, Martinos Center at Massachusetts General Hospital) is supported by the National Institute of Dental and Craniofacial Research (NIDCR), the National Institute of Mental Health (NIMH) and the National Institute of Neurological Disorders and Stroke (NINDS). Collectively, the HCP is the result of efforts of co-investigators from the University of California, Los Angeles, Martinos Center for Biomedical Imaging at Massachusetts General Hospital (MGH), Washington University, and the University of Minnesota.

Skyra 3T scanner with $TE = 89.5$ ms, $TR = 5520$ ms, where the modification involved customization of enhanced gradient coils noted to be beneficial towards diffusion imaging [64]. The q-space samples were collected over three concentric spherical shells defined using $b = \{1000, 2000, 3000\}$ s/mm², where 270 q-space samples were distributed equally between each of the three shells using the acquisition scheme proposed by [18] to achieve uniform sampling over each spherical shell. Finally, 15 b_0 images (e.g., T_2 MRI images without applying diffusion gradients) were provided for normalization purposes. For our testing purposes only the subset of b_0 and $b = 2000$ s/mm² data was used since the GCSR algorithm was designed for single shell HARDI data. Applying GCSR to a three dimensional $32 \times 32 \times 5$ FOV near the left and right ventricles lead to the following results depicted in Figure 7.4 (results are overlaid on top of fractional anisotropy maps to highlight fiber geometries).




Figure 7.4: ODFs captured from HCP *in vivo* data near the left and right ventricles, where correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction

From these results we can see that the *in vivo* results are fairly close to results from the computer simulated experiments, in the sense that fiber trajectories can be inferred accurately from observing the ODFs. It is very interesting to note that the noise present in the ground truth is significantly attenuated in the reconstructions. The similarity between both GCSR and GMM Reconstructions results suggests that this effect is a direct consequence of the regularization and increased spatial resolution as opposed to differences between choice of sparse representation of HARDI signals. More *in vivo* results can be seen in Figure 7.5 where a similar phenomenon is noted. To more effectively evaluate the discrepancies between the ground truth and reconstructed data, let us now move on to more concrete quantitative metrics.




Figure 7.5: ODFs captured from HCP *in vivo* data in regions exhibiting bending and curving geometries. Correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction

7.2 Quantitative metrics for HARDI data

As seen in the previous section, some minor discrepancies are observed between the ODFs in ground truth data versus both compressed sensing approaches. However, there were two important unanswered questions thus far, namely (1) how severe is the discrepancy in terms of quantitative metrics and (2) to what extent does smoothing of ODFs affect the practical application of modelling of neural fiber tracts in the brain? Unsurprisingly, these questions have also been brought up in the literature. Among different quantitative metrics that have been proposed, the following two are most often used in the field. They are known as Normalized Mean Square Error (NMSE) and Average Angular Error (AAE)

and are designed to measure different sources of discrepancies between datasets. Let us first discuss the more intuitive NMSE metric, which is defined as

$$\text{NMSE} = \frac{1}{N_x N_y N_z} \sum_{\mathbf{r}} \frac{\|s(\mathbf{r}) - \hat{s}(\mathbf{r})\|_2^2}{\|s(\mathbf{r})\|_2^2} \quad (7.2)$$

where the product $N_x N_y N_z$ collectively describes the number of voxels in the spatial domain, $s(\mathbf{r})$ and $\hat{s}(\mathbf{r})$ denotes the reference and estimated HARDI signal respectively for a given vector \mathbf{r} denoting the spatial location. To account for measurement noises causing potential instability issues when $\|s(\mathbf{r})\|_2^2$ approaches 0, outliers in the reference data were replaced by the median value within its minimum size local neighbourhood for NMSE calculations.

Similar to the notion of MSE in image processing, the integrity of NMSE is heavily compromised when the reference and estimated signals are misaligned or suffer deformation, which would call for image registration routines to remedy this problem. Fortunately, these problems do not arise through the frameworks used to recover these reconstructions.

Nevertheless, NMSE does not provide any means of inferring the primary directions of water molecule displacement from diffusion, the primary motivation of HARDI. While an ODF is very useful at providing qualitative evaluations of this feature, it would be helpful to assess this numerically. Towards that end, the metric Average Angular Error (AAE) was proposed to measure (in degrees) the orientation discrepancy in terms of the principal positive angle between two diffusion profiles through

$$\text{AAE} = \frac{180}{\pi} \text{E}[\cos^{-1}(\mathbf{u}_0 \cdot \hat{\mathbf{u}})] \quad (7.3)$$

where \mathbf{u}_0 and $\hat{\mathbf{u}}$ are unit vectors denoting the principal orientation of diffusion within a voxel in reference and estimated data respectively, and $\text{E}[\cdot]$ denotes the expected value operation which in practical applications is replaced by the sample mean among all the voxels. In practice, \mathbf{u}_0 and $\hat{\mathbf{u}}$ are estimated by using the direction corresponding to ODF maxima; this hinges on the assumption that ODF maximas denote the principal direction of diffusion at any given voxel.

The discrepancy between the ground truth and reconstructed data from both GCSR and GMM was measured through NMSE and AAE in the computer simulated data and *in vivo* data. To provide a broader assessment of different regimes of diffusion processes (e.g., fast versus slow), the computer simulated data was generated for b-values (see section 2.7) of $b \in \{1000, 2000, 3000\} \text{ s/mm}^2$. In addition, the synthesized Rician noise level was also

varied such that the SNR lied between $[5, 40]$ to account for acquired datasets in practice having different degrees of signal integrity. On the other hand, the *in vivo* metrics were computed for the regions depicted in Figures 7.4 and 7.5 and can be found in Figure 7.7.

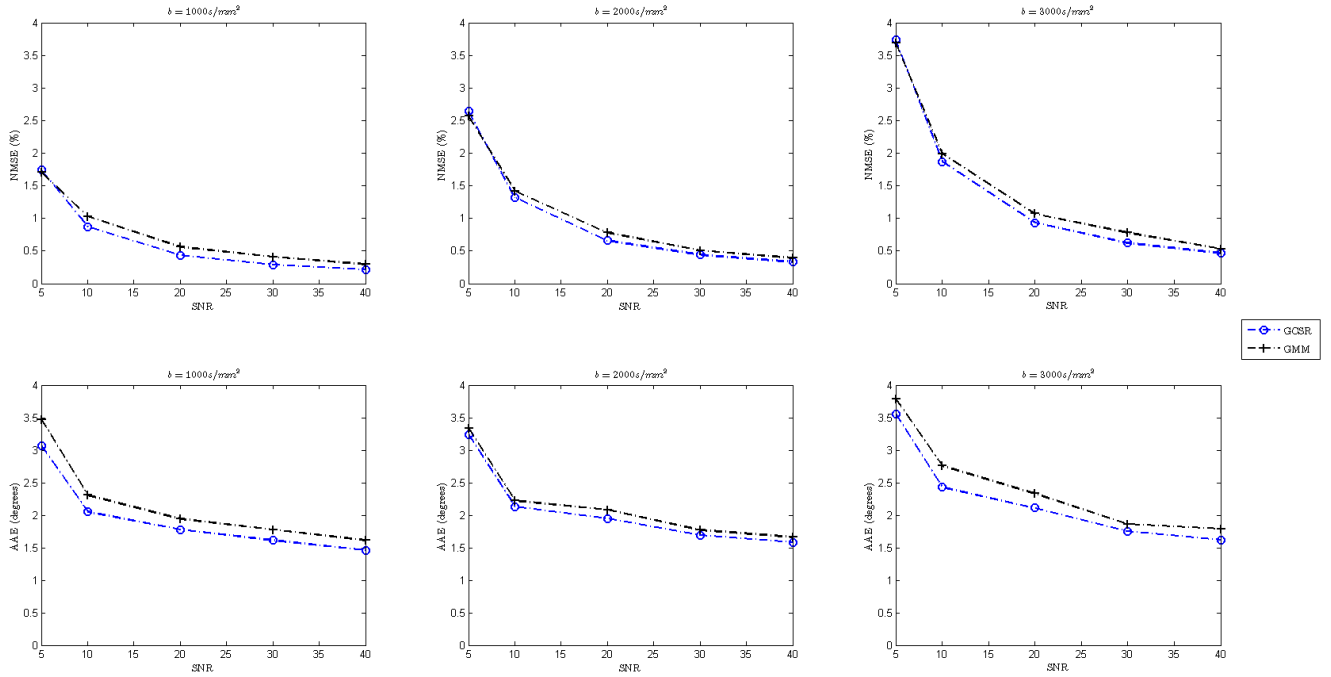


Figure 7.6: NMSE (top row) and AAE metrics (bottom row) for reconstructions of computer simulated data at various SNR levels and b-values.

In all of the plots of Figure 7.6, it can readily be seen that having a higher b-value increases the error metrics. This is due to increased signal attenuation in the presence of stronger applied diffusion gradients, and therefore the measurement noises becomes increasingly comparable in magnitude relative to the desired signal. Also, it is observed that GCSR yields lower measurement errors for nearly all of the experiments, with the exception of NMSE at low SNR levels. Last but not least, as expected both error metrics decrease in the presence of higher SNR levels.

On the other hand, the bar graphs of Figure 7.7 show that the NMSE is acceptable at less than 4% for all regions of interest. Unsurprisingly, the region containing smooth fiber trajectories have lowest errors and are in line with the earlier qualitative observations. It is interesting to note that the region with bending fibers has a higher NMSE than the one




Figure 7.7: NMSE (left) and AAE metrics (right) for reconstructions of HCP *in vivo* data in various regions. Fiber geometries of smooth, bending and curving correspond to Figures 7.4, 7.5 (top) and 7.5 (bottom) respectively.

exhibiting curving fibers, but on the other hand boasts a lower AAE. The discrepancy is most likely due to bending fibers exhibiting a sharper rate of change in orientation.

While these results show that a promising estimate is provided through applying either of these compressed sensing frameworks, it is beneficial to assess these results through a more practical approach of directly modelling fiber trajectories within the brain. Since such applications would have intrinsic methods of attenuating the effects of measurement noises, this provides an effective alternative of validating the reconstructed data. Specifically, this field is known as *Fiber Tractography* and will be discussed in more detail below.

7.3 Fiber Tractography

In this section we discuss how neural tracts in the brain can be modelled by looking at fiber tractography results. There are two distinct flavours of fiber tractography designed to overcome the shortcomings of the other; streamline and probabilistic tractography. While streamline tractography models the fiber trajectory directly, its weakness is its inability to probe possible but non-primary fiber directions. Conversely, probabilistic tractography

seeks to provide the greatest coverage of possible fiber orientations by generating connectivity maps, which computes the probability that the dominant neural tract through a pre-defined seed point passes through any given voxel. While being an immensely useful tool in practical applications, unfortunately connectivity maps cannot quantitatively assess the accuracy of reconstructed data. For this reason, fiber tractography experiments in this thesis will be limited to streamline tractography.




Figure 7.8: Various bundles of fiber tracts generated via streamline tractography using the Camino software.

7.3.1 Streamline Tractography

In streamline tractography, a region of interest is chosen as the set of seed points from which fiber trajectories are originated. The fiber trajectories are typically found by examining fiber orientation distribution function (fODF) maximas at each voxel location. These fODFs are not to be confused with the previously mentioned diffusion ODFs. Whereas diffusion ODFs classify the probability of diffusivity along a given direction, the resulting distribution is typically too fanned out to precisely trace out a given fiber trajectory. This is particularly problematic in the case of multiple fibers crossing at small angles, in which the sharp peaks of fODFs are invaluable towards identifying the number of fibers and their respective orientations. Towards that end, tractography algorithms will need to first estimate the fODFs from a dMRI dataset, through techniques such as spherical deconvolution

[67]. From this, one can see streamline tractography provides a very convenient method of comparing reconstruction to ground truths due to the tractography output being fiber tracts. Provided that there is a one-to-one correspondence between the fibers in ground truth and reconstructions, we can perform quantitative analysis by directly calculating metrics such as spatial displacement between the two fibers. An example of streamline tractography generated with the Camino software [22] using data from HCP can be seen in Figure 7.8, whereby different bundles of fibers can be distinguished by the colour coding. Even within this small segment, various fiber geometries such as the bending blue fibers or the relatively straight dark green fibers can be observed. Because tractography outputs only the neural tracts, the results are typically overlaid on top of DTI fractional anisotropy maps for visual guidance.

Qualitative results from fiber tractography on the HCP data can be seen in Figure 7.9, where the enclosed region contains a multitude of bending fibers and is identical to the one found in the first row of Figure 7.5. It can be seen that tracing ODF maximas from Figure 7.5 yields a similar result to the generated fiber tracts of Figure 7.9. Although the ODFs from reconstructed data were smoothed as seen in the previous experiments, this effect is largely absent in fiber tractography results as the three fiber bundles are distinct and can be clearly identified. Last but not least, it is interesting to note that fiber bundles are denser in both reconstructed datasets.




Figure 7.9: Fiber tractography using HCP *in vivo* data (see Figure 7.5), where correspondence is defined as: (Left) Ground Truth, (Middle) GCSR, (Right) GMM Reconstruction

In the dMRI community, tractography competitions are held regularly to stimulate the development of algorithms both for fODF generation and the actual fiber tracking. Researchers are invited to submit their tractography algorithms, which are subsequently tested on dMRI scans of pre-established phantoms and benchmarked against one another.

Regardless of whether the phantoms were synthetically constructed or computer made, it provided a ground truth basis for comparison between the actual tract layout and the result from various tractography algorithms. One such competition was the FiberCup competition (<http://www.lnao.fr/spip.php?rubrique79>), which not only provided a synthetically constructed phantom illustrating different commonly encountered fiber configurations but also presented several quantitative metrics for evaluating each algorithm. Considerable insights from that competition were borrowed towards the experiments conducted in this thesis. While explanation of key ideas are provided in this thesis, the reader is also encouraged to consult the manuscript for a deeper justification for their design decisions [28].

7.3.2 Quantitative metrics in Streamline Tractography

The streamline tractography results from Camino can be interpreted as a list of points indicating coordinates for which the fiber trajectory passes through. As this corresponds to a discrete sampling of the fiber, the immediate question that comes to mind is a proper means of interpolation between the points. The FiberCup competition elected to fit cubic B-splines over the sampling points to model the fiber trajectory, which is appropriate given neural tracts typically do not exhibit heavily jagged paths. For this reason (with the exception of one pathological test case, to be discussed later), the experiments of this thesis also use this method of interpolation. Following this, each fiber trajectory is resampled at a higher number of points.

The next issue to be addressed is appropriate quantitative metrics for comparing two distinct fiber trajectories. The fiber trajectories are typically given as a collection of spatial locations through which they traverse through. An intuitive metric is to measure the average spatial displacement between the two fibers. Note that the aforementioned resampling process ensures each fiber contains precisely N points. Denoting the n^{th} point along two fibers by $(x_{1,n}, y_{1,n}, z_{1,n})$ and $(x_{2,n}, y_{2,n}, z_{2,n})$ respectively, the average spatial displacement is simply the Euclidean Distance between each point represented as

$$\text{Spatial Displacement} = \frac{1}{N} \sum_{n=1}^N \sqrt{(x_{2,n} - x_{1,n})^2 + (y_{2,n} - y_{1,n})^2 + (z_{2,n} - z_{1,n})^2} \quad (7.4)$$

However, it can be easily seen that for fibers in close proximity with one another, this metric will yield good results but will fail to account for differences in fiber geometries.

Since fiber tractography is performed based on fODF reconstruction at each voxel, needless to say examining fiber geometry is critical for assessing reconstruction fidelity. Towards that end, the tangent metric was proposed for measuring the differences in shape between two fibers. The tangent metric computes the tangent vector to the fiber at each of the sampled points on the fiber. Subsequently, the angle between the vectors is calculated and averaged over all N sampling points. Defining the normalized tangent vectors at the first fiber’s n^{th} point by $t_{1,n} = \frac{f'(n)}{\|f'(n)\|}$ (second fiber by $t_{2,n}$ respectively), this can be mathematically presented as

$$\text{Tangent Metric} = \frac{1}{N} \sum_{n=1}^N \frac{180}{\pi} |\cos^{-1}(t_{1,n} \cdot t_{2,n})| \quad (7.5)$$

where once again N denotes the number of sampling points on the fiber. Whereas fibers in parallel would have a significant spatial displacement, the tangent metric is able to identify a large degree of similarity between parallel fibers. In tandem, these two metrics provide a fairly deep assessment of the difference between two neural tracts.

7.3.3 Experimental results from Fiber Tractography

As seen in Figure 7.8, the presence of fibers with different geometric characteristics is prevalent in many different parts of the human brain. It is clear that arbitrarily selecting fibers and comparing metrics can yield vastly differing results. Therefore, it is important to first classify available fiber geometries to avoid the trap of comparing apples to oranges.

Towards that end, five distinct test cases have been chosen based on their perceived fiber geometry. Three of the test cases are designed to test single fiber geometries that are dominant in areas of the brain such as the corpus callosum, while the other two test cases deal with more complex geometries involving multiple intertwined fiber trajectories. The selected regions of the brain as well as examples of fiber geometry are shown in 7.10. For each of the five fiber geometries presented in Figure 7.10, results from 20 fibers exhibiting similar geometry was averaged to reduce uncertainty.

Finally, a word on the datasets used for this experiment is in order. As with the previous analysis of HARDI data, the super resolution method proposed by Mani [46] (henceforth also referred to as the reference algorithm) was adopted as a benchmark for GCSR. Table 7.1 reports the spatial displacement and tangent metrics for each case.

Before comparing the performance of the GCSR and GMM reconstruction, it is helpful to first understand Camino’s philosophy of performing streamline tractography as well as

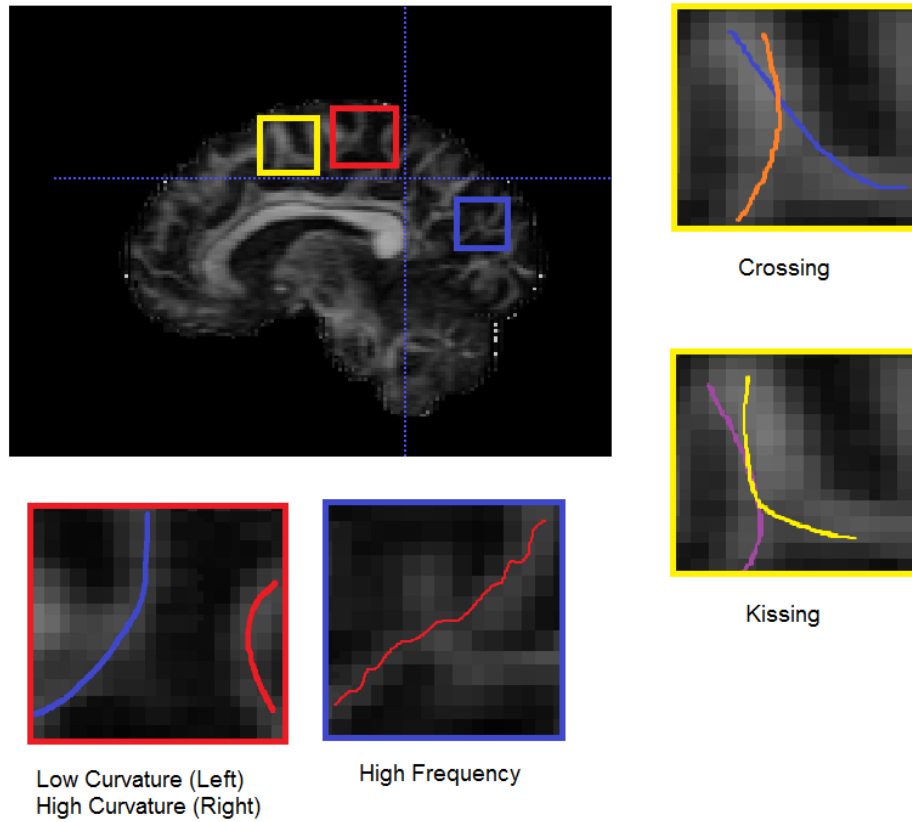


Figure 7.10: Selective fibers in the designated ROIs are chosen for error metrics based on their distinctive geometries.

the proposed test cases. Firstly, tractography algorithms have a tendency to favour fiber trajectories with low curvatures. This is due to its frequent occurrence in the human brain and should be treated as the best case. Indeed, low curvature fibers have very little spatial and tangential error between the reconstructions and the ground truth. This is further reinforced by looking at the results for high curvature fibers, which exhibit a significantly increased error in both metrics. Related to the first point is how fiber trajectories usually do not exhibit high frequency twisting and bending behaviour. To provide a more complete streamline tractography analysis one such test case was included to reflect more pathological situations. Last but not least, the last two test cases look at scenarios where two fibers are in such close proximity that the presence of both are detected within a single voxel. As discussed previously in the dMRI section this is a critical case

Table 7.1: Comparison of tractography metrics between GCSR and GMM Reconstructed data

Geometry	GCSR		GMM Reconstruction	
	Spatial (mm)	Tangent (°)	Spatial (mm)	Tangent (°)
Low Curvature	0.033	0.014	0.048	0.016
High Curvature	0.19	1.32	0.26	1.39
High Frequency	0.84	1.83	0.86	1.94
Crossing	1.61	0.16	1.77	0.18
Kissing	3.15	0.29	3.39	0.35

that cannot be inferred accurately with DTI.

With a fundamental understanding of how Camino’s tractography algorithm behaviour, one can subsequently gain considerable insight by analysing the discrepancies between reconstructed data and ground truth. The performance between GCSR and the Reference algorithm are relatively similar for (1) Low Curvature and (3) High Frequency cases. This can be attributed to tractography algorithm behaviour because they are well and badly equipped to handle those two respective cases; thus the reconstruction fidelity becomes masked. On the other hand, for (2) High Curvature as well as the multi fiber cases GCSR is ahead by a clear margin. These results show that GCSR provides a superior reconstruction, something that is difficult to assess based on visually examining the qualitative results.

7.3.4 Comparison with denoised measurements

While the ground truth data provided an invaluable reference point, a non-negligible degree of measurement noise was observed in the ground truth dataset. In the following experiment, the performance of GCSR is compared to ground truth data processed with a suitable denoising algorithm. In the dMRI community, the Non-Local Means (NLM) Algorithm [76] has been demonstrated by many to be effective at removing the Rician noise contamination induced by the quadrature receiver coils in an MRI scanner. It is interesting to note that the NLM Algorithm was initially conceived with a Gaussian noise model, which is an acceptable approximation for high SNR data. However as discussed earlier HARDI datasets tend to require higher b-values which in turn leads to lower SNR due to the increased presence of noise when the base MR signal is strongly attenuated. In

Table 7.2: Comparison of tractography metrics between GCSR and NLM filtered version of ground truth data

Geometry	GCSR		NLM Filtered Ground Truth	
	Spatial (mm)	Tangent (°)	Spatial (mm)	Tangent (°)
Low Curvature	0.033	0.014	0.029	0.008
High Curvature	0.19	1.32	0.14	1.27
High Frequency	0.84	1.83	0.83	1.84
Crossing	1.61	0.16	1.33	0.15
Kissing	3.15	0.29	2.78	0.26

this experiment, the NLM algorithm was adapted from a freely available MATLAB implementation provided by Pierrick Pouper [47] and applied on the ground truth. Results can be seen in Table 7.2. As we can see from the results of GCSR reconstructed and filtered data in Table 7.2, the discrepancy between them is comparable with one another in all test cases. This suggests that the error incurred during GCSR reconstruction as seen in the previous experiments can be attributed to the noise present in measured ground truth data.

The series of streamline tractography experiments provided another means of validation between ground truth and reconstructed data. Even though the previous experiments with ODFs showed some discrepancy between ground truth and reconstructed datasets, the fiber tractography experiments provides sufficient evidence that GCSR can provide data at increased spatial resolution without affecting the ability to accurately recover the underlying fiber tracts. It was also demonstrated that smaller degrees of spatial displacement and tangent metrics resulted from GCSR as opposed to GMM Reconstruction.

Chapter 8

Summary and Future Works

In this thesis we have started with the basic physics associated with diffusion and Nuclear Magnetic Resonance, and saw how they can be applied to medical imaging. Studies such as CHARMED and NODDI have showed that dMRI is a promising tool for exploring pathways of neural tracts within the human brain. Among various dMRI models, HARDI provided a balance between acquisition time and effectively recovering diffusion orientations through sampling diffusion gradients over a spherical shell. The diffusion orientation profile was consolidated into a spherical probability distribution function known as the diffusion ODF, which was evaluated at each voxel in the spatial domain.

Revisiting the goal of dMRI to model neural fiber tracts in the brain leads to the revelation that in practice, the typical level of spatial resolution is insufficient. This can be seen from the partial volume effect, where different types of tissue are present within a single voxel. A similar argument was made in the q-space, where the long acquisition time associated with dense q-space sampling would inevitably result in motion artifacts induced by patient movement. These concerns stimulated the research in dMRI towards methods of increasing the spatial resolution while maintaining reasonable acquisition times. While both hardware implementations and post-processing techniques were explored in the literature, post-processing techniques were found to be more feasible in practice.

Of particular importance to this thesis was the theory of compressed sensing, allowing data with a sparse representation in some basis to be accurately reconstructed from very few samples. Using that fact that HARDI signals could be sparsely represented using spherical ridgelets, the problem of inherently low spatial resolution in HARDI data could be addressed by means of super resolution. The proposed GCSR algorithm used appropriate regularization functions such as minimization of the TV semi-norm across each DWI and

ridgelet coefficient continuity to recover SR coefficients at increased spatial resolution, which was formulated as a convex optimization problem. Subsequently, a specific form of proximity algorithm known as ADMM utilized the powerful technique of variable splitting to reformulate minimization of the complicated cost function into simpler sub-problems.

Testing of our approach with computer simulated and *in vivo* data lead to the realization that the reconstructed ODFs were comparable to the ground truth. Comparisons using quantitative metrics reflected that reconstructed data provided a reasonable estimation to the ground truth. To further assess the viability of reconstructed data in practical applications, validation was made through fiber tractography experiments under the hypothesis that the observed discrepancy was attributed to measurement noises in the ground truth data. Streamline tractography experiments using different fiber geometries showed that the associated discrepancy between ground truth and reconstructed data was negligible. Moreover, comparison between the reference method showed that our proposed compressed sensing using spherical ridgelets was slightly superior. Finally, it was seen through tractography with denoised ground truth data that the reconstruction error was reduced, validating the previous hypothesis of measurement errors due to the presence of noise.

That being said, there is still room for improvement in future works. For example, the *in vivo* data for our experiments were acquired using EPI, which boasted fast acquisition times but came at the expense of spatial resolution. Spiral sampling retains the advantage of accelerated acquisition while featuring less adverse artifacts in the data, likely improving the quality of reconstruction. Another direction would be incorporating datasets from patients with known disorders (such as Alzheimer’s Disease) and determining whether tract based statistics from performing fiber tractography on more localized super resolution data would yield deeper insight towards diagnosis. Similarly, the quantitative metrics used for assessing HARDI data could be improved by considering increased weighting in areas of importance as opposed to uniform weighting across all voxels. Last, but not least, other forms of regularizing the data should also be explored to improve accuracy of the reconstructions. For example, the diffusion encoded images could have a sparse representation in a multi-resolution transform such as framelet frames [25].

References

- [1] Iman Aganj, Christophe Lenglet, Guillermo Sapiro, Essa Yacoub, Kamil Ugurbil, and Noam Harel. Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. *Magnetic Resonance in Medicine*, 64(2):554–566, 2010.
- [2] Andrew L Alexander, Khader M Hasan, Mariana Lazar, Jay S Tsuruda, and Dennis L Parker. Analysis of partial volume effects in diffusion-tensor mri. *Magnetic Resonance in Medicine*, 45(5):770–780, 2001.
- [3] Yaniv Assaf and Peter J Basser. Composite hindered and restricted model of diffusion (charmed) mr imaging of the human brain. *Neuroimage*, 27(1):48–58, 2005.
- [4] Haz-Edine Assemlal, David Tschumperlé, Luc Brun, and Kaleem Siddiqi. Recent advances in diffusion mri modeling: Angular and radial reconstruction. *Medical image analysis*, 15(4):369–396, 2011.
- [5] Suyah P Awate and Edward VR DiBella. Compressed sensing hardi via rotation-invariant concise dictionaries, flexible k-space undersampling, and multiscale spatial regularity. In *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*, pages 9–12. IEEE, 2013.
- [6] Alan Barnett. Theory of q-ball imaging redux: Implications for fiber tracking. *Magnetic Resonance in Medicine*, 62(4):910–923, 2009.
- [7] Peter J Basser, James Mattiello, and Denis LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1):259–267, 1994.
- [8] TEJ Behrens, H Johansen Berg, Saad Jbabdi, MFS Rushworth, and MW Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34(1):144–155, 2007.

- [9] TEJ Behrens, MW Woolrich, M Jenkinson, H Johansen-Berg, RG Nunes, S Clare, PM Matthews, JM Brady, and SM Smith. Characterization and propagation of uncertainty in diffusion-weighted mr imaging. *Magnetic resonance in medicine*, 50(5):1077–1088, 2003.
- [10] Berkin Bilgic, Kawin Setsompop, Julien Cohen-Adad, Anastasia Yendiki, Lawrence L Wald, and Elfar Adalsteinsson. Accelerated diffusion spectrum imaging with compressed sensing using adaptive dictionaries. *Magnetic Resonance in Medicine*, 68(6):1747–1754, 2012.
- [11] Kai Tobias Block, Martin Uecker, and Jens Frahm. Undersampled radial mri with multiple coils. iterative image reconstruction using a total variation constraint. *Magnetic resonance in medicine*, 57(6):1086–1098, 2007.
- [12] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [13] Mark A Brown and Richard C Semelka. *MRI: basic principles and applications*. John Wiley & Sons, 2011.
- [14] Fernando Calamante, Jacques-Donald Tournier, Graeme D Jackson, and Alan Connelly. Track-density imaging (tdi): super-resolution white matter imaging using whole-brain track-density mapping. *Neuroimage*, 53(4):1233–1243, 2010.
- [15] Paul T Callaghan. *Translational dynamics and magnetic resonance: principles of pulsed gradient spin echo NMR*. Oxford University Press, 2011.
- [16] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [17] Emmanuel J Candès, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8):1207–1223, 2006.
- [18] Emmanuel Caruyer, Jian Cheng, Christophe Lenglet, Guillermo Sapiro, Tianzi Jiang, Rachid Deriche, et al. Optimal design of multiple q-shells experiments for diffusion mri. In *MICCAI Workshop on Computational Diffusion MRI-CDMRI’11*, 2011.

- [19] Mara Cercignani, Matilde Inglese, Elisabetta Pagani, Giancarlo Comi, and Massimo Filippi. Mean diffusivity and fractional anisotropy histograms of patients with multiple sclerosis. *American Journal of Neuroradiology*, 22(5):952–958, 2001.
- [20] Antonin Chambolle and Pierre-Louis Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76(2):167–188, 1997.
- [21] Stuart Clare. *Functional MRI: Methods and Applications*. PhD thesis, University of Nottingham, October 1997.
- [22] PA Cook, Y Bai, SKKS Nedjati-Gilani, KK Seunarine, MG Hall, GJ Parker, and DC Alexander. Camino: open-source diffusion-mri reconstruction and processing. In *14th scientific meeting of the international society for magnetic resonance in medicine*, volume 2759, 2006.
- [23] Pierrick Coupé, José V Manjón, Maxime Chamberland, Maxime Descoteaux, and Bassem Hiba. Collaborative patch-based super-resolution for diffusion-weighted images. *NeuroImage*, 83:245–261, 2013.
- [24] Alessandro Daducci, E Canales-Rodriguez, Maxime Descoteaux, Y Gur, M Mani, S Merlet, A Ramirez-Manzanares, P Rodrigues, M Reisert, and F Seppehrband. Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion mri. *Medical Imaging, IEEE Transactions on*, 33(2):384–399, 2013.
- [25] Ingrid Daubechies, Bin Han, Amos Ron, and Zuowei Shen. Framelets: Mra-based constructions of wavelet frames. *Applied and Computational Harmonic Analysis*, 14(1):1–46, 2003.
- [26] Maxime Descoteaux, Rachid Deriche, TR Knosche, and Alfred Anwander. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *Medical Imaging, IEEE Transactions on*, 28(2):269–286, 2009.
- [27] Maxime Descoteaux, Rachid Deriche, Denis Le Bihan, Jean-François Mangin, and Cyril Poupon. Multiple q-shell diffusion propagator imaging. *Medical image analysis*, 15(4):603–621, 2011.
- [28] Pierre Fillard, Maxime Descoteaux, Alvina Goh, Sylvain Gouttard, Ben Jeurissen, James Malcolm, Alonso Ramirez-Manzanares, Marco Reisert, Ken Sakaie, Fatima Tensaouti, et al. Quantitative evaluation of 10 tractography algorithms on a realistic diffusion mr phantom. *Neuroimage*, 56(1):220–234, 2011.

- [29] Lawrence R Frank. Anisotropy in high angular resolution diffusion-weighted mri. *Magnetic Resonance in Medicine*, 45(6):935–939, 2001.
- [30] Lawrence R Frank. Characterization of anisotropy in high angular resolution diffusion-weighted mri. *Magnetic Resonance in Medicine*, 47(6):1083–1099, 2002.
- [31] Paul Funk. Über eine geometrische anwendung der abelschen integralgleichung. *Mathematische Annalen*, 77(1):129–135, 1915.
- [32] Xiujuan Geng, Thomas J Ross, Hong Gu, Wanyong Shin, Wang Zhan, Yi-Ping Chao, Ching-Po Lin, Norbert Schuff, and Yihong Yang. Diffeomorphic image registration of diffusion mri using spherical harmonics. *Medical Imaging, IEEE Transactions on*, 30(3):747–758, 2011.
- [33] Alexander C Guo, James R MacFall, and James M Provenzale. Multiple sclerosis: Diffusion tensor mr imaging for evaluation of normal-appearing white matter 1. *Radiology*, 222(3):729–736, 2002.
- [34] Robin M Heidemann, Alfred Anwander, Thorsten Feiweier, Thomas R Knösche, and Robert Turner. k-space and q-space: combining ultra-high spatial and angular resolution in diffusion imaging using zooppa at 7t. *Neuroimage*, 60(2):967–978, 2012.
- [35] William R Hendee and E Russell Ritenour. *Medical imaging physics*. John Wiley & Sons, 2003.
- [36] Christopher P Hess, Pratik Mukherjee, Eric T Han, Duan Xu, and Daniel B Vigneron. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magnetic Resonance in Medicine*, 56(1):104–117, 2006.
- [37] Hao Huang, Jiangyang Zhang, Hangyi Jiang, Setsu Wakana, Lidia Poetscher, Michael I Miller, Peter van Zijl, Argye E Hillis, Robert Wytik, and Susumu Mori. Dti tractography based parcellation of white matter: application to the mid-sagittal morphology of corpus callosum. *Neuroimage*, 26(1):195–205, 2005.
- [38] Bing Jian and Baba C Vemuri. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted mri. *Medical Imaging, IEEE Transactions on*, 26(11):1464–1471, 2007.
- [39] Heidi Johansen-Berg and Timothy EJ Behrens. *Diffusion MRI: From quantitative measurement to in-vivo neuroanatomy*. Academic Press, 2009.

- [40] Derek K Jones, Thomas R Knösche, and Robert Turner. White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion mri. *Neuroimage*, 73:239–254, 2013.
- [41] Irina Kezele, Maxime Descoteaux, Cyril Poupon, Fabrice Poupon, and J-F Mangin. Spherical wavelet transform for odF sharpening. *Medical image analysis*, 14(3):332–342, 2010.
- [42] Cheng Guan Koay, Lin-Ching Chang, John D Carew, Carlo Pierpaoli, and Peter J Basser. A unifying theoretical and algorithmic framework for least squares methods of estimation in diffusion tensor imaging. *Journal of Magnetic Resonance*, 182(1):115–125, 2006.
- [43] Denis Le Bihan. Looking into the functional architecture of the brain with diffusion mri. *Nature Reviews Neuroscience*, 4(6):469–480, 2003.
- [44] C Liu, R Bammer, and M Moseley. High resolution multi-shot sense dti using self-navigated interleaved spirals (snails). In *Proc. Intl. Soc. Mag. Reson. Med*, volume 13, page 10, 2005.
- [45] Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic resonance in medicine*, 58(6):1182–1195, 2007.
- [46] Merry Mani, Mathews Jacob, Arnaud Guidon, Chunlei Liu, A Song, Vincent Magnotta, and Jianhui Zhong. Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 326–329. IEEE, 2012.
- [47] José V Manjón, Pierrick Coupé, Luis Martí-Bonmatí, D Louis Collins, and Montserrat Robles. Adaptive non-local means denoising of mr images with spatially varying noise levels. *Journal of Magnetic Resonance Imaging*, 31(1):192–203, 2010.
- [48] Marion I Menzel, Ek T Tan, Kedar Khare, Jonathan I Sperl, Kevin F King, Xiaodong Tao, Christopher J Hardy, and Luca Marinelli. Accelerated diffusion spectrum imaging in the human brain using compressed sensing. *Magnetic Resonance in Medicine*, 66(5):1226–1233, 2011.
- [49] Sylvain Merlet, Rachid Deriche, et al. Compressed sensing for accelerated eap recovery in diffusion mri. *CDMRI'10*, 2010.

- [50] Oleg Michailovich and Yogesh Rathi. On approximation of orientation distributions by means of spherical ridgelets. *Image Processing, IEEE Transactions on*, 19(2):461–477, 2010.
- [51] Susumu Mori. *Introduction to diffusion tensor imaging*. Elsevier, 2007.
- [52] Shahrum Nedjati-Gilani, Daniel C Alexander, and Geoffrey JM Parker. Regularized super-resolution for diffusion mri. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, pages 875–878. IEEE, 2008.
- [53] Thoralf Niendorf, Rick M Dijkhuizen, David G Norris, Menno van Lookeren Campagne, and Klaas Nicolay. Biexponential diffusion attenuation in various states of brain tissue: Implications for diffusion-weighted imaging. *Magnetic Resonance in Medicine*, 36(6):847–857, 1996.
- [54] Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in optimization*, 1(3):123–231, 2013.
- [55] Sharon Peled and Yehezkel Yeshurun. Superresolution in mri: application to human white matter fiber tract visualization by diffusion tensor imaging. *Magnetic resonance in medicine*, 45(1):29–35, 2001.
- [56] Carlo Pierpaoli, Peter Jezzard, Peter J Basser, Alan Barnett, and Giovanni Di Chiro. Diffusion tensor mr imaging of the human brain. *Radiology*, 201(3):637–648, 1996.
- [57] Alonso Ramirez-Manzanares, Mariano Rivera, Baba C Vemuri, Paul Carney, and Thomas Mareci. Diffusion basis functions decomposition for estimating white matter intravoxel fiber geometry. *Medical Imaging, IEEE Transactions on*, 26(8):1091–1102, 2007.
- [58] Marco Reisert and Valerij G Kiselev. Fiber continuity: An anisotropic prior for odf estimation. *Medical Imaging, IEEE Transactions on*, 30(6):1274–1283, 2011.
- [59] Ralph Rockafellar. On the maximal monotonicity of subdifferential mappings. *Pacific Journal of Mathematics*, 33(1):209–216, 1970.
- [60] Edward B Saff and A BJ Kuijlaars. Distributing many points on a sphere. *The Mathematical Intelligencer*, 19(1):5–11, 1997.
- [61] Benoit Scherrer, Ali Gholipour, and Simon K Warfield. Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. *Medical image analysis*, 16(7):1465–1476, 2012.

- [62] Liewei Sha, Hua Guo, and Allen W Song. An improved gridding method for spiral mri using nonuniform fast fourier transform. *Journal of Magnetic Resonance*, 162(2):250–258, 2003.
- [63] Stamatios N Sotiropoulos, Saad Jbabdi, Jesper L Andersson, Mark W Woolrich, Kamil Ugurbil, and Timothy EJ Behrens. Rubix: Combining spatial resolutions for bayesian inference of crossing fibers in diffusion mri. *IEEE transactions on medical imaging*, 32(6):969, 2013.
- [64] Stamatios N Sotiropoulos, Saad Jbabdi, Junqian Xu, Jesper L Andersson, Steen Moeller, Edward J Auerbach, Matthew F Glasser, Moises Hernandez, Guillermo Sapiro, Mark Jenkinson, et al. Advances in diffusion mri acquisition and processing in the human connectome project. *Neuroimage*, 80:125–143, 2013.
- [65] EO Stejskal and JE Tanner. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The journal of chemical physics*, 42(1):288–292, 1965.
- [66] Paul Tofts. *Quantitative MRI of the brain: measuring changes caused by disease*. John Wiley & Sons, 2005.
- [67] J-Donald Tournier, Fernando Calamante, David G Gadian, and Alan Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted mri data using spherical deconvolution. *NeuroImage*, 23(3):1176–1185, 2004.
- [68] Joel A Tropp and Anna C Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *Information Theory, IEEE Transactions on*, 53(12):4655–4666, 2007.
- [69] David S Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.
- [70] David S Tuch, Timothy G Reese, Mette R Wiegell, Nikos Makris, John W Belliveau, and Van J Wedeen. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48(4):577–582, 2002.
- [71] David S Tuch, Timothy G Reese, Mette R Wiegell, and Van J Wedeen. Diffusion mri of complex neural architecture. *Neuron*, 40(5):885–895, 2003.
- [72] David S Tuch, RM Weisskoff, JW Belliveau, and VJ Wedeen. High angular resolution diffusion imaging of the human brain. In *Proceedings of the 7th Annual Meeting of ISMRM, Philadelphia*, volume 321, 1999.

- [73] L Wald. MR Image Encoding. 2006.
- [74] Van J Wedeen, Patric Hagmann, Wen-Yih Isaac Tseng, Timothy G Reese, and Robert M Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54(6):1377–1386, 2005.
- [75] Van J Wedeen, RP Wang, Jeremy D Schmahmann, T Benner, WYI Tseng, Guangping Dai, DN Pandya, Patric Hagmann, Helen D’Arceuil, and Alex J de Crespigny. Diffusion spectrum magnetic resonance imaging (dsi) tractography of crossing fibers. *Neuroimage*, 41(4):1267–1277, 2008.
- [76] Nicolas Wiest-Daesslé, Sylvain Prima, Pierrick Coupé, Sean Patrick Morrissey, and Christian Barillot. Rician noise removal by non-local means filtering for low signal-to-noise ratio mri: applications to dt-mri. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008*, pages 171–179. Springer, 2008.
- [77] BJ Wilm, Jonas Svensson, A Henning, KP Pruessmann, P Boesiger, and SS Kollias. Reduced field-of-view mri using outer volume suppression for spinal cord diffusion imaging. *Magnetic Resonance in Medicine*, 57(3):625–630, 2007.
- [78] Yu-Chien Wu and Andrew L Alexander. Hybrid diffusion imaging. *NeuroImage*, 36(3):617–629, 2007.
- [79] Pew-Thian Yap, Hongyu An, Yasheng Chen, and Dinggang Shen. Fiber-driven resolution enhancement of diffusion-weighted images. *NeuroImage*, 84:939–950, 2014.
- [80] Wenxing Ye, Sharon Portnoy, Alireza Entezari, Stephen J Blackband, and Baba C Vemuri. An efficient interlaced multi-shell sampling scheme for reconstruction of diffusion propagators. *Medical Imaging, IEEE Transactions on*, 31(5):1043–1050, 2012.
- [81] Hui Zhang, Torben Schneider, Claudia A Wheeler-Kingshott, and Daniel C Alexander. Noddi: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage*, 61(4):1000–1016, 2012.
- [82] Quan Zhou. Spatially regularized reconstruction of fibre orientation distributions in the presence of isotropic diffusion. Master’s thesis, University of Waterloo, April 2014.