Acquisition- and Modeling-independent Resolution Enhancement of Brain dwMRI Volumes

by

Mohammad Salem Bajammal

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE

in

The Faculty of Graduate and Postdoctoral Studies

(Electrical and Computer Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

August 2016

Abstract

Diffusion-weighted magnetic resonance imaging (dwMRI) provides unique capabilities for non-invasive imaging of neural fiber pathways in the brain. dwMRI is an increasingly popular imaging method and has promising diagnostic and surgical applications for Alzheimer's disease, brain tumors, and epilepsy, to name a few.

However, one limitation of dwMRI (specifically, the more common diffusion tensor imaging scheme, DTI) is that it suffers from a relatively low resolution. This often leads to ambiguity in determining location and orientation of neural fibers, and therefore reduces the reliability of information gained from dwMRI.

Several approaches have been suggested to address this issue. One approach is to have a finer sampling grid, as in diffusion spectrum imaging (DSI) and high-angular resolution imaging (HARDI). While this did result in a resolution improvement, it has the side effects of lowering the quality of image signal-to-noise ratio (SNR) or prolonging imaging time, which hinders its use in routine clinical practice.

Subsequently, an alternative approach has been proposed based on super-resolution methods, where multiple low resolution images are fused into a higher resolution one. While this managed to improve resolution without reducing SNR, the multiple acquisitions required still resulted in a prolonged imaging time.

In this thesis, we propose a processing pipeline that uses a super res-

olution approach based on dictionary learning for alleviating the dwMRI low resolution problem. Unlike the majority of existing dwMRI resolution enhancement approaches, our proposed framework does not require modifying the dwMRI acquisition. This makes it applicable to legacy data. Moreover, this approach does not require using a specific diffusion model.

Motivated by how functional connectivity (FC) reflects the underlying structural connectivity (SC), we use the Human Connectome Project and Kirby multimodal dataset to quantitatively validate our results by investigating the consistency between SC and FC before and after super-resolving the data. Based on this scheme, we show that our method outperforms interpolation and the only existing single image super-resolution method for dMRI that is not dependent on a specific diffusion model. Qualitatively, we illustrate the improved resolution in diffusion images and illustrate the revealed details beyond what is achievable with the original data.

Preface

The work performed in this thesis has resulted in the following publications:

- (In progress) Bajammal, M. and Ng, B. and Abugharbieh, R. "High Resolution Diffusion MRI Data without Acquisition Modifications"

  This paper was based on a collaboration between Bajammal and BiSICL alumni Dr. Ng under the supervision and guidance of Prof. Abugharbieh. Dr. Ng. contributed the ideas of: using online dictionary to enable method scalability, using two different databases for building the dictionary and testing to show generalizability, and using affinity propagation to find prototype gradient volumes to reduce computational load. Bajammal contributed the code implementations, generated the results, and the idea of using a clustering approach on multi-shell data. The paper manuscript is still in its early stages and will be edited by all co-authors.

- Bajammal, M. and Yoldemir, B. and Abugharbieh, R. "Comparison of Structural Connectivity Metrics for Multimodal Brain Image Analysis", International Symposium on Biomedical Imaging (ISBI), Brooklyn-USA, Pages: 934–937, April 2015.

  This paper was based on a collaboration between Bajammal and BiSICL PhD candidate Yoldemir under the supervision and guidance of Prof. Abugharbieh. Yoldemir contributed the paper idea, preprocessing of

the data and parcellation, and implementation of the tractography method and three of the four anatomical connectivity metrics. Bajammal contributed the implementation of the fourth anatomical connectivity metric and the validation code as well as generated the results. In terms of paper writing, both students contributed equally. The paper was edited by the supervisor. Parts of this paper are included in Chapter 4.

- Yoldemir, B. and Bajammal, M. and Abugharbieh, R. "Dictionary Based Super-Resolution for Diffusion MRI", MICCAI Workshop on Computational Diffusion MRI (CDMRI), Cambridge-USA, Pages: 194–204, September 2014.

  This paper was based on a collaboration between Bajammal and BiSICL PhD candidate Yoldemir under the supervision and guidance of Prof. Abugharbieh. Bajammal contributed the algorithmic idea conception, implementation of the method and validation scheme, as well as generation of the results. Yoldemir contributed the application idea conception, preprocessing of the data and parcellation, and the validation scheme. In terms of manuscript writing, Yoldemir contributed the majority of the effort. The paper was edited by the supervisor. Parts of this paper are included in Chapter 3.

Table of Contents

Abstract ii

Preface iv

Table of Contents vi

List of Figures viii

List of Acronyms xi

Acknowledgements xiii

Dedication xiv

1 Introduction 1

1.1 Motivation and Problem Statement 1

1.2 Magnetic Resonance Imaging of the Brain 2

1.3 Diffusion Modeling 10

1.4 Limitations and Alternative Imaging Schemes 12

1.4.1 Diffusion Spectrum Imaging 12

1.4.2 High Angular Resolution Diffusion Imaging 14

1.5 Thesis Objectives and Proposed Approach 15

1.6 Thesis Organization 17

1.7 Thesis Contributions 17

2 Related Literature 19
2.1 Overview 19
2.2 Shifting-based dwMRI Super-resolution 20
2.2.1 Translational Field of View Shifting 21
2.2.2 Orientational Field of View Shifting 22
2.3 Model-based dwMRI Super-resolution 24
2.3.1 Super-resolved Diffusion Tensor 25
2.3.2 Super-resolved Spherical Deconvolution 26

3 Proposed Framework 28
3.1 Overview 28
3.2 Processing Pipeline 28
3.2.1 Diffusion Shell Clustering 30
3.2.2 Training Matrix Preconditioning 36
3.2.3 Joint-dictionary Learning 47

4 Validation Methodology 57
4.1 Overview 57
4.2 Functional Connectivity Estimation 58
4.3 Structural Connectivity Estimation 59

5 Results 66
5.1 Datasets 66
5.2 Quantitative Performance 67
5.3 Qualitative Results 76

6 Conclusions 82
6.1 Discussion 82
6.2 Thesis Contributions and Future Work 83

Bibliography 85

List of Figures

1.1 Basic elements of NMR experiment 4
1.2 Spin alignment in external B_0 field. 6

2.1 Spatial and orientational field-of-view shifts 23

3.1 An overview of the proposed processing pipeline. 29
3.2 Illustration of diffusion shells 32
3.3 Shell clustering correlation matrix. Volumetric patch dimension: $d = 19$ 37
3.4 Shell clustering correlation matrix. Volumetric patch dimension: $d = 11$ 38
3.5 Shell clustering correlation matrix. Volumetric patch dimension: $d = 3$ 39
3.6 Effect of volumetric patch dimension on correlation matrix range 40
3.7 Sample diffusion clustering result. 41
3.8 Comparison of condition number of sparse coding matrix. . 44
3.9 The proposed training matrix preconditioning process. . . . 45
3.10 Local structure detection using non-linear approximation . . 48
3.11 Result of non-linear thresholding 49
3.12 Binary thresholding for atom initialization 50
3.13 Coding matrix condition numbers after preconditioning . . 51

3.14 Effect of dictionary optimization weight and size on reconstruction quality 55
3.15 Effect of dictionary atom size on reconstruction quality . . . 56

4.1 Comparison of four common SC metrics in terms of SC-FC (resting state fMRI) correlation for 38 subjects from the Human Connectome Project. 63
4.2 Comparison of four common SC metrics in terms of SC-FC (task-fMRI) correlation for 38 subjects from the Human Connectome Project. The shaded bands represent the standard deviation of SC-FC across 7 different tasks. 64
4.3 Parcels with highest connectivity to posterior cingulate cortex as obtained from: (a) Functional connectivity (FC), (b) normalized fiber count, and (c) average FA. Note how the arrangement of connected parcels obtained by normalized fiber count has a better resemblance to the parcels obtained by FC, compared to average FA. 65

5.1 Relative reconstruction error η for HCP dataset [91] test subjects. The vertical bars show the standard deviation of error across diffusion directions at each shell. The vertical bars for the $b0$ volume show the standard deviation of error across repeated $b0$ acquisitions. 68
5.2 Relative reconstruction error η for Kirby dataset [51] test subjects using dictionaries trained on HCP dataset [91]. The vertical bars show the standard deviation of error across diffusion directions. 71

5.3 Relative reconstruction error η for Kirby dataset [51] test subjects using dictionaries trained on the remaining subjects in the same dataset. The vertical bars show the standard deviation of error across diffusion directions. 73
5.4 Relative reconstruction error η for Kirby dataset [51] test subjects, trained on the remaining subjects in the dataset. The selection of test versus training subjects in this figure is the opposite of Figure 5.3. 74
5.5 Quantitative SC-FC performance of resolution enhancement 77
5.6 Quantitative SC-FC performance of resolution enhancement 78
5.7 Quantitative SC-FC performance of resolution enhancement 80
5.8 Qualitative comparison, raw diffusion images 81

List of Acronyms

AP	Affinity propagation
DSI	Diffusion spectrum imaging
DTI	Diffusion tensor imaging
DW-MRI	Diffusion-weighted Magnetic Resonance Imaging
EPI	Echo planar imaging
FA	Fractional anisotropy
FC	Functional connectivity
fMRI	Functional magnetic resonance imaging
FOV	Field of view
GFA	Generalized fractional anisotropy
HARDI	High angular resolution diffusion imaging
HCP	Human Connectome Project dataset
k-space	Space of spatial position
MRI	Magnetic resonance imaging
MSE	Mean-squared Error
NLM	Non-local means

NMR	Nuclear magnetic resonance
ODF	Orientation distribution function
PGSE	Pulsed Gradient Spin Echo
PDF	Probability density function
QBI	Q-ball imaging
q-space	Space of spin displacement
ROI	Region of interest
SC	Structural connectivity
SNR	Signal to noise ratio

Acknowledgements

I would like to express my gratitude and appreciation to my supervisor Prof. Rafeef Abugharbieh for her continuous support and guidance. Her help and advice was key to my success.

I would also like to thank all current and previous members of the Biomedical Signal and Image Computing Laboratory (BiSICL) at the University of British Columbia for their companionship and support.

I also wish to thank Dr. Pierrick Coupé from the Laboratoire Bordelais de Recherche en Informatique, France, for running comparative experiments using their CLASR method.

Dedication

... to my family

Chapter 1

Introduction

1.1 Motivation and Problem Statement

Diffusion-weighted magnetic resonance imaging (dwMRI) is an increasingly common approach of performing brain imaging, with many applications in both research and clinical practice. For instance, dwMRI has been utilized in the assessment and therapy planning of brain tumors [30, 58], in the neurological analysis and modeling of schizophrenia [6, 23] and Alzheimer's disease [56, 77], for the prognosis and treatment monitoring of multiple sclerosis [25, 40], the diagnosis and abnormalities detection in traumatic brain injuries [18, 53], as well as the assessment and planning of surgical interventions in epilepsy [33, 96], among many other applications.

A major motivation behind the increasingly common utilization of dwMRI is that it provides powerful capabilities for non-invasive imaging of neural structures in the brain. The accurate estimation of these structures enables a more precise understanding of the structural connectivity in the brain.

However, the accuracy of estimating these neural structures is often hampered by the inherently low resolution of dwMRI. A single pixel (or voxel, for 3D data) can therefore contain many distinct fibers with differing orientations, especially in the commonly used diffusion tensor imaging scheme (DTI). At such locations, the orientation typically becomes ambiguous, which leads to erroneous information about brain structure.

Therefore, increasing the resolution of dwMRI data holds great promise

towards more accurate delineation of fibers. Accordingly, there has been a number of modified dwMRI imaging approaches aiming for increased resolution, which will be explored in the next chapter. However, they tend to have practical limitations such as reduced image quality and a long imaging time. Such limitations motivate the search for another approach for increasing resolution. This work will present an alternative approach of achieving this goal.

1.2 Magnetic Resonance Imaging of the Brain

Diffusion-weighted magnetic resonance imaging (dwMRI) is one of the subcategories of magnetic resonance imaging (MRI). It is an imaging technique that uses water diffusion strength as a contrast in MRI images. As such, a research investigation involving dwMRI imaging would benefit from an overview of the underlying principles of MRI imaging.

MRI is an imaging method for creating an image of magnetic properties of the nuclei of objects being imaged. More specifically, MRI is based on the physical phenomenon of nuclear magnetic resonance (NMR), which describes the interaction of external electromagnetic radiation with nuclei in a magnetic field. It is this phenomenon that makes MRI imaging possible, via external electromagnetic radiation probing of the nuclei of an imaged object.

The interaction of a nucleus with external fields depends on the *spin* of the nucleus. The spin is a quantum mechanical measure of angular momentum. The value of a quantum mechanical spin depends on the number of protons and neutrons in the nucleus. Accordingly, each atom and isotope (atom variants having different number of neutrons) has a particular spin value.

The quantum mechanical spin of a nucleus is a major determinant of

whether or not a material composed of that nucleus can be imaged using MRI. A nucleus that has a spin value of zero is not affected by magnetic fields and can not be imaged in MRI. This is illustrated in Figure 1.1. A nucleus has a zero spin when the number of protons and neutrons are both even numbers.

In order to be able to detect a nucleus using MRI, it should have a non-zero integer or half-integer quantum mechanical spin, which is the case for odd values of number of protons or neutrons. Fortunately, a large number of biological tissues are composed of materials whose spin values are integer or half-integer. In practice, almost all medical MRI imaging is based on the *hydrogen nucleus* because of the large proportion of water in body tissues and the fact that hydrogen's spin value is half-integer.

Another property of the nucleus that affects its interaction with external fields is its *magnetic moment* μ. Like the quantum spin, the magnetic moment also depends on the number of protons and neutrons in the nucleus. As such, each atom or isotope has a magnetic moment value. For the commonly used hydrogen nucleus, the magnetic moment is $\mu = 2.79$ N-m/T. The value of the moment indicates the amount of torque a nucleus will experience when a force is exerted on it by an external field.

The magnetic moment is not a scalar value, but rather a vector quantity. The orientation of the vector is aligned with the axis of rotation of the nucleus. Most objects of medical interest are relatively large and macro-scale, containing billions of nuclei. Therefore, it is more practical to define the *net magnetization* vector as the combination of the individual magnetic moments in the nuclei.

In a tissue in its normal state (i.e. without any external excitations), the individual magnetic moments are randomly distributed due to the random locations and orientations of nuclei. Therefore, the net magnetization vector is practically zero. On the other hand, when a tissue is placed in an exter-

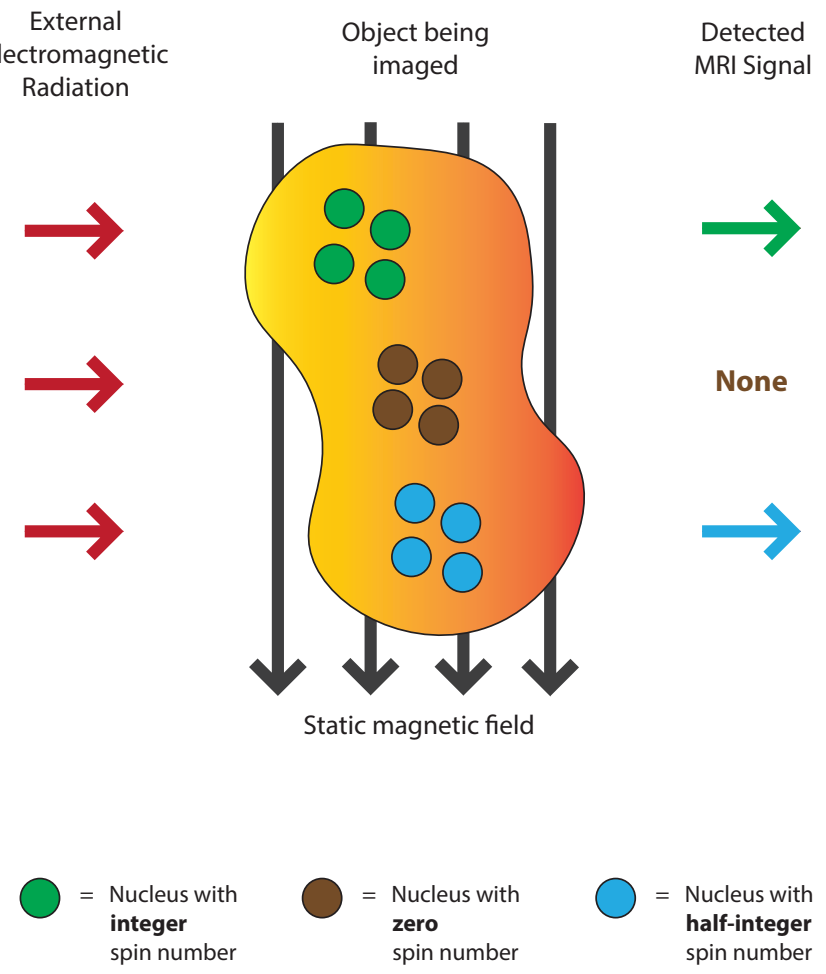

Figure 1.1: An illustration of the basic elements of an NMR experiment, and the required properties of the imaged object.

nal magnetic field, a much more interesting behavior arises. The magnetic moments of the nuclei start to align with the external magnetic field. This is illustrated in Figure 1.2.

However, the alignment of a given nucleus occurs in one of two opposite directions: alignment that is parallel to the external field, and another that is anti-parallel to the field. That is, the magnetic moment can be either one of two opposite vectors. The number of nuclei which are oriented in each of the two opposite directions is not equal. The parallel direction is in a lower energy state compared to the anti-parallel direction. As such, there are slightly more nuclei assuming the parallel direction. Fortunately, even this slight difference results in a detectable net magnetic moment, especially considering the fact that the scale of interest in imaged tissues contains hundreds of billions of nuclei.

When the magnetic moments of the nuclei begin to align with the external field, a rotational force is at work during this alignment. Due to angular momentum, this rotational force causes the magnetic moment to resonate back and forth around the new external field axis. The angular frequency of this periodic rotation is given by *Larmor's equation*, one of the cornerstone equations in magnetic resonance imaging

$$\omega = \gamma B_0 \tag{1.1}$$

where ω is the Larmor frequency: the angular frequency of resonance (also referred to as frequency of precession) in units of radians, B_0 is the strength of the external magnetic field in units of Tesla (T), and γ is a material property constant referred to as the *gyromagnetic ratio*. For hydrogen, $\gamma = 42.57$ MHz/T.

As long as the external B_0 field is on, the precession continues without dampening as there are no frictional forces. Accordingly, the external field B_0 results in a steady state condition where the magnetic moments

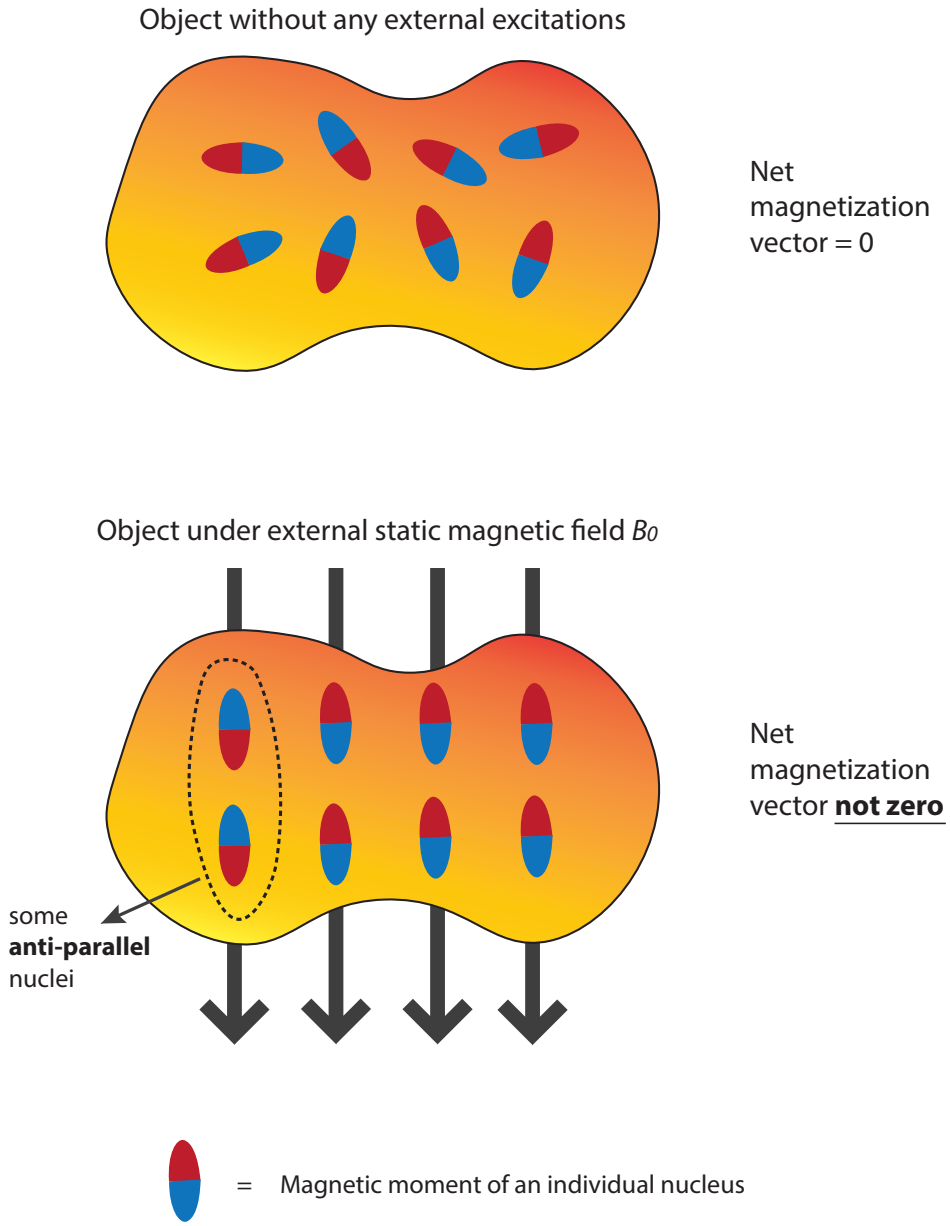

Figure 1.2: Alignment of magnetic moments under an external static magnetic filed B_0. Parallel nuclei are more common and favorable due to lower energy state compared to anti-parallel.

are of known precession frequency. Even though the individual moments are steadily rotating about the z-axis (the axis of B_0), the net magnetization vector has a steady orientation parallel to the z-axis, and we write $M_z = M_0$. However, because the individual precessing moments are not in-phase, the transverse component $M_{x,y}$ is zero.

Therefore, the net magnetization vector $\mathbf{M}$ when B_0 is applied is given by the steady time-invariant expression

$$\mathbf{M} = M_0\hat{\mathbf{z}} \tag{1.2}$$

However, in order to induce a signal in a detector coil, a time-varying magnetic field is necessary (based on Faraday's induction). Accordingly, the individual magnetic moments need to be perturbed and the resultant perturbed net magnetization creates a signal in the detector.

The perturbation of the magnetic moments is achieved through an additional external field that is time-varying, referred to as the B_1 or RF pulse. Because the magnetic moments are precessing at a frequency of ω, the applied B_1 pulse also needs to have a frequency of ω in order to magnetically couple (i.e. to resonate) with the precessing moments. The direction of the B_1 pulse is perpendicular to B_0. Assuming that B_0 is oriented along the z-axis and B_1 along the x-axis, the result of applying the pulse is to tip the orientation of the net magnetization from the longitudinal z-axis towards the transverse x,y-plane. The net change in orientation is referred to as the *tip angle*.

The tip angle depends on the magnitude and duration of the RF pulse. Longer and more intense pulses yield larger tip angles. The intensity is typically much smaller than the static B_0, because magnetic resonance coupling allows for cumulative perturbation, and therefore a small pulse applied for a given amount of time can accumulate and yield sufficient perturbation.

Under this process of tipping the net magnetization vector, if a detector

coil is positioned such that its axis lies in the transverse (x,y) plane, an AC voltage signal oscillating at the Larmor frequency is generated in the coil based on Faraday's induction. This detected signal is the MR signal. This detection provides an external readout of the changes in net magnetization occurring within tissue.

However, the detected MR signal corresponds to all precessing nuclei. In other words, the detected signal is a combination of all magnetic moments from all nuclei regardless of their location. So far, the process does not specify the location from which the detected signal is generated. This issue must be addressed because, after all, the association of an MR signal with a location is a necessary information in order to form an image.

The main approach of establishing spatial information in an MR signal is the use of magnetic field *gradients*. Such field gradients have spatially varying intensities within the region of interest in the imaged object. The gradient is typically linear with relatively small intensity variation in units of milliTesla per meter.

Applying a gradient to a region of interest results in having a linearly varying B_0 within the region. The result of a linearly varying B_0 is a linearly varying precession frequency ω. That is, each location is now associated with a particular frequency. This is in contrast to the initial case without gradients, where all locations had the same frequency and therefore no spatial information was available.

However, it is important to note that, so far, the applied gradient only encodes spatial information along one direction only. That is, if we are to imagine the region of interest to be an image slice, the applied gradient so far only provides information about which image column the signal is coming from. Therefore, an additional gradient is required in order to fully encode the spatial information in an image slice.

While the first spatial direction was determined using frequency-encoding,

information about the second direction is defined through *phase-encoding* direction. This is exactly similar to frequency-encoding, in the sense that it is also based on a spatially varying field. However, it is applied in a second direction that is perpendicular to the frequency-encoding direction. The result of applying the phase-encoding gradient is that the precession phases are now linearly varying along the new direction. The effect of the gradient is to encode the phase information in the detected signal, hence the name "phase-encoding gradient".

With the frequency-encoding direction dividing the image into columns and the phase-encoding direction dividing the image into rows, the spatial unit at the intersection of both gradients is referred to as a *voxel*. The x-dimension of the voxel is specified from the frequency information and its y-dimension is specified from the phase information. The intensity of the voxel is determined through Fourier transform. After the application of both gradients, the detected MR signal in the coil is composed of different frequencies at different phases. An image is then formed by taking the Fourier transform of the detected MR signal, which result in phase and magnitude information of the various frequency components in the signal.

So far, the applied sequence of gradients can provide information about an image slice. However, no aspect of the process determines *which* slice should be imaged within a region of interest. This part of the imaging process, which enables the acquisition of different slices that can span a 3D volume, is referred to as the *slice-select direction* encoding.

A common method of encoding the slice-select direction involves a two step process that applies an RF pulse on top of a linear gradient. First, a linear gradient is created along the slice-select direction (i.e. the direction perpendicular to the slices). Then, instead of applying an RF pulse that resonates with the precessing moments produced by B_0 as before, the RF pulse is designed to have frequency components that correspond to only a small

range that lies within the applied gradient. The center frequency of the RF pulse and the bandwidth of the frequency range it contains determine the position and thickness of the selected slice, respectively.

Therefore, the creation of an MRI volume involves the application of all three gradient directions. The sequence in which the gradients and pulses are applied is referred to as the *pulse sequence*, and the process and steps that have been described in this section represent the main parts of a pulse sequence.

1.3 Diffusion Modeling

The purpose of diffusion modeling is to fit the DW-MRI dataset into a model that enables extraction of various useful information about the microstructure of the brain. The diffusivity information in a DW-MRI dataset is measured using the pulse sequence described in the previous section.

One of the earliest and most common approaches of modeling diffusion information is diffusion tensor imaging (DTI) [9]. DTI uses at least six diffusion weighted volumes to create a tensor that describes the diffusion of water within each voxel.

Accordingly, the cornerstone element in DTI is the diffusion tensor. This can be expressed in a matrix form as follows

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \tag{1.3}$$

which, due to the symmetry of diffusion (i.e. equality between both sides the diagonal; $D_{xy} = D_{yx}$, etc), can be fully determined using only six of the quantities in the matrix. The diagonal elements represent the diffusivities of the three axes. The off-diagonal elements are the correlation (or covariance) between any two axes.

Thus, **D** represents a diffusion covariance matrix. The eigenvectors of **D** represent the axes of an ellipsoid representing diffusion in the voxel. The eigenvalues (λ_1, λ_2, λ_3) of **D** represent the three diffusivities along the three axes of the ellipsoid. The diffusion tensor orientation is taken to be the orientation of the principal eigenvector, which is the eigenvector that corresponds to the largest eigenvalue.

Another important property of the diffusion tensor that is based on eigenvalues is the *fractional anisotropy* (FA) [19], defined as follows

$$\mathrm{FA} = \sqrt{\frac{3}{2}\left(\frac{(\lambda_1 - \mathrm{ADC})^2 + (\lambda_2 - \mathrm{ADC})^2 + (\lambda_3 - \mathrm{ADC})^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}\right)} \tag{1.4}$$

where ADC is the *apparent diffusion coefficient* (also referred to as mean diffusivity), and is proportial to the tensor trace: $\mathrm{ADC} = (D_{xx} + D_{yy} + D_{zz})/3$, which is equivalent to the average of eigenvalues. ADC is also sometimes used to refer to each direction separately, in which case the ADCs correspond to the eigenvalues.

FA is a normalized quantitative indicator of the degree of anisotropy in a voxel. A near-zero FA value indicates an isotropic diffusion within the voxel. That is, the diffusivities along all three orthogonal axes are similar. This provides an indication that water is diffusing equally along all directions, and therefore the local tissue structure is unrestrictive and non-directional. On the other hand, FA values closer to one indicate a highly anisotropic diffusion. This is a diffusion that is very large along one direction only, and very small in the other two directions. As such, high FA results from a voxel which has a local tissue structure that is restricting water diffusion along one direction.

1.4 Limitations and Alternative Imaging Schemes

While DTI offers useful insight into brain structure through the use of tensors, it has a number of limitations. The accurate construction of fiber tracts is important in gaining insights into brain function since fiber tracts act as the infrastructure enabling communication between brain regions [39]. However, accuracy of the reconstructed fiber tracts is often hampered by the often low resolution of DTI data. A voxel can thus comprise several distinct fiber bundles with differing orientations, leading to *partial volume effect* [4]. At such locations, diffusion information typically becomes ambiguous, and tractography is often falsely terminated.

Fortunately, the dwMRI acquisition process has a multitude of parameters that affect the quality of the final volume data. These include factors such as magnetic field strength, voxel size, diffusion orientations, among other factors. This has resulted in a significant body of research devoted to optimizing such factors with the aim of achieving more accurate fiber estimation.

The following sections will review some of the key alternative imaging schemes (mainly Diffusion Spectrum Imaging - DSI, and High Angular Resolution Imaging - HARDI) that are addressing the limitations of DTI.

1.4.1 Diffusion Spectrum Imaging

One approach that has been proposed to capture more detailed diffusion information is diffusion spectrum imaging (DSI) [93]. DSI acquires full diffusion information in a substantial region of the q-space, which is the 3D space representing the extent and orientation of spin diffusions along the three axes of motion (x,y,z). This is in contrast with the more common DTI imaging scheme, which does not aim to fully sample the q-space.

Accordingly, in order to sample a large region of the q-space, the im-

age acquisition in DSI is different from DTI. A major difference between the two techniques is that DSI adds additional gradients during imaging. That is, in addition to the typical slice selection, frequency-encoding, and phase-encoding gradients, three new gradients are added to fully encode diffusion in a 3D Cartesian grid comprising many diffusion magnitudes and orientations. In other words, DSI is essentially a fully 6D imaging modality, which samples both the k-space (spatial position sampling) and q-space (spin displacement sampling) simultaneously.

As a result of the additional three dimensions of q-space, DSI provides an explicit representation of diffusion. In other words, DSI does not require diffusion modeling. This is because at each voxel, the q-space volume has been shown [93] to be a direct representation (through a Fourier transform) of the spin-displacement PDF (probability density function) at each voxel. This is in contrast with DTI, which can not be performed without modeling diffusion using tensors.

Experimental studies comparing DSI relative to DTI have shown improved detection of fiber crossings in human and rat brains [55, 94]. In a number of ROIs, the fibers reconstructed from DSI have shown better resemblance to known anatomical pathways in the brain.

Despite the improvement in fiber crossing detection, DSI has a significant drawback. The sampling of q-space requires a significant amount of data, with typically a 3D grid of 500 q-values used in scanning, representing DW gradients over a multitude of orientations and magnitudes. This results in a much longer acquisition time. While a typical DTI acquisition takes about 3-5 minutes, DSI acquisition time is usually around 40 minutes [16, 57], and is typically performed with a smaller number of repetitions compared to DTI in order to reduce scan time, and therefore the SNR is generally lower. Accordingly, this is a significant drawback that prevents DSI from being a practical imaging scheme.

1.4.2 High Angular Resolution Diffusion Imaging

In an effort to reduce the significant burden of sampling a 3D q-space grid in DSI, an alternative approach that uses only angular samples was proposed [5, 27, 88]. This approach, which came to be known as *high angular resolution diffusion-weighted imaging* (HARDI), only samples a spherical subset of the diffusion space, which is referred to as a *diffusion shell*. That is, the acquisition is simplified to a single (or, sometimes, a few [1, 21]) diffusion shells instead of fully sampling the entire 3D diffusion space as in DSI.

By restricting the acquisition to a high-resolution spherical shell instead of a full Cartesian grid, HARDI simplifies the acquisition while maintaining some of the information provided by DSI. This simplification is further supported by the fact that the typical subsequent processing pipeline is based on tracking the locally estimated orientations. Accordingly, HARDI's simplification strategy of focusing more on angular diffusion data relative to radial data can be argued to be an efficient strategy.

The choice of the magnitude(s) of shells is a trade off between the benefits of HARDI (few shells and orientations) and DSI (full Cartesian grid). However, it is known that higher magnitudes and more orientations yield better performance [74].

However, because full information on the q-space is no longer available, HARDI requires a diffusion modeling step, which was not required in the case of DSI. Several modeling approaches have been proposed in literature. However, one of the most common approaches of reconstructing diffusion ODF from HARDI data is *Q-ball Imaging* (QBI) [87]. An attractive feature of QBI is that it is a model-free method. In QBI, the measured spherical diffusion signal is directly used to reconstruct the diffusion ODF. This is performed through the Funk-Radon transform, which extends the planar Radon transform to spherical tomographic reconstruction. QBI has been

shown to yield better reconstruction of fiber crossings compared to DTI [87, 89].

Despite the improved reconstruction afforded by QBI, the technique has some drawbacks. QBI has been shown to require high diffusion-weighting factors (b-values) in order to be able to resolve fiber crossings [48, 86], with suggested values of about $b = 3000 - 4000 s mm^{-2}$. In comparison, typical DTI values are around $b \sim 1000 s mm^{-2}$. Due to the higher b-values in QBI and the denser gradient directions sampling, the resultant SNR can be very low [62]. Furthermore, the higher number of gradient orientations (~ 60 to few hundreds) results in a significant acquisition time of around ~ 30 minutes to a few hours [48], which hampers practical clinical use.

1.5 Thesis Objectives and Proposed Approach

The aim of this thesis is to propose an approach for increasing the spatial resolution of brain dwMRI data. Increasing the resolution will help in enabling more detailed extraction of information from dwMRI data, and therefore help improve the estimation of brain fiber structures. In this section, we describe the main objectives that this work is aiming to address.

Recently, a powerful approach of enhancing the resolution of images using dictionary learning has been shown to yield good performance in natural images [46, 97]. Dictionary learning is a process in which a signal or image is represented using as few learned basis functions as possible. The details of this process are discussed at length in Chapter 3.

In the dictionary learning based super-resolution process, two dictionaries are created: a high-resolution dictionary and a corresponding low-resolution dictionary. Through these joint dictionaries, a high-resolution image can be generated from a new unseen low-resolution image.

In a similar fashion, we propose a processing pipeline in this thesis that

is built on top of the same joint-dictionary learning approach and extend it to multi-shell dwMRI data. We chose to create a pipeline that adopts this approach in order to be able to create a resolution enhancement of dwMRI without resorting to acquisition modifications. The details of the proposed processing pipeline are described in Chapter 3.

An objective of the proposed approach is to not require modifications to the acquisition scheme. This stems from the fact that most of the proposed alternatives to DTI require acquisition methods which require a very long imaging time to obtain acceptable image quality, hampering its utilization in routine clinical use as explained earlier in section 1.4.

To this end, we aim to propose an approach that should not require modifications to the dwMRI acquisition process. This ensures that the proposed processing pipeline maintains its applicability in a wide variety of clinical and research settings. This also allows the proposed approach to be used on *legacy data* in pre-existing clinical or research data sets. These objectives can not be attained using other approaches that are based on acquisition modifications

Another objective of the proposed approach is to be independent from the choice of diffusion modeling. In other words, the proposed approach is not attempting to change or improve the diffusion model. This ensures that the proposed approach can be used with whatever diffusion model used in various acquisitions. It is also not a tractography method, but rather can be used as an input to tractography. This makes it suitable for use with any tractography method preferred by the end user.

Furthermore, as a result of the preceding features, the proposed approach should be modular and flexible. In other words, it can be used with any other method that aims to enhance structural information. For instance, it can be used as *an additional* step after other approaches that use acquisition modifications.

1.6 Thesis Organization

The rest of this thesis is organized as follows. In Chapter 2, we examine key related literature aiming to achieve similar objectives of increasing dwMRI data resolution. We categorize them into main categories and explore the advantages and disadvantages of each category.

In Chapter 3, the proposed framework is described. We examine in details the rationale and design of each block in the proposed framework, and determine the various parameters of the methods in order to yield good performance. We begin by describing the preconditioning steps employed to improve the condition of the coding matrix, after which we determine the parameters of the preconditioning process, followed by determining the parameters of dictionary learning.

In Chapter 4, we describe the proposed validation methodology. Chapter 5 presents quantitative and qualitative results examining the performance of the proposed framework, and also examines the dependency of performance on variations in datasets, as well as comparisons against other methods.

1.7 Thesis Contributions

The following is a brief overview of the contributions of this thesis. This will be explained in detail in the following chapters.

We propose a dwMRI data processing pipeline (built on a dictionary learning approach) that enhances the resolution of dwMRI after the data has been acquired. The pipeline does not require modifications to the dwMRI acquisition process, and therefore is more practical in clinical conditions where simple acquisition methods are typically used.

Due to the absence of acquisition requirements, the pipeline can be read-

ily used with existing legacy databases of dwMRI data. This can be useful when trying to utilize or reuse datasets which has been acquired with older acquisition technologies. Absence of acquisition requirements also makes it usable with different imaging schemes, such as DTI or HARDI (single, and multiple-shell). Some of the existing super-resolution approaches are specific to HARDI or DTI.

We also note that the proposed pipeline does not require repeated dwMRI acquisitions, unlike classical shifting-based super-resolution methods. Another contribution is that the pipeline performs a resolution enhancement directly on raw dwMRI data, resulting in a model-free pipeline.

This lack of required modeling is useful because it does not restrict the end user to a particular model, which gives the freedom of using different diffusion models in different research or clinical situations, depending on the specific research conditions or objectives. Finally, we note that the model-independence also allows for *cumulative* enhancements, in which any other diffusion-specific or tractography-specific enhancements can be readily applied on top of the pipeline.

Chapter 2

Related Literature

2.1 Overview

Obtaining structural information about the brain in a non-invasive fashion is one of the major applications of brain dwMRI, among other applications such as using diffusion anisotropy to diagnose brain infarction [95] as well as brain development and aging [72].

This structural information takes the form of a set of connected fiber segments (referred to as tracts) that are reconstructed from dwMRI. The fiber tracts in turn give raise to a network of inter-connected spatial regions in the brain.

The cornerstone of obtaining structural reconstruction is to perform a *tractography* on the dwMRI data. In its most common form, tractography generates a set of fiber tracts by connecting the largest eigenvectors of the diffusion tensor at each voxel, starting from multiple seed points. This approach is referred to as *deterministic streamline tractography*. Many other tractography approaches have been proposed, such as methods based on globally-consistent reconstructions as well as probabilistic reconstructions.

The accurate construction of fiber tracts is important in gaining insights into brain function since fiber tracts act as the infrastructure enabling communication between brain regions [39]. However, accuracy of the reconstructed fiber tracts is often hampered by the inherently low resolution of dwMRI data. Currently achievable spatial dwMRI resolution is around 2

mm^3, while the actual neuronal fiber diameter is on the order of 1 μm. A voxel can thus comprise several distinct fiber bundles with differing orientations, leading to *partial volume effect* [4]. At such locations, diffusion information typically becomes ambiguous, and tractography is often falsely terminated.

Several methods have been explored in literature to address the partial-volume effect. The common goal of these methods is to reduce the ambiguity present in the imaged voxels. While there is a wide variety of approaches and strategies proposed in literature, we grouped them into two major categories: shifting-based dwMRI super-resolution, and model-based dwMRI super-resolution.

At this point it is important to note that the enhancement of structural reconstruction can also be achieved via improvements in tractography algorithms. While research in this field is very active, the goal of the current investigation is to explore structural enhancement approaches that are independent of tractography. The rationale is that such approaches are beneficial because any structural enhancements they offer may then be accumulated *in addition* to improvements in tractography algorithms.

2.2 Shifting-based dwMRI Super-resolution

One of the earlier attempts to improve resolution of dwMRI was through the use of multiple shifted acquisition, in which the imaging process acquires multiple volumes, each representing a slightly shifted region of physical space. This is followed by a subsequent super-resolution reconstruction method [10, 41]. The main difference between this approach and more traditional Q-space imaging enhancements is that the latter aims to resolve more details through a modified imaging acquisition, while super-resolution entails a post-imaging processing approach.

As implied by the name, super-resolution is a method of increasing the resolution of an image. The central concept of super-resolution is the reconstruction of high-resolution details from multiple lower-resolution image acquisitions. The main benefit of super-resolution is providing a higher resolution image without requiring modifications to the imaging hardware or optics. This can be beneficial in situations where such hardware modifications may be difficult or impractical.

The rationale for super-resolution is based on an information fusion process. The super-resolution process assumes that, if multiple low-resolution images of a scene are available, and if such images were related to each other by relatively small shifts, then a higher resolution image can be obtained by fusing the information from the set of shifted low resolution images into a single high resolution image [90].

2.2.1 Translational Field of View Shifting

One of the first of investigations in translational-shifting super-resolution of dwMRI was by Peled and Yeshurun [71]. In this work, super-resolution reconstruction was directly applied to multiple acquisitions of dwMRI data. Specifically, each slice was acquired eight times, with each acquisition at a subvoxel spatial shift relative to the first acquisition. This was achieved by changing the FOV (field of view) in the frequency-encoding and the phase-encoding directions (i.e. within the in-plane directions in a multi-slice acquisition).

However, this approach has been proven later to have a fundamental problem [75]. This is because spatial sub-voxel shifts in the FOV in the in-plane directions simply correspond to linear phase modulations in the k-space. This means that, the multiple acquisitions correspond to the same k-space points, and no new points or new information is acquired. Any observable improvement may therefore be attributed to increased SNR be-

cause of the use of more averaging acquisitions, which is already a standard practice in most acquisitions.

Accordingly, subsequent research efforts have focused on super-resolving the through-plane (i.e. the slice-select) direction of the acquisition. One approach [32] performs sub-voxel shifts in the FOV along the slice-select direction.

2.2.2 Orientational Field of View Shifting

An alternative approach to translational sub-voxel shifting has been proposed based on orientational FOV shifting [76, 79]. For each low resolution acquisition, the FOV was rotated around the frequency-encoding direction, resulting in a series of rotations (a total of six equidistant orientations were acquired, 30^o apart).

The orientational FOV shifting approach has been shown to yield better resolution enhancement relative to the sub-voxel spatial shift [80] for the same number of low-resolution acquisitions. The methods adopting this approach are motivated by the rationale that orientational shifting would arguably result in a more efficient sampling of the k-space for the same number of shifting. This has not been explored quantitatively in literature. The difference between the two FOV shifting approaches is illustrated in Figure 2.1.

Despite the reported resolution enhancements using shift-based super-resolution, the process has a number of limitations. First of all, it is very important to note that the vast majority of super-resolution literature in MRI have focused on standard (non-diffusion weighted) images. Its use in DTI (or other Q-space based modalities for that matter) received little attention. In most super-resolution studies, a total of around 6-8 FOV shifts were acquired (whether spatial or rotational), which were based on empirical examination and no theoretical limit has been investigated. Strictly speaking,

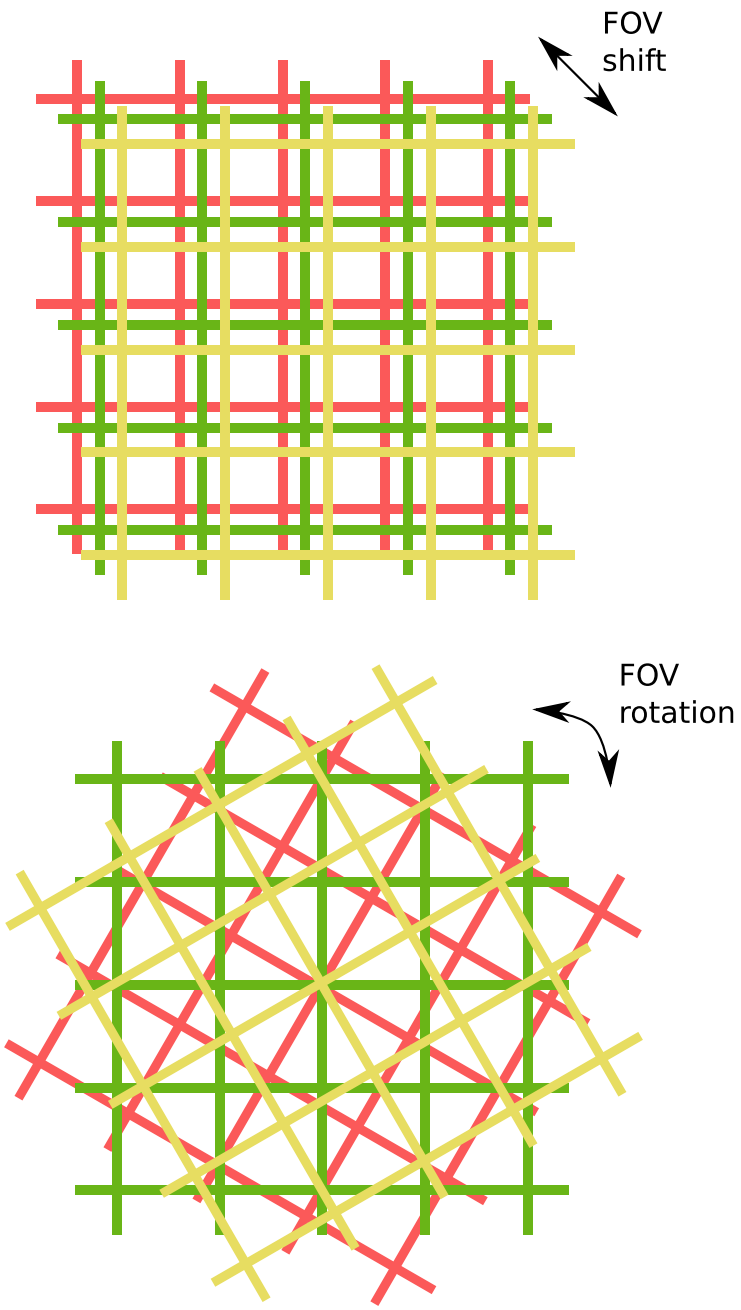

Figure 2.1: Illustration of the spatial sub-voxel FOV shifting super-resolution in comparison to the orientational shifting super-resolution.

these shifted acquisitions are in-addition to the few additional repeated acquisitions performed to average the signal and boost the SNR. Accordingly, around 10 acquisitions can be expected to be performed for each gradient direction in a DTI acquisition. And with the recommended optimal number of gradients being around 30 [43], this results in the requirement of having around 300 acquisitions for super-resolution. Recently, one of the first DTI studies [73] used only 12 gradients (total acquisition time of around 9 minutes) with no repetitions, which is much lower than the recommended number of gradients. More research needs to be done in investigating the use of shift-based super-resolution for DTI.

Another limitation of shifting-based super-resolution approaches is that it is not possible to super-resolve information within the between-planes direction, thus limiting any possible resolution enhancement to the two other dimensions only, due to limitations imposed by the encoding scheme [75].

2.3 Model-based dwMRI Super-resolution

An alternative class of super-resolution approaches may be identified as involving the construction of more complex models of dwMRI data. Specifically, methods belonging to this approach focus on increasing the resolved details in various diffusion models. This then enables exploring the reconstructed neural structures at a greater resolution.

This approach is different from the methods described in the previous section in the sense that it does not require multiple shifted acquisitions. However, due to the often mutual dependency between model and acquisition, some of these methods may also include certain acquisition requirements. This section presents a review of key literature adopting this model-based super-resolution approach.

2.3.1 Super-resolved Diffusion Tensor

A model-based approach aiming to super-resolve diffusion tensors has been proposed [35, 66]. This approach essentially proposes an alternative diffusion tensor reconstruction method such that the diffusion tensors are constructed at a higher resolution grid.

The basic framework is based on expressing the diffusion signal in terms of a diffusion tensor at a higher resolution. This is in contrast with the standard method of diffusion tensor construction, where the diffusion signal is expressed in terms of a diffusion tensor at the same resolution (i.e. same grid as dwMRI data). This can be expressed in a generic form as follows

$$S^{\mathrm{LR}}\left(x_i^{\mathrm{LR}}\right) = f\left(\mathbf{D}^{\mathrm{HR}}\left(x_j^{\mathrm{HR}}\right)\right) \tag{2.1}$$

where S^{LR} is the original (low-resolution) diffusion signal, x_i^{LR} represents the original (low-resolution) grid, f indicates a function of $\mathbf{D}^{\mathrm{HR}}$, which is the diffusion tensor at the higher resolution grid x_j^{HR}. Further details on equation parameters can be found in [35].

Accordingly, a relation is established between the low (original) resolution dwMRI data and high-resolution tensors. These tensors are then constructed using an inverse problem approach. An energy function is therefore created, and then minimized to generate the tensors.

While better tensor resolution has been demonstrated [35], a major limitation of this approach is that it is inherently restricted to using diffusion tensors as the model, which has been shown to be suboptimal in modeling diffusion in complex neural fiber structures [8]. Therefore, this approach can not be used with other diffusion models such as ODF (orientational density function), for instance. Its applicability is limited due to the requirement of a tensor model.

2.3.2 Super-resolved Spherical Deconvolution

Another modeling method has been proposed which aims to construct a high resolution orientational probability density function (ODF) based on a spherical deconvolution of the diffusion-weighted data [84, 85]. As implied by the name, the ODF assigns a probability for the existence of a fiber at various orientations on a sphere.

In this model, the diffusion-weighted signals (as measured on a diffusion shell sphere) are modeled as the convolution of a certain transfer function with an ODF

$$S(\theta, \phi) = F(\theta, \phi) * R(\theta) \tag{2.2}$$

where θ, ϕ are the elevation and azimuthal angles in a sphere and $S(\theta, \phi)$ is the measured diffusion signal at a voxel at the orientation (θ, ϕ). $F(\theta, \phi)$ is the desired unknown ODF, and $R(\theta)$ is the transfer function. The transfer function would typically be determined from regions in the data where the ODF is anatomically known to be a single coherent orientation of fibers.

Therefore, given a diffusion-weighted signal $S(\theta, \phi)$ and the transfer function $R(\theta)$, a deconvolution process would then enable the estimation of the ODF $F(\theta, \phi)$. Since the functions involved in the convolution are spherical, the deconvolution is also performed over a sphere [38].

The super-resolution approach used by these methods is based on performing a spherical deconvolution to estimate the ODF at resolutions higher than what is present in the measured data $S(\theta, \phi)$ [84]. This is achieved by introducing non-negativity constraints on the ODF, which stems from the anatomical impossibility of having negative fiber density. This has the dual benefit of reducing background noise which may cause the ODF to take negative values, and subsequently allows a finer estimation of the ODF due to the reduced background noise.

While this approach has demonstrated higher resolution in the gener-

ated ODFs [84], it is specifically restricted to choosing ODFs as a model. This makes it unusable with the wide variety of other modeling methods. Furthermore, the proposed method requires the use of HARDI data. As such, this super-resolution approach can not be used with the more common DTI data, which is also often the type of data of many legacy dwMRI databases.

Chapter 3

Proposed Framework

3.1 Overview

The proposed framework for enhancing the resolution of dwMRI differs from the existing shifting-based and model-based dwMRI super-resolution approaches discussed in Chapter 2. In the proposed framework, there is no requirement for a modified acquisition sequence, nor a choice of a particular diffusion model. The proposed framework makes no assumptions about acquisition or modeling.

In this chapter, we will describe the framework in detail. The framework has two components: a processing pipeline and a validation methodology. The processing pipeline is described in section 3.2, and the validation methodology is described in Chapter 4.

3.2 Processing Pipeline

Recently, a powerful approach of super-resolving data using a dictionary learning approach has been shown to yield good performance in super-resolving natural images [46, 97]. Dictionary learning is a process in which a signal or image is represented using as few basis functions as possible. The key concept of dictionary learning is that the basis functions are learned from the data, rather than being constructed from a generating function.

In the process of super-resolution via dictionary learning, two dictionaries are typically created: a high-resolution dictionary and a corresponding

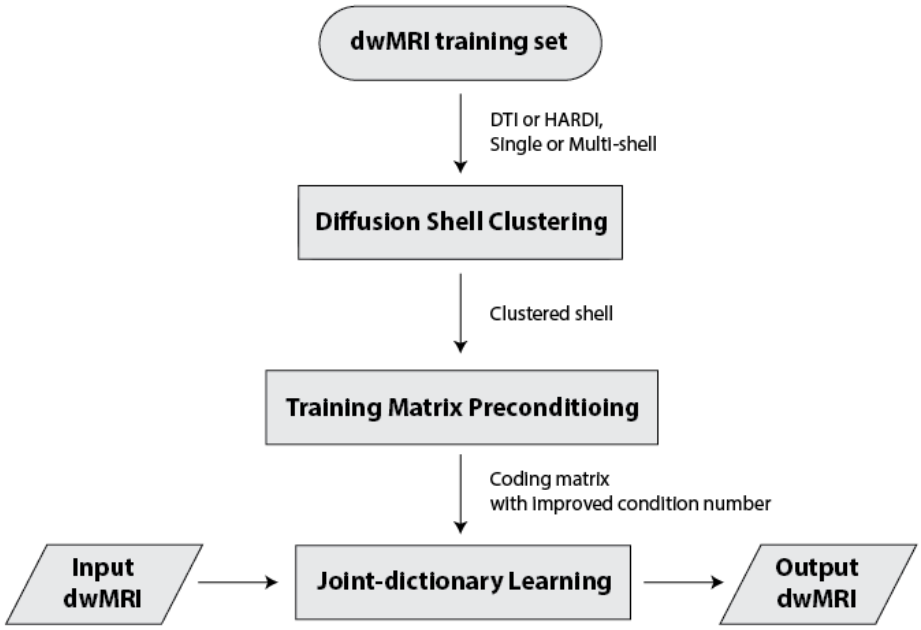

Figure 3.1: An overview of the proposed processing pipeline.

low-resolution dictionary. Through these joint dictionaries, a high-resolution image can be generated from a new, previously unseen, low-resolution image.

In a similar fashion, we propose a processing pipeline in this thesis that is built on top of the same joint-dictionary learning approach and extend it to multi-shell dwMRI data. We chose to create a pipeline that adopts this approach in order to be able to create a resolution enhancement of dwMRI without resorting to acquisition modifications, which was the case in the alternative super-resolution methods described earlier in Chapter 2.

The proposed processing pipeline is shown in Figure 3.1. The process starts with selecting a dwMRI training dataset. The data is then clustered into a single diffusion shell. We recall from Chapter 1 that a diffusion shell is simply a spherical subset of the diffusion space instead of the entire 3D diffusion space. This will be discussed in more details in section 3.2.1.

The clustering is then followed by a preconditioning process that aims to improve the condition number of a coding matrix. Finally, the process terminates with a dictionary learning step applied on the dwMRI data.

3.2.1 Diffusion Shell Clustering

In general, dwMRI volumes may be described as residing on diffusion shells. Each diffusion shell describes a certain strength of diffusion weighting. Subsequently, each acquired diffusion direction represents a sample on the corresponding diffusion shell. At the limit, as more directions are acquired in each shell and as inter-shell gaps are reduced, the acquired data set approaches that of a DSI acquisition. Figure 3.2 shows an illustration of the shell structure in a typical dwMRI dataset.

Historically, most acquisitions were typically limited to single shells. This was due to practical limitations in the acquisition hardware and corresponding software pipeline algorithms. Increasingly, this is being gradu-

ally replaced by multi-shell acquisitions. This form of acquisition allows for higher angular contrast between the various diffusion directions [13, 17, 45] and also enables the use of richer diffusion modeling [2, 22, 47].

However, the benefits afforded by multiple-shells are offset by the more elaborate sampling scheme and the resulting substantial increase in the size of the dataset. While the design of multi-shell acquisitions is still an active field of research, some studies [12, 81] have provided suggestions that an optimal acquisition can be achieved using 3 shells and around 160 - 280 total measurements (across all shells).

Such multi-shell acquisitions posit a number of difficulties for the proposed resolution enhancement framework. First, there is a practicality issue of loading and learning on the large number of volumes in a multi-shell acquisition, as can be noted by observing the large number of dots in Figure 3.2. For instance, a typical acquisition following the aforementioned scheme (for example, a typical subject data in the Human Connectome Project – HCP) requires around ~ 8 GB of memory space, not including any other software or system resources requirements. While this is gradually becoming a non-issue on some modern computers with ample resources, it may still present a practical limitation on processing multi-shell acquisitions for a large majority of users (which was the case on the machines used to conduct this research).

Another, more fundamental, issue of training on multi-shell acquisitions is the increased ill-conditioning of the training matrix. This is due to finer angular sampling in typical multi-shell acquisitions which increases the odds of having highly-similar patches. The resulting ill-conditioned matrix prevents the creation of a joint hi-res/low-res dictionary, which is necessary for the proposed resolution enhancement framework, as will be described in more details in section 3.2.2.

We propose to address these issues through a clustering approach. More

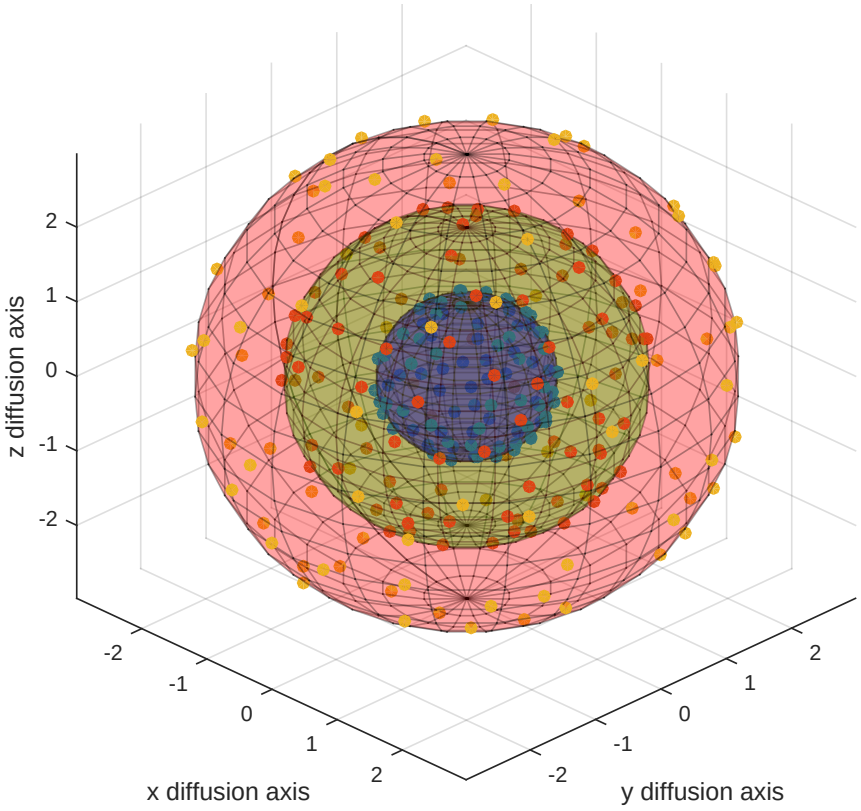

Figure 3.2: Illustration of the diffusion shells structure in a typical dwMRI data. Each shell represents a certain strength of diffusion weighting. Each dot represents a 3D diffusion volume.

specifically, we propose to cluster the measurements across all acquired shells into a set of representative dwMRI volumes. By design, this proposed *diffusion shell clustering* should be able to condense a multi-shell acquisition into a relatively smaller and *heterogeneous* single shell that would still capture the bulk of information present in the multi-shell dataset. While such a clustering approach, by design, reduces the amount of information, we adopt this approach to reduce the computational load and the ill-conditioning of the training matrix as explained earlier.

There are a number of existing clustering methods that are widely used to cluster many types of data. Common methods include K-means clustering, K-medoids, and Gaussian mixture models, to name a few. However, most of these clustering methods require identifying the number of clusters before hand. Identifying the correct number of clusters remains a challenge, and may be different for different acquisitions. This is especially true for multi-shell data, where there can be great variability between one acquisition and another.

At this point, we recall that one of the main objectives of this work is to propose a framework that is independent of the acquisition scheme. As such, we aim for the processing pipeline to be able to handle older and common DTI data, as well as various forms of the richer and more recent multi-shell HARDI data.

In order to achieve this goal, we propose to utilize *affinity propagation* (AP) [28] to perform diffusion shell clustering. In AP, the dataset is iteratively analyzed to generate a set of exemplars: data points that are most representative of their respective clusters. The affinity propagation process begins with a similarity matrix, which indicates how well a certain point serves as an exemplar to other points. The exemplars then minimize the pairwise error or distance between themselves and potential cluster members. We refer the reader to [28] for more details.

Most importantly, the final number of clusters is determined automatically. In addition, there is no initial set of clusters. Instead, all points are candidate exemplars. All points are represented as nodes on a network, and messages are passed between nodes depending on the degree of similarity between the pair, denoted by $s(i,j)$. The iterative propagation of such affinity messages throughout the network results in a final set of clusters and exemplars. Briefly, the iterations initialize and subsequently update two matrices: responsibility matrix $r(i,j)$ indicates suitability of j to be an exemplar for i, and availability matrix $a(i,j)$ indicates how suitable it is for i to choose j as an exemplar. We refer the reader to [28] for algorithm details.

We now show how we will utilize the AP method to implement diffusion shell clustering. We begin by constructing a feature vector ϕ_m for each dwMRI volume (regardless of overall shell structure) to be the standard deviation of DW volumetric patches, as follows:

$$p_{l,n,m} = \mathrm{vec}(\mathbf{P}_{i,j,k}\mathbf{V}_m) \tag{3.1}$$

$$\phi_m = \left(\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(p_{i,1,m} - \overline{p}_{i,1,m}\right)^2}, \ldots, \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(p_{i,n,m} - \overline{p}_{i,n,m}\right)^2} \right) \tag{3.2}$$

where $\mathbf{P}$ is defined as an operator extracting isotropic volumetric patches centered at the voxel (i,j,k), $\mathbf{V}_m$ is the m-th volume in the acquired data, and $p_{l,n,m}$ is the vectorization of the n-th volumetric patch of the m-th volume at the l-th index.

We then build a DW measurements correlation matrix $\mathbf{\Psi}$ from the feature vectors, as follows:

$$\Sigma_{i,j} = \mathrm{cov}(\phi_i, \phi_j) \tag{3.3}$$

$$\mathbf{\Psi} = \mathrm{diag}(\Sigma)^{-\frac{1}{2}}\,\Sigma\,\mathrm{diag}(\Sigma)^{-\frac{1}{2}} \tag{3.4}$$

after which we utilize $\mathbf{\Psi}$ to provide similarity measurements for use during affinity propagation:

$$r(i,j)|_{t=t_0} = \Psi_{i,j} - \max_{j^* \neq j} \left\{ \Psi_{i,j^*} \right\} \tag{3.5}$$

which is then propagated across the nodes of the network as explained earlier in this section. All of the resultant cluster exemplars (regardless of their count, i.e. no thresholding is performed) are then used for the remainder of the learning process.

We recall that the goal of using a clustering approach was to provide a reduction in the computational load and to reduce the occurrence of highly similar atoms. The training matrix preconditioning step of section 3.2.2 performs the rest of this reduction, and hence we opted for a simpler and easier to adjust feature vector. Accordingly, our feature vector ϕ_m has a single parameter: the isotropic patch dimension, d. The choice of this parameter's value is expected to change the correlation matrix $\mathbf{\Psi}$. As such, we need to determine an optimal value of d.

In our case, the optimality condition would be a matrix $\mathbf{\Psi}$ that is as discriminative as possible between different diffusion clusters, while still maintaining high correlation within the same cluster. Accordingly, an optimal value of d should try to maintain as much dispersion as possible in the correlation values in $\mathbf{\Psi}$. In Figures 3.3 to 3.5 we show examples of the resulting $\mathbf{\Psi}$ for different values of d. These example figures were generated from a randomly selected subject in the HCP dataset (the dataset is described in section 5.1).

The results of the effect of d on $\mathbf{\Psi}$ are summarized in Figure 3.6. The figure shows the range of the resultant $\mathbf{\Psi}$ for various values of d for 10 random subjects from the HCP data. We note that the range is monotonically decreasing, and that its maximum occurs at a value of $d = 3$. This gives the largest discriminative power for our chosen feature vector. Accordingly, the

value of $d = 3$ will be used for the remainder of this work. An example of the result of clustering is shown in Figure 3.7. Note how the correlation matrix after clustering clearly shows a clustered grouping of diffusion indices. Each block or group of diffusion indices can now be represented using any index within that cluster.

3.2.2 Training Matrix Preconditioning

Dictionary learning methods are typically highly non convex [3, 59]. Strictly speaking, the optimizations involved in these methods are generally combinatorial in nature:

$$\min_{\alpha} \ \|\alpha\|_0 \ \text{ subject to } \ y = \mathbf{D}\alpha \tag{3.6}$$

for a full-rank $\mathbf{D} \in \mathbb{R}^{n \times m}$ ($n < m$). This has been shown to be an NP-hard problem [65]. An approach that has been used to address this issue is to perform a relaxation from an $l0$ to an $l1$ norm, which results in a convexification of the problem that yields a result that approximates the true sparse coding vector [15]. However, when the sparse coding is combined with a dictionary update (which is a required combination in dictionary learning methods [3, 59]),

$$\min_{\mathbf{D}, \{\alpha_i\}_{i=1}^{M}} \ \sum_{i=1}^{M} \|y_i - \mathbf{D}\alpha_i\|_2^2 \ \text{ subject to } \ \|\alpha_i\|_1 \leq k \tag{3.7}$$

the problem becomes a nested optimization, in which the first optimizes the sparsity of each coding vector α_i for a given $\mathbf{D}$, and the second optimizes over $\mathbf{D}$. The nested nature of the optimization has naturally resulted in algorithms that solve the problem by also alternating between two optimizations [3, 24]. Most of the variations between the algorithms lie in the choice of heuristic for each of the two steps.

However, due to their high non-convexity, such approaches will often fall into local minima or saddle points [3]. In fact, even when assuming a

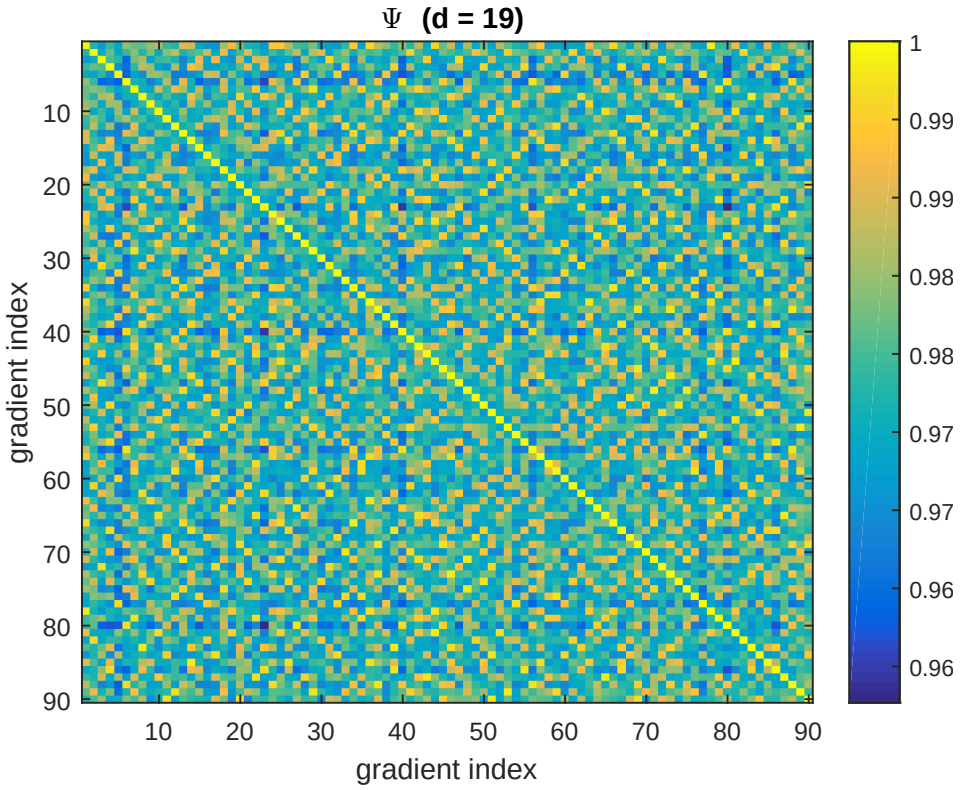

Figure 3.3: An illustration of the resultant correlation matrix $\mathbf{\Psi}$ obtained with a feature-vector volumetric patch dimension $d = 19$.

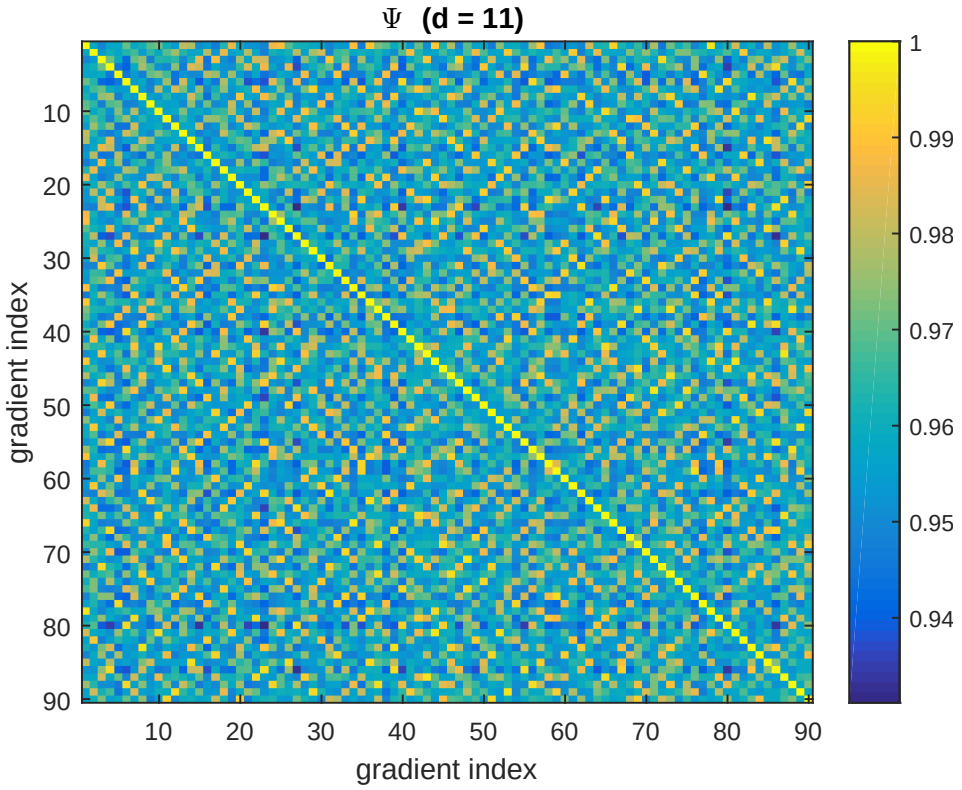

Figure 3.4: An illustration of the resultant correlation matrix $\mathbf{\Psi}$ obtained with a feature-vector volumetric patch dimension $d = 11$.

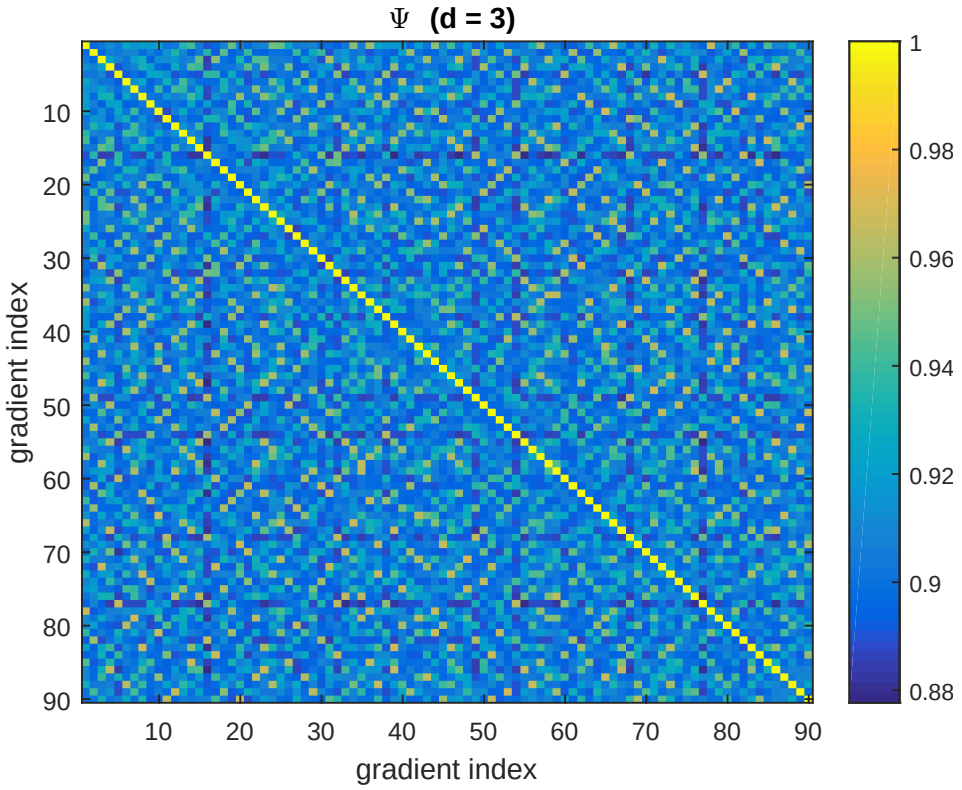

Figure 3.5: An illustration of the resultant correlation matrix $\mathbf{\Psi}$ obtained with a feature-vector volumetric patch dimension $d = 3$.

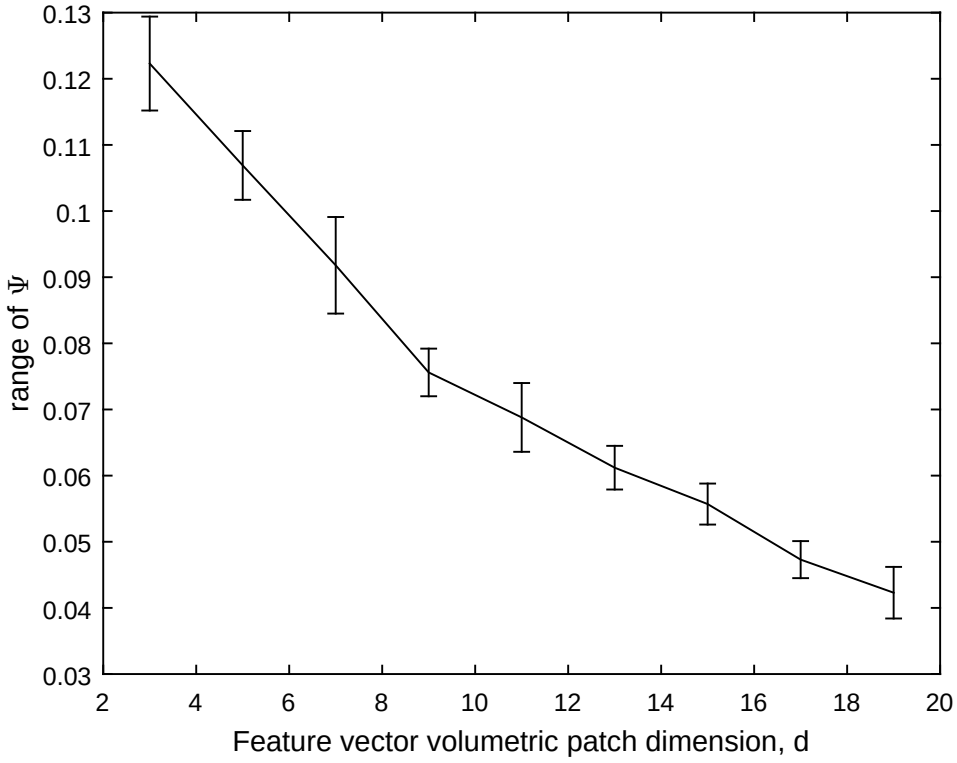

Figure 3.6: The effect of the feature-vector volumetric patch dimension, d, on the range of the resultant correlation matrix $\mathbf{\Psi}$. Larger values are better. Vertical bars show standard deviation across 10 HCP volumes.

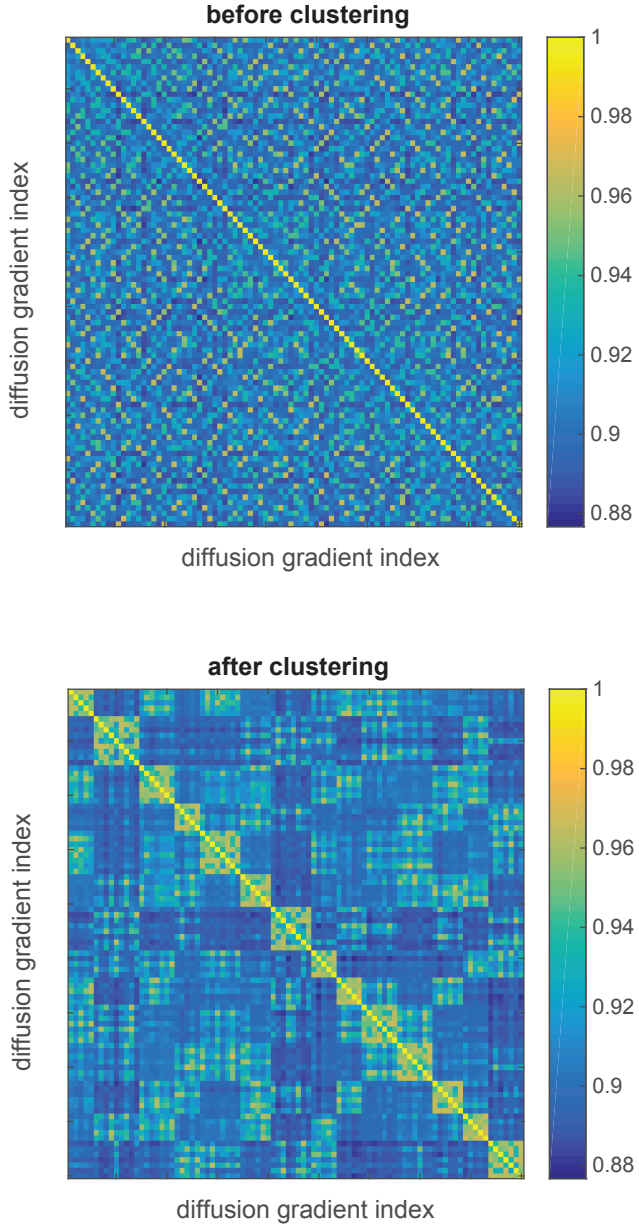

Figure 3.7: An example result of diffusion shell clustering. The upper figure shows the Ψ matrix before clustering, and bottom figure shows the matrix after clustering. Note how the bottom figure clusters the data, therefore provides information on which gradient indices are relatively identical and therefore replaceable by any index within a given cluster.

perfect sparse coding step, only a convergence to a local minimum is guaranteed. This guarantee does not hold when approximate sparse coding methods are used, which is often the case due to the combinatorial nature of solving an exact sparse coding as explained earlier.

As a result, the choice of training matrix initialization will have an effect on the convergence of the training. Different initializations will converge to different local minima. In previous works utilizing dictionary learning methods in image processing, the training matrix was typically initialized to include all patches from the supplied training set [54, 69, 78]. The results obtained from this initialization approach provided acceptable performance in different applications, such as face recognition, image denoising, and remote sensing.

However, this approach presents a challenge when used with dwMRI data. Volumetric patches constructed from dwMRI data have a greater degree of similarity compared to natural images. This can be attributed to a couple of reasons. First, most existing applications work with 2D data while dwMRI is a 4D dataset. The increased dimensionality allows for greater overlap between patches, hence increasing possibility of similarity between patches. This increased similarity results in a higher chance of generating multicollinear atoms when a coding matrix is constructed. In addition, natural training images often represent a much larger variety of choices for patches compared to the more monotonous dwMRI data. Furthermore, especially in the increasingly more common HARDI imaging schemes, the angular resolution of acquisition is very high. At a given physical location, this results in very similar volumetric patches across acquisition orientations, which further increases the chances of having multicollinear atoms.

To illustrate the aforementioned points, we show an example comparing the condition number of a dictionary coding matrix from both natural images and dwMRI data. This is illustrated in Figure 3.8. We observe

that, for the same dictionary size, dwMRI data yield significantly more ill-conditioned coding matrices compared to natural images. As such, this makes it impossible to use the dictionary learning framework to perform super-resolution since the coding matrix becomes singular, preventing inversion operations. This is explained in section 3.2.3 in more detail.

However, we propose to solve the problem using a preconditioning approach. We perform an alternative construction of the training matrix in order to improve the condition number. In this alternative construction, we seek to populate the training matrix with initial atoms that have reduced overlap and thus reduced collinearity. We propose to achieve this by populating the initial atoms along salient 3D structures in a dwMRI volume. Our motivation for following this approach is that including more salient structures as initial training atoms would be expected to reduce overlap compared to allowing smooth, non-structured, patches into the training matrix. An overview of the proposed preconditioning approach is shown in Figure 3.9.

We propose to detect the salient local structures using a *shearlet decomposition* [50]. This method decomposes a volume into a set of coefficients that represent 3D surface-like discontinuities at various locations, scales, and orientations. Each coefficient is associated with a particular combination of scaling, shearing, and translation of a generating function. By capturing the spatial structure of high dimensional discontinuities (instead of capturing a 1D discontinuity in each direction), this system provides a better representation of high dimensional structures.

The shearlet system has been shown to yield optimally sparse representations compared to similar frameworks, such as curvelets and contourlets [34, 49]. In this work, we use the implementation provided by the authors [50]. As for the filter parameters, we used the parameters that has been shown in a previous study [70] to yield good reconstruction in MRI, which

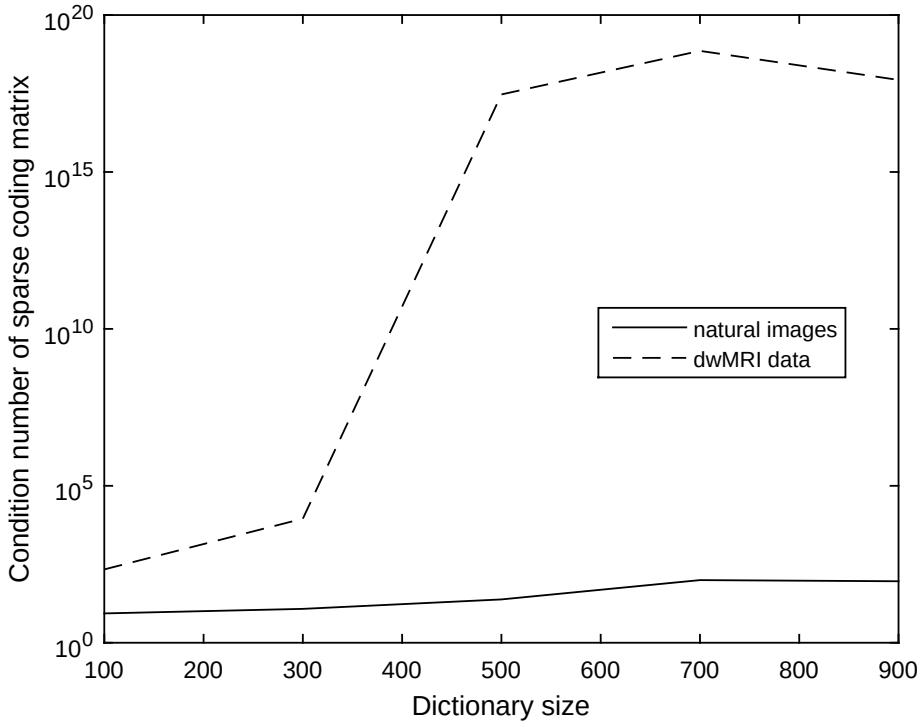

Figure 3.8: An example comparison of the condition number of the coding matrix in both natural and dwMRI data. The coding matrix of dwMRI data exhibit orders of magnitude higher condition number compared to natural images, which presents a challenge in using the matrix in the necessary inversion operations.

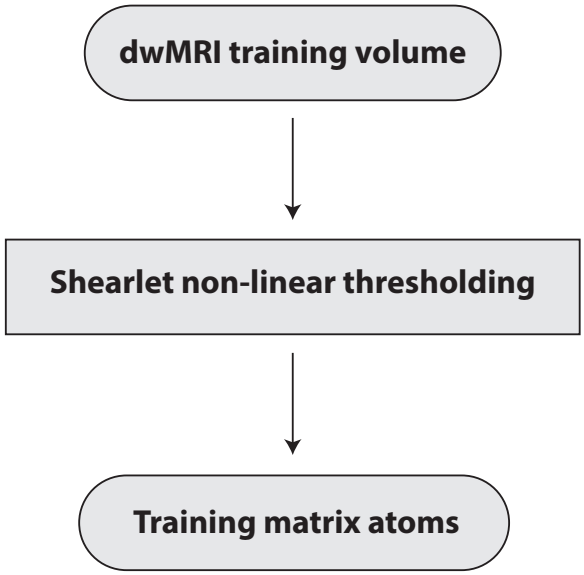

Figure 3.9: The proposed training matrix preconditioning process.

were diamond flat filters, with 4 scalings and 4 shear levels in the first two scales and 8 shear levels in remaining two scales.

We are now at a stage of being able to decompose the dwMRI volume into a multi-scale geometrical representation system. What we need to do is to determine a method for extracting the salient local structures in order to populate the initial training matrix.

We perform this task by a non-linear thresholding of the decomposed shearlet coefficients. Using this approach, by selecting only the highest M coefficients, we extract the most 3D-surface-like structures in the dwMRI volume. We also note that since the thresholding is done non-linearly, the selected M coefficients out of all N coefficients have no specific scale or orientation, and therefore the structures that are extracted are not at a predetermined location, scale, or orientation.

Accordingly, we run a non-linear thresholding experiment to determine the most suitable threshold for detecting structures. The result is shown in Figure 3.10. We begin by observing that, as expected, the curve starts with a large error due to the absence of most coefficients. Then there is a large drop in error, after which the approximation gradually plateaus.

We emphasize that our goal is *not* to faithfully reconstruct the volume, and hence large M/N is not desirable. Instead, we are looking for only the most geometrically salient structures. A good point to look for structures would be at the onset when we barely include any coefficients but the approximation error drops sharply. This indicates that, at this point, the volume has now gained strong salient structures. In Figure 3.10, this occurs around the value 0.03. As such, this will be our choice for non-linearly thresholding the coefficients. Figure 3.11 shows an example of the result obtained from using this thresholding value.

Finally, we binary threshold the reconstructed volume at different values. We then assign all resulting voxels of this operation as center points of

the atoms to be initialized into the training matrix. The dictionary is then learned on the training matrix, and the condition number of the resulting coding matrix is plotted. The result of this experiment is shown in Figure 3.12. We can clearly observe a minimum point around 2%, which is then monotonically increasing afterwards. As such, we choose the thresholding value of 2%.

Using the above parameters, we now examine the end result of the preconditioning process by observing the new condition numbers. This is shown in Figure 3.13. The figure confirms that we have achieved a low condition number for the coding matrix.

3.2.3 Joint-dictionary Learning

In this section, we describe details pertaining to construction of the joint-dictionaries, and the process of using the dictionaries to super-resolve data. We begin by generating a training set using the approach proposed in section 3.2.2. Let the training set be denoted by $\mathbf{\Omega}$, and defined as follows

$$\mathbf{\Omega} = \{ (\rho_L^i, \rho_H^i) \mid \rho_L \in \mathbb{R}^d, \rho_H \in \mathbb{R}^{8d} \} \tag{3.8}$$

where ρ_L indicates a low-resolution patch, ρ_H indicates a high-resolution patch, i is the patch index within the training set, and d is the dimension of the low-resolution patch. Next, the low-resolution dictionary is created using the same minimization as in [46, 97]:

$$\mathbf{D}_L, \{\alpha^i\} = \underset{\mathbf{D}_L, \{\alpha^i\}}{\operatorname{argmin}} \sum_i \|\rho_L^i - \mathbf{D}_L \alpha^i\|_2^2 + \lambda \|\alpha^i\|_1 \tag{3.9}$$

where $\mathbf{D}_L \in \mathbb{R}^{n \times m}$, with m atoms of size n, and $\alpha^i \in \mathbb{R}^m$ is the sparse coding vector of the i-th patch, and λ is an optimization weight controlling the sparsity of the coding vector.

The existing dictionary-based super-resolution methods [46, 97] typically use the K-SVD (K-singular value decomposition) algorithm [3] to im-

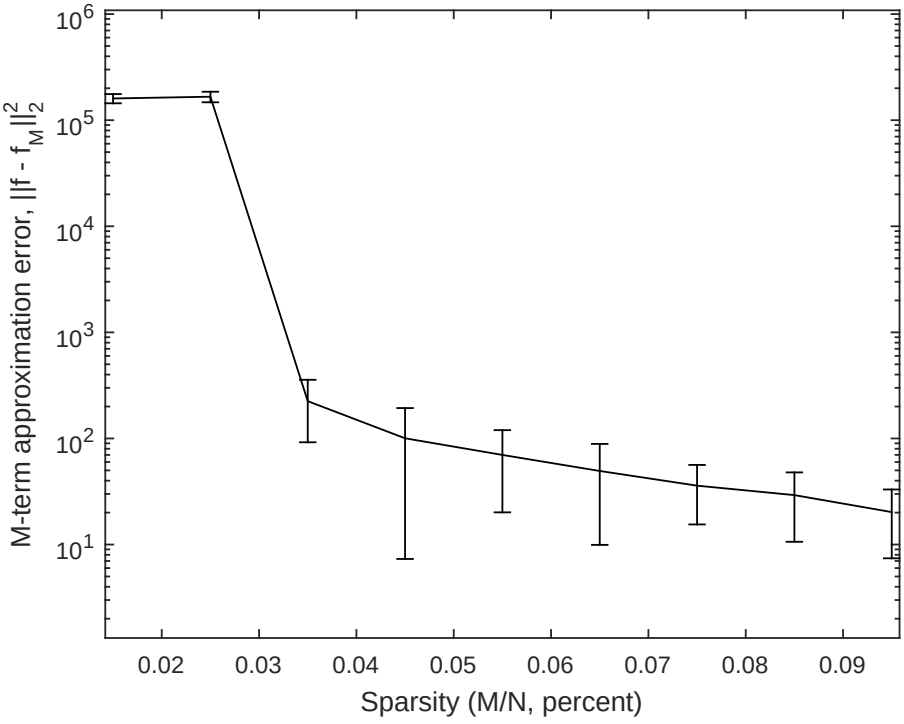

Figure 3.10: The M-term approximation of a shearlet decomposition of a dwMRI data. Note the onset of a drop in approximation error occurring around 0.03, which gradually plateaus afterwards. This indicates that, at this point, the volume has now gained strong salient structures. This will therefore be the M/N ratio used to detect structures. Vertical bars show standard deviation across 10 random HCP volumes.

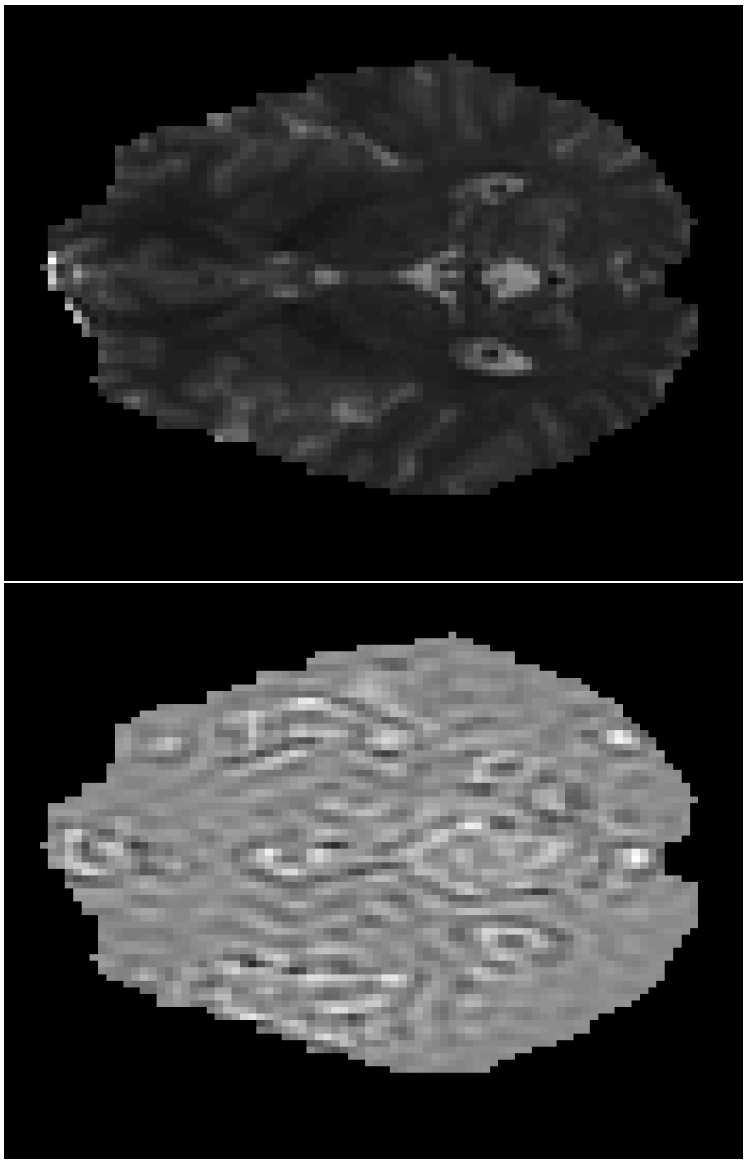

Figure 3.11: Example result of non-linearly thresholding the shearlet coefficients. The threshold is set to the point 0.03 in Figure 3.10. The intensities in the upper figure reflect diffusion signal, while intensities in the bottom figure reflects salient shearlet locations.

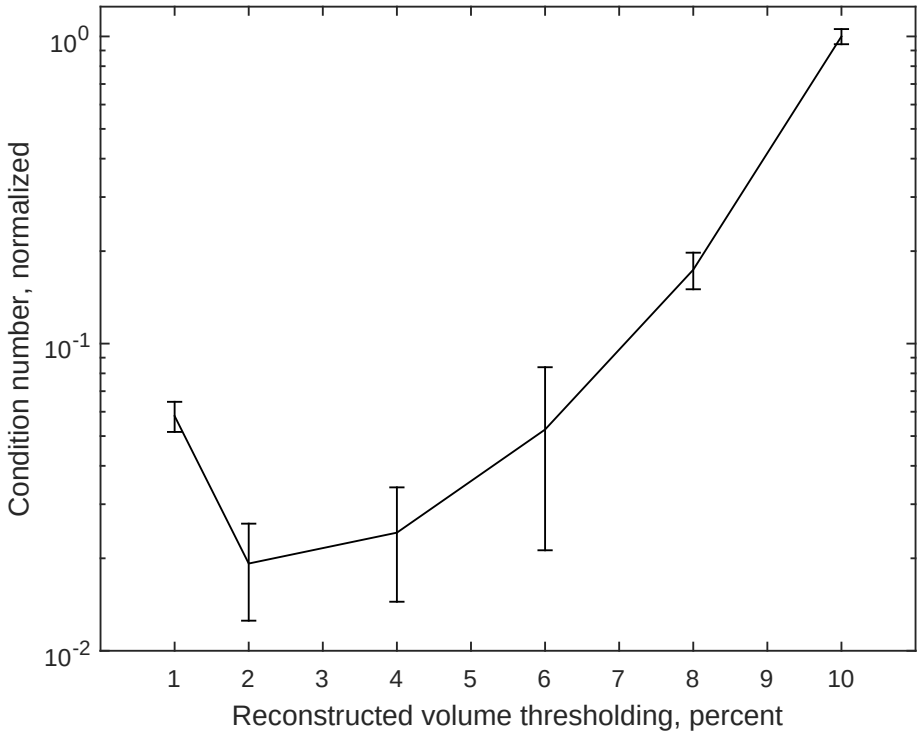

Figure 3.12: The condition number of the resulting coding matrix for various levels of binary thresholding of the shearlet reconstructed volume. Vertical bars show standard deviation across 10 random HCP volumes.

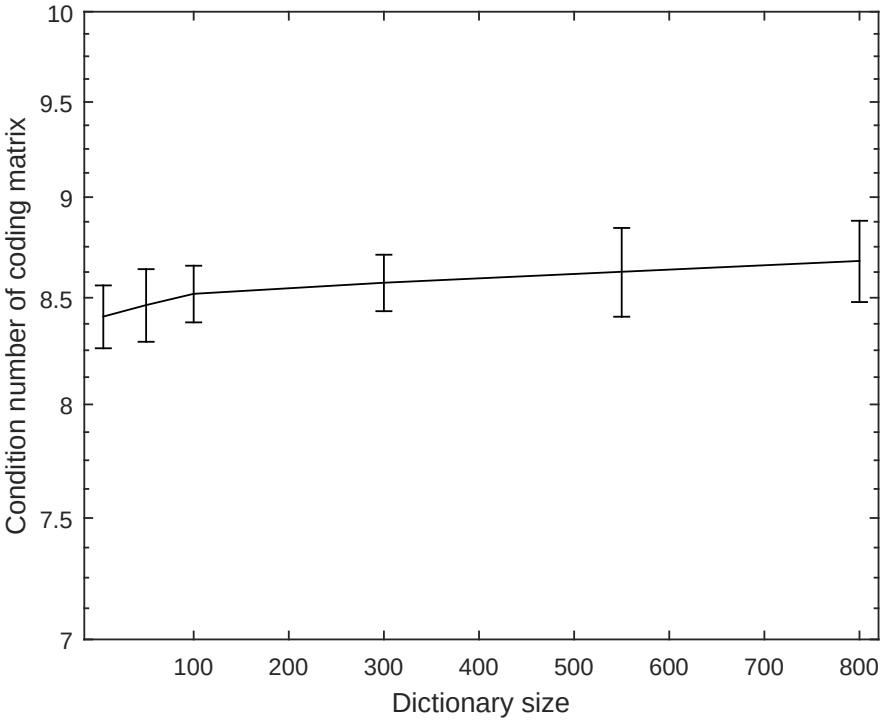

Figure 3.13: The condition numbers of dictionary coding matrix after the training matrix preconditioning step. The condition numbers are low at a various dictionary sizes.

plement the joint-dictionary construction. While K-SVD has demonstrated good performance, it posits a challenge for implementing the super resolution method for dwMRI data.

This is due to the batch-based nature of the algorithm, which requires access to the entire training set. For dwMRI data (especially the increasingly common multi-shell acquisitions), the training set can be in the range of tens of millions of samples, which represents a computational burden if required to be loaded at once. Furthermore, it has been shown that batch-based methods such as K-SVD can not effectively handle large training sets [60].

At this point, we recall that one of the objectives of this thesis is to provide an acquisition-independent framework. With the above limitations of K-SVD regarding large training sets (which hinders the use of newer acquisitions like multi-shell), it is no longer a good choice for implementing the joint-dictionary dwMRI super-resolution.

In contrast, in this work we propose to implement the joint-dictionary super resolution process via an *online learning* approach using the SPAMS (SPArse Modeling Software) algorithm [60]. Using this algorithm, the training set is only minimally loaded (one or a few training samples at a time), which makes the learning online. Our choice of this implementation stems from the fact that dwMRI data is orders of magnitude larger in size compared to natural images that have been the target of existing dictionary-based super-resolution methods [46, 97]. As such, we use an online learning approach in order to *gradually* construct the dictionaries from a sequence of *portions* of the dwMRI training set.

After computing the sparse coding vector α^i in (3.9), the high-resolution dictionary $\mathbf{D}_H$ is constructed as in [46, 97]

$$\mathbf{D}_H = \underset{\mathbf{D}_H}{\operatorname{argmin}} \sum_i \|\rho_H^i - \mathbf{D}_H \alpha^i\|_2^2 \tag{3.10}$$

which can be solved as in [97] using pseudo-inversion:

$$\mathbf{D}_H = \mathbf{P}_H \mathbf{A}^+ = \mathbf{P}_H \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \tag{3.11}$$

where $\mathbf{P}_H$ is a matrix of the patch set ρ_H^i and $\mathbf{A}$ is a matrix of the coding coefficient vectors α^i.

For dwMRI data, this is presents a problem due to the spatial and angular overlap between patches as described in sections 3.2.1 and 3.2.2, leading to multicollinearity and ill-conditioning of $\mathbf{A}$ which hinders the stability of the inversion in (3.11). The approach we propose to solve this issue is described in section 3.2.2, which preforms a preconditioning on the dwMRI training matrix in order to improve the condition number of the resultant coding matrix $\mathbf{A}$. This enables the successful computation of the high-resolution dictionary $\mathbf{D}_H$ for dwMRI data.

After the two dictionaries $\mathbf{D}_L$ and $\mathbf{D}_H$ are constructed, the training phase is concluded. At this stage, a new un-seen dwMRI dataset can now be super-resolved. We now describe the super-resolution process utilizing the two learned dictionaries, which follows the same approach as [46, 97].

First, the input volume is converted into a patch set that matches the size of low resolution dictionary $\mathbf{D}_L$. Next, each patch in the set is sparse coded against the low resolution dictionary. After that, the resultant coding vector is multiplied against the high resolution dictionary, which finally generates the high resolution patch set. The patch set is then finally assembled back into a high-resolution volume, averaging overlapping areas.

We now turn to the problem of determining the set of dictionary learning parameters for super-resolving dwMRI. We start by examining the effect of the dictionary learning optimization weight λ on the reconstruction quality. That is, the reconstruction of a given volume from its downsampled version. The quality of this reconstruction is shown in Figure 3.14.

From the figure, we make a number of observations. First, we note that

decreasing the weight λ increases the reconstruction quality. This plateaus at around $\lambda = 10^{-3}$. We also note that, for a given value of λ, larger dictionaries result in better reconstruction. This holds until around 700 - 1100 atoms, where performance plateaus. This behavior is reasonable since a larger dictionary of atoms allows for learning more example atoms. Accordingly, we choose $\lambda = 10^{-3}$ with a dictionary of 900 atoms as parameters for the rest of this paper. The effect of the choice of images used in training, and whether performance depends on the type of image used, will be explored in great detail in Chapter 5.

Next, we determine the atom size parameter. In this analysis, we examine the reconstruction quality at different isotropic length dimensions of the atom. The performance is shown in Figure 3.15. The figure shows that better performance is achieved with smaller atom sizes, and that the best atom size is 3 voxels (isotropic). This performance is reasonable since a training matrix with smaller atoms would provide more consistency to the dictionary learning method.

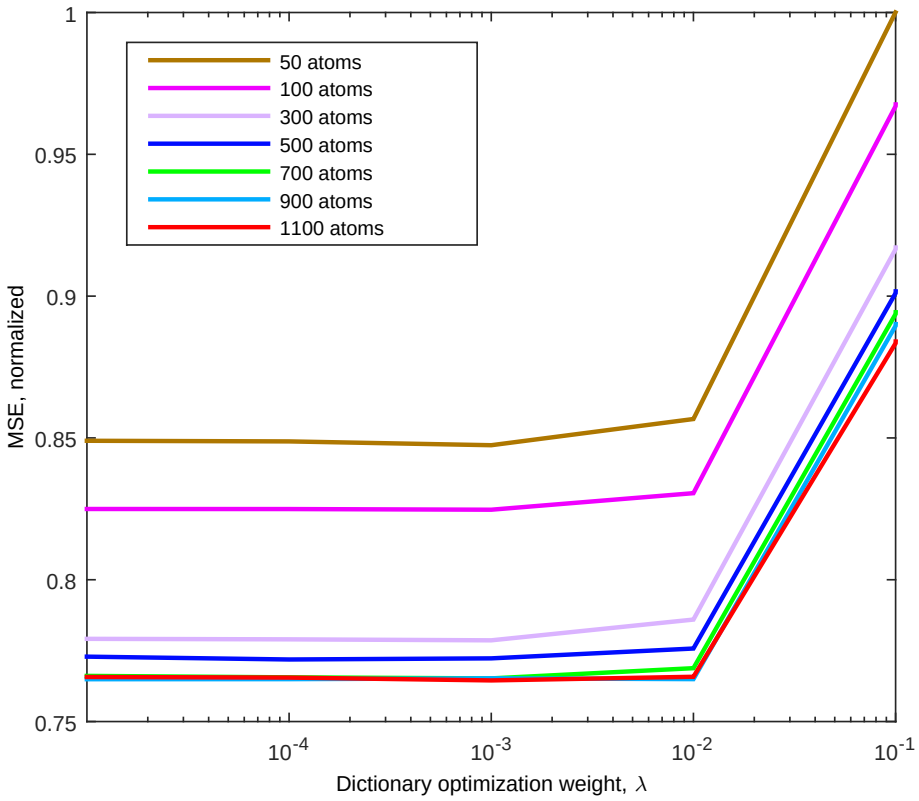

Figure 3.14: The effect of dictionary optimization weight λ, at different dictionary sizes, on reconstruction performance. Note the plateau occurring around 700 - 1100 atoms.

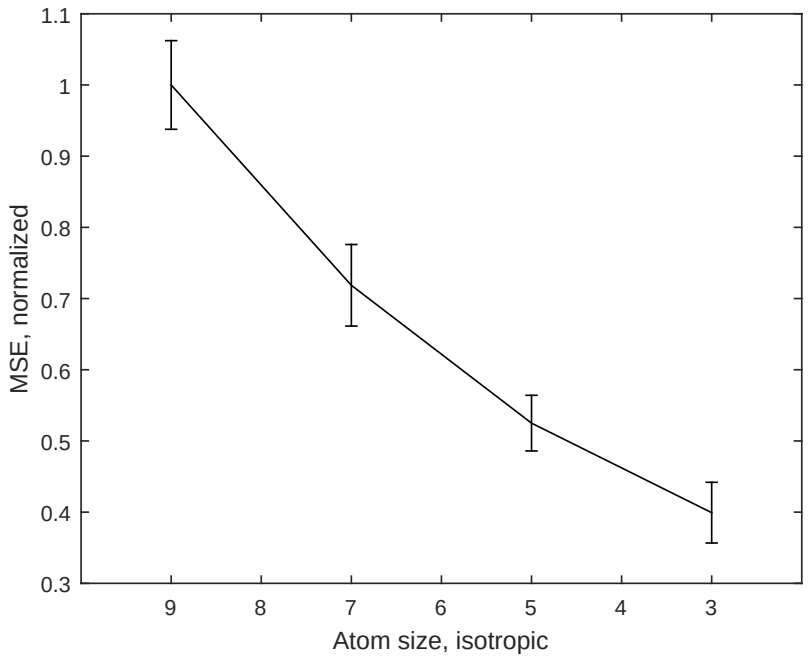

Figure 3.15: The effect of dictionary atom size (isotropic) on reconstruction performance. Vertical bars show standard deviation across 10 HCP volumes.

Chapter 4

Validation Methodology

4.1 Overview

Increasing the resolution of a dwMRI volume beyond what is available in the original data presents a challenge in terms of validation. This is because of the lack of ground truth information at the higher resolution. dwMRI remains the only technique for obtaining non-invasive in vivo information about fiber tracts in the brain. As such, we have no other high resolution source of information to directly validate the generated high resolution dwMRI volumes against.

In order to address this issue, we propose to use a validation approach that builds on the existence of a strong dependence between structure and function in the brain [39, 52]. Accordingly, brain regions that have strong structural connections also have a strong functional connection [92, 99].

Motivated by this relationship between how functional connectivity (FC) reflects the underlying structural connectivity (SC), we quantitatively validate our results by investigating the consistency between SC and FC before and after super-resolving the data. In other words, a greater SC-FC consistency indicates a better reconstruction quality, compared to a low SC-FC.

This validation approach is beneficial for a number of reasons. First, fMRI data is typically readily included with many dwMRI acquisitions, reflecting the same subject at the same point of time. This has the benefit of having an independent modality (fMRI, in this case) used to validate an-

other modality (dwMRI) of the same subject. Furthermore, fMRI is typically not resolution limited. In fact, fMRI voxels are rarely meaningful individually without being parcellated into a larger group. For these reasons, improvements in dwMRI reconstructions would translate into greater SC-FC consistency.

In the next sections, we start by describing the process of obtaining FC from fMRI data. We then proceed to estimate SC from dwMRI. After the two estimates are obtained, we finally calculate the SC-FC correlation as a quantitative metric for dwMRI resolution enhancement.

4.2 Functional Connectivity Estimation

Functional information on the brain is estimated from functional-MRI data (fMRI). For each voxel, the intensity of the data represents the level of neural activity, as measured through oxygen-level changes (referred to as the blood-oxygen-level dependent contrast, *BOLD*). In addition, due to the dynamic nature of brain function, the data is continuously recorded for a period of time. The result is a functional time course per voxel.

We begin the estimation of functional connectivity by parcellating the brain voxels into regions. Parcellation is commonly performed on fMRI data because functional information is rarely specific to a certain voxel, but rather occur synchronously with a larger group of voxels.

We parcellate the brain into 200 regions using Ward clustering [63], which has been shown to perform better than other fMRI parcellation approaches [82]. The number of parcellation regions was set based on the results of a previous study [83] which recommended the use of 200 or more parcels.

Accordingly, let $\mathbf{Z}$ be a $t \times d$ matrix of fMRI time courses, where t is the number of time points and d is the number of brain regions. We then

estimate the functional connectivity (FC) using Pearson's correlation:

$$\mathbf{FC} = \mathbf{Z}^T\mathbf{Z}/(t-1) \tag{4.1}$$

which, in a comparative study with other common FC metrics, has been shown to yield better performance [26].

4.3 Structural Connectivity Estimation

While the neuroimaging research community has largely settled on Pearson's correlation for FC connectivity measurement [26], investigations of structural connectivity (SC) metrics have received less attention. An original contribution of this thesis includes presenting a comparison of structural connectivity metrics assessed from the perspective of the largely accepted inherent relationship between brain structure and function [7].

Quantifying structural connectivity in the brain is most commonly based on quantifying one or more aspects of the streamlines reconstructed using deterministic tractography, though computationally expensive approaches based on probabilistic tractography techniques were also explored [11]. The choice of which streamline property to measure and of how to map it into a structural connectivity metric are key aspects affecting the structural connectivity estimates.

Arguably the most common SC metric is the number of reconstructed streamlines between pairs of brain regions, commonly referred to as fiber count. A variant of this approach involves a normalization of the fiber count by the total volume of the region pairs they connect to account for the variable size of the brain regions [36]. Besides metrics based on fiber count, use of the total length of reconstructed streamlines has also been suggested as a measure of structural connectivity [61] aiming to correct for the fact that longer tracts have larger accumulated error, leading to lower fiber counts.

Another metric, the average fractional anisotropy (FA) along streamlines connecting regions, has also been proposed as a proxy for structural connectivity strength [14].

It is important to acknowledge that all of the aforementioned SC metrics are confounded by several factors, limiting their interpretability. First, tractography can only delineate bundles of fibers in the brain, and not individual fibers. The term fiber count can thus be misleading. Indeed, using the term streamline count has been recently proposed as an alternative [44]. Nonetheless, we use the term fiber count for easier interpretation and to conform to the jargon used in existing literature, with the understanding that it is the streamlines that are actually being counted. Moreover, we note that the number of fibers is dependent on the number of seeds used for tracking the fibers, the tractography method used, and several features of the pathway such as curvature, length and width [44]. Additionally, we highlight that FA not only depends on the reliability of local diffusion information, but also on a large number of modulating factors such as axonal ordering, axonal density, amount of myelination, and increase in extracellular or intracellular water [44]. Such confounding factors did not impede the adoption of a variety of SC metrics, driven by a practical need for quantifying the degree of connection between brain regions.

We propose that reconciling the presence of confounding factors with the practical need for connectivity estimation calls for a detailed analysis, in a quantitative comparison, to determine which SC metric has the highest potential of being of practical use in multimodal brain image analysis efforts. To this end, we compare four commonly used SC metrics in terms of their impact on the relationship between estimates of SC and FC. On 38 subjects from the Human Connectome Project (HCP) database [91] (which is described in detail in the datasets section – section 5.1), we show that region volume-normalized fiber count best correlates with FC. We also show

that total fiber length has the least bias towards distance between brain regions. We further demonstrate that these results hold across seven different tasks and resting-state data.

We now describe the calculation of the four SC metrics. Let $r_{i,j}^k$ be the k^{th} reconstructed fiber between a pair of structurally connected regions P_i and P_j. We consider four widely used structural connectivity measures in this work: fiber count ($f_{i,j}$), fiber count normalized by the total volume of the connected regions ($N_{i,j}$), total length of fibers connecting region pairs ($L_{i,j}$), and average FA along the fibers. For each subject, we compute the Pearson's correlation between the FC and SC estimates to quantify the SC-FC relationship for each SC metric. More formally, the metrics can be expressed as follows:

$$N_{i,j} = \frac{f_{i,j}}{V(P_i) + V(P_j)} \tag{4.2}$$

$$L_{i,j} = \sum_k l(r_{i,j}^k) \tag{4.3}$$

where $V(\cdot)$ is the volume of the corresponding region, and $l(r_{i,j}^k)$ is the length of $r_{i,j}^k$.

Prior to the computation of SC metrics, we reconstruct the fibers via global tractography on constant solid angle orientation distribution function (ODF) using MITK [68]. Global tractography was chosen over the more common streamline tractography since it was recently shown to facilitate higher SC-FC consistency [98]. In global tractography, short fiber segments are connected together to generate the set of fiber tracts that best explains the measured dMRI data. As such, at regions with unreliable local diffusion information, the geometry of the surrounding fibers drives the tracking process to prevent premature termination of fibers as is commonly observed in streamline tractography.

The results of the SC-FC correlation for resting-state and task fMRI are

shown in Figures 4.1 and 4.2, respectively. As observed from these figures, average FA has lower correlation with FC compared to the rest of the examined SC metrics (fiber count, volume-normalized fiber count, and total fiber length). This is true for both resting-state and task fMRI. We speculate that the reason for the observed low average FA correlation can be attributed to the large number of factors affecting local diffusion anisotropy [44].

Figure 4.1 also shows that the volume-normalized fiber count has the highest correlation with FC compared to the rest of the examined SC metrics for both resting-state and task fMRI. The pairwise differences between SC-FC correlation assessed using normalized fiber count and other SC metrics were found to be statistically significant at $p < 0.001$ based on the Wilcoxon signed rank test. Our results thus imply that the compensation (due to normalized fiber count) for the differences in number of fibers due to the variable size of brain regions yields better depiction of structural networks.

We also note that the relatively consistent SC-FC correlation levels across a variety of tasks and resting-state data support the notion that SC forms the backbone of the brain connectivity around which functional reorganization occurs to respond to different tasks. A diverse repertoire of functional brain connectivity patterns can thus arise constrained by the same structural substrate.

Figure 4.3 shows a qualitative comparison between functional and structural connectivity patterns. We averaged subject-specific connectivity matrices to compute group results. Specifically, the top 10 parcels having strongest connectivity to posterior cingulate cortex (PCC) are overlaid onto the brain using these group-level connectivity matrices. PCC was selected as the seed as it is known to be a structural and functional hub facilitating efficient communication in the brain [37]. This figure shows that SC patterns estimated using normalized fiber count resemble FC patterns more than those estimated using average FA.

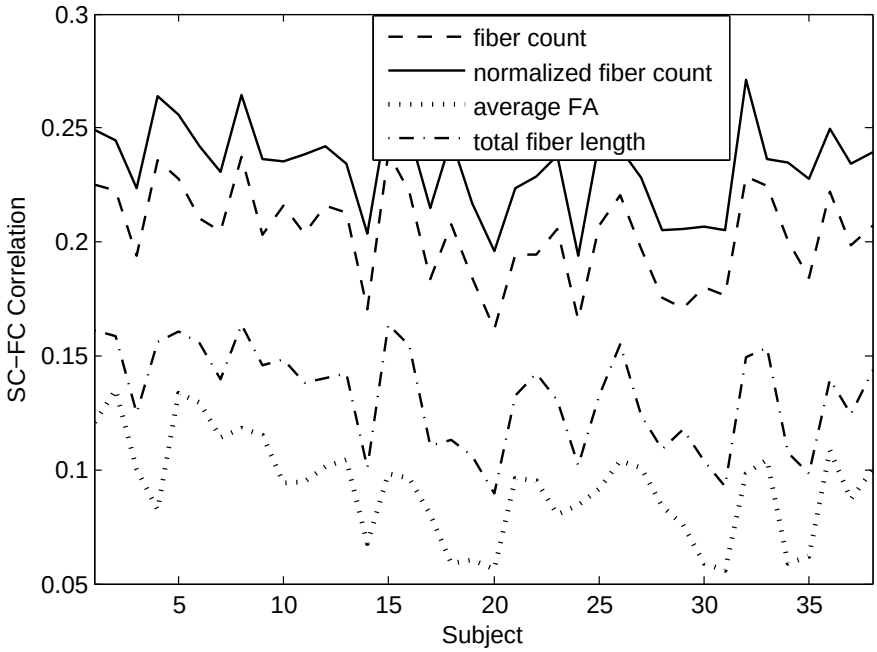

Figure 4.1: Comparison of four common SC metrics in terms of SC-FC (resting state fMRI) correlation for 38 subjects from the Human Connectome Project.

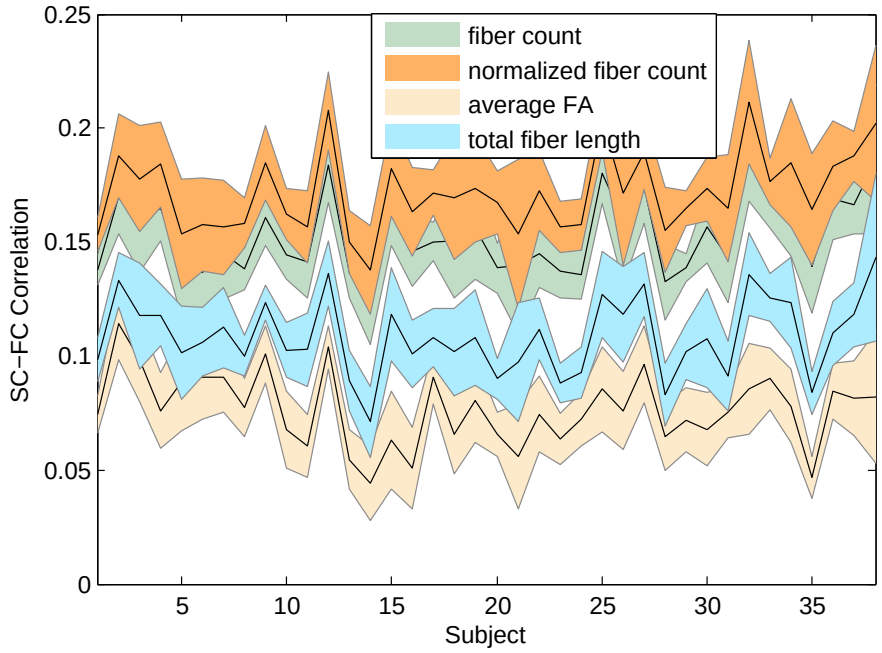

Figure 4.2: Comparison of four common SC metrics in terms of SC-FC (task-fMRI) correlation for 38 subjects from the Human Connectome Project. The shaded bands represent the standard deviation of SC-FC across 7 different tasks.

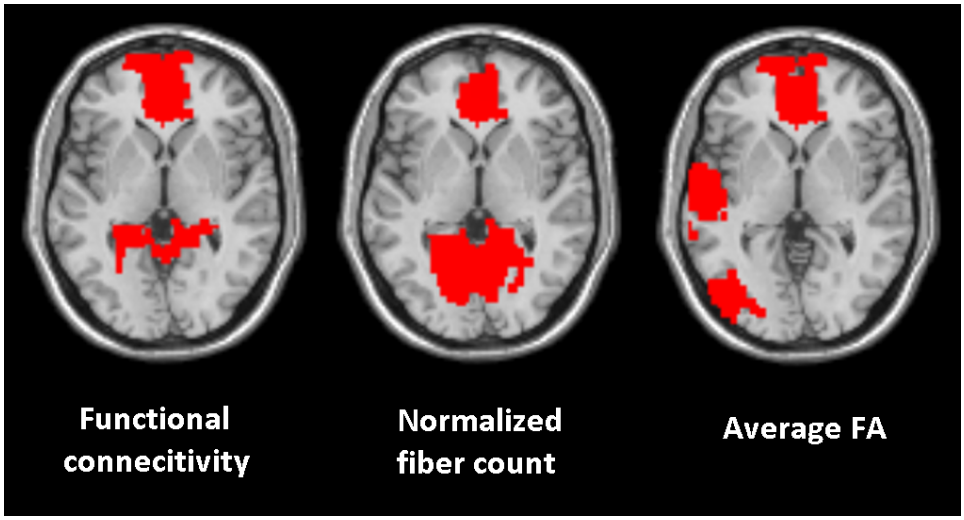

Figure 4.3: Parcels with highest connectivity to posterior cingulate cortex as obtained from: (a) Functional connectivity (FC), (b) normalized fiber count, and (c) average FA. Note how the arrangement of connected parcels obtained by normalized fiber count has a better resemblance to the parcels obtained by FC, compared to average FA.

Chapter 5

Results

5.1 Datasets

Two different publicly available datasets were used for the experiments in this chapter. The first dataset consists of dwMRI data from 38 subjects (17 males and 21 females, ages ranging from 22 to 35 years) from the Human Connectome Project (HCP) Q2-13 dataset [91]. This release of the dataset has 40 subjects for which dwMRI data was available. We excluded two subjects from the dataset (subjects #209733 and #528446) as per HCP's recommendation, due to reported structural brain abnormalities. The dwMRI data had a voxel size of 1.25 mm (isotropic), 3 diffusion shells (at $b = 1000$, 2000 and 3000 s/mm^2) and a total of 288 gradient indecies. The HCP dwMRI data used in this chapter includes the suggested minimal preprocessing pipeline already applied by the HCP team, including corrections for EPI distortion, eddy current, gradient nonlinearity and motion artifacts [31]. Further details on the dwMRI acquisition can be found in [91].

The second set of data is the Kirby 21 dataset [51]. This dataset comprises scans of 21 subjects (11 males and 10 females, 32 ± 9.4 years old). We used two modalities from this dataset: rs-fMRI data and dwMRI data. The rs-fMRI data consisted of a 7-minute acquisition with a TR of 2 s and a voxel size of 3 mm (isotropic). We then preprocessed the data for motion correction and bandpass filtered from 0.01 and 0.1 Hz using an in-house MATLAB code. We then divide the brain into 150 parcels using Ward clustering [64]

applied on the voxel time courses. The dwMRI data consisted of 32 diffusion gradients with a b-value of 700 s/mm^2, in addition to a single $b0$ image. The voxel size was $0.83 \times 0.83 \times 2.2$ mm^3. However, since anisotropic voxels were previously shown to be suboptimal for further processing of dwMRI data in the sense that fiber branching is less detectable [67], we resampled the data to 2 mm isotropic voxels prior to any subsequent analysis and processing. Finally, we then warped the functionally derived group parcellation map to the $b0$ volume of each subject using FSL [42] in order to facilitate comparisons of structural and functional metrics.

5.2 Quantitative Performance

We perform quantitative assessments in three categories of tests: tests on the HCP data, tests on the Kirby data, and tests that combine both datasets (i.e. training on one and testing on the other). In order to quantify the quality of the constructed volumes at higher resolutions, we performed two sets of experiments. The first experiment measures the similarity between a ground truth volume and a reconstructed volume from its downsampled version. The quality is quantified using a reconstruction error metric $\eta = MSR_{SR}/MSE_{Interp.}$, which are the MSE from the super-resolved reconstruction and spline interpolation, respectively.

The second experiment aims to quantify the quality of the constructed volumes at a higher resolution than that of the ground truth. The quantification metric of this experiment follows the SC-FC correlation measurements described earlier in Chapter 4.

Multi-shell results

Figure 5.1 shows the reconstruction error η for a number of HCP test subjects. These results were generated from training the dictionaries on 30 ran-

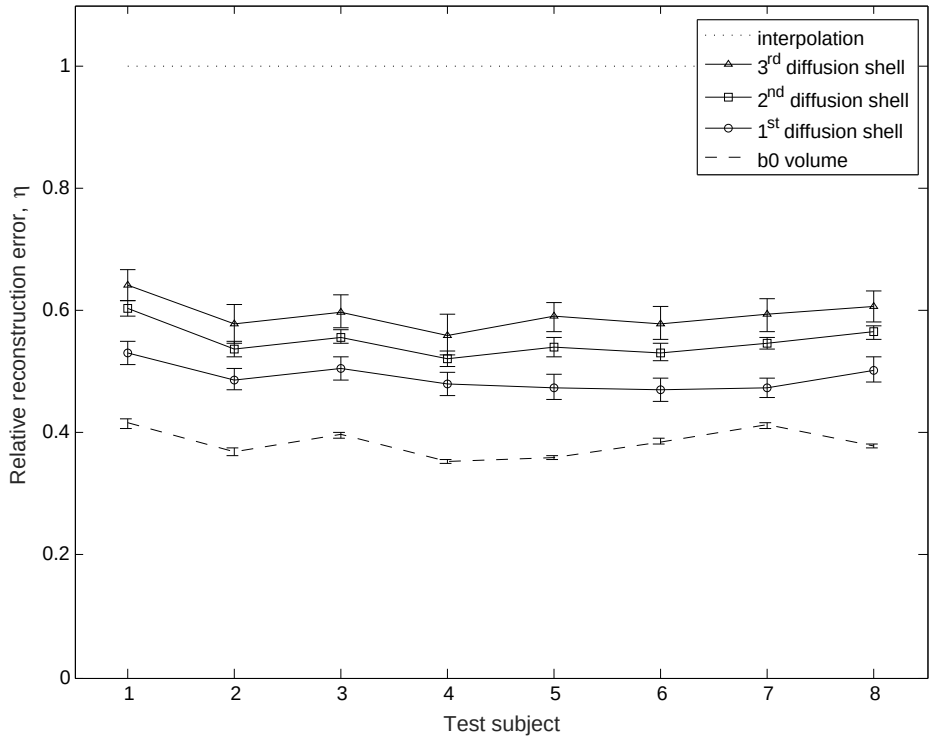

Figure 5.1: Relative reconstruction error η for HCP dataset [91] test subjects. The vertical bars show the standard deviation of error across diffusion directions at each shell. The vertical bars for the $b0$ volume show the standard deviation of error across repeated $b0$ acquisitions.

dom subjects from the HCP data, and then testing the reconstruction quality on the remaining 8 subjects from the same dataset. We observe that the figure shows η values in the range of ~ 0.4 for $b0$ volumes, indicating an improvement of around 60%. We also observe a highly consistent reconstruction quality of $b0$ volumes as indicated by the small standard deviation of error.

The figure also shows the reconstruction quality of the diffusion shells. We observe that the quality of diffusion shell reconstruction is lower than that of the $b0$ volumes. This performance is not unexpected and we attribute it to the lower SNR of the diffusion shell acquisitions compared to the $b0$ volumes and also the lower number of repetitions compared to $b0$ volumes. This is further supported by the observation that, within the diffusion shell results, higher shells show lower reconstruction quality compared to lower shells, recalling that higher shells have lower SNR compared to lower shells. For the same reasons, we also observe that the consistency of reconstruction quality for diffusion shells is generally lower than that of the $b0$ volumes.

Next, the results of Figure 5.2 use the same dictionary trained for Figure 5.1, but examines the generalization ability of the dictionary by testing it on unseen data from the Kirby dataset. That is, the results of this figure are based on training the dictionary on the HCP dataset and testing on the Kirby dataset, thereby assessing generalization performance of the learned dictionary to a different acquisition type.

Accordingly, we make a number of observations on Figure 5.2. First, we observe that the range of reconstruction quality values are relatively within the same range as that of Figure 5.1. This shows a good generalization performance of the learned dictionaries. However, we still note that there is a slight reduction in the reconstruction quality of $b0$ volumes between Figures 5.1 and 5.2. We emphasize that this reduction is only partially related to the generalizability of the dictionary. The reason for this is that the HCP

$b0$ volumes were acquired 18 times and averaged in order to improve SNR, while Kirby $b0$ volumes were only acquired once. Hence, the reconstruction quality of HCP $b0$ volumes is expectedly better than Kirby $b0$ volumes, regardless of generalizability of the dictionary. Nonetheless, the quality of Kirby $b0$ volumes constructed from the dictionary is still relatively similar to that of HCP $b0$ volumes.

In addition, we also note that the gap between the $b0$ and diffusion volumes is smaller in Figure 5.2 compared to Figure 5.1. We suggest that this may be attributed to the fact the difference between $b0$ and diffusion volumes is smaller for Kirby data compared to HCP. For Kirby data, this is 700 s/mm^2 while for HCP data the difference is 1000 s/mm^2 and higher.

Single-shell results

Next, we examine the effect of changing training data on the quality of reconstructions. This is shown in Figure 5.3. As in the previous figure, this experiment also uses the Kirby dataset for test subjects. However, the training data is changed from Kirby to HCP. The previous experiment used HCP data for training and Kirby data for testing, while this experiment in Figure 5.3 uses Kirby data for both training and testing. The difference between the two experiments is the choice of training dataset.

The training data is constructed from a randomly selected list of 11 subjects from the Kirby dataset. The testing dataset is then chosen to be the remaining 10 subjects in the dataset. The results are shown in Figure 5.3.

We make a number of observations regarding this experiment. First, we note that the overall range of reconstruction error is relatively similar to that of Figure 5.2. This shows that, for the same test dataset, the performance is not highly sensitive to training data. Nonetheless, given that the training and test subjects now belong to the same dataset, we do observe an improvement in the consistency of reconstruction. More specifically, while

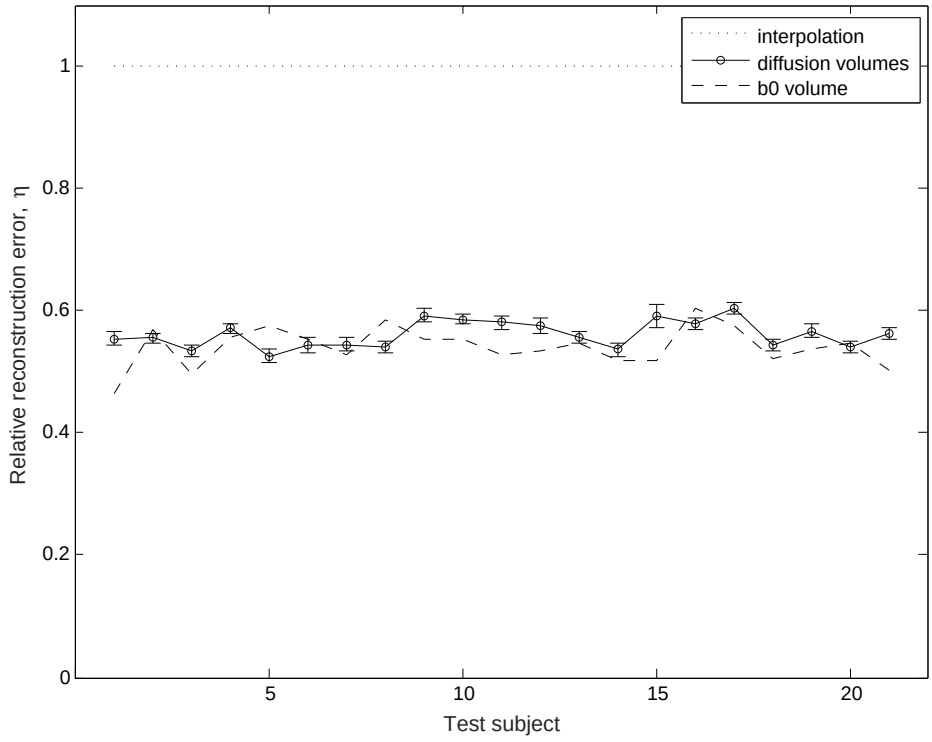

Figure 5.2: Relative reconstruction error η for Kirby dataset [51] test subjects using dictionaries trained on HCP dataset [91]. The vertical bars show the standard deviation of error across diffusion directions.

the reconstruction error ranged between 0.5 to 0.6 in Figure 5.2, the error range in Figure 5.3 is more consistent at around 0.5. We also note that the gap between $b0$ and diffusion volumes is now considerably smaller. This may be attributed to the fact that the dictionary is now learning from a single diffusion shell compared to learning from a cluster of shells in the HCP data. Therefore, there is now expectedly less ambiguity in reconstructing diffusion volumes, which helps improve reconstruction quality. This also is expected to result in more consistent reconstruction quality across diffusion volumes for the same shell, which is indeed the case in Figure 5.3.

Next, we perform another experiment on the sensitivity to the choice of training data. In this experiment, we reverse the roles of the training and testing data used in the previous experiment. The test data of Figure 5.3 are now used as training data, and the training data of that figure are now used as test data. The result of this experiment is shown in Figure 5.4. We observe that the results of both experiments are clearly very similar. The aforementioned observations about the range of error and reconstruction consistency in the previous experiment also clearly hold in this experiment. This shows that, for a given dataset, the results do not exhibit a sensitivity to the choice of training subjects.

SC-FC correlation We then proceed to the next set of experiments where we compare SC-FC correlations in order to assess the improvement in extending the resolution beyond ground truth.

To the best of our knowledge, the only previous work that tackled the problem of super-resolving dMRI data from a single acquisition independent of the diffusion model was by Coupe et al [20] (referred to as CLASR – collaborative and locally adaptive super-resolution). Specifically, the authors showed that super-resolving $b0$ image using a locally adaptive patch-based strategy, and using this high-resolution $b0$ image to drive the recon-

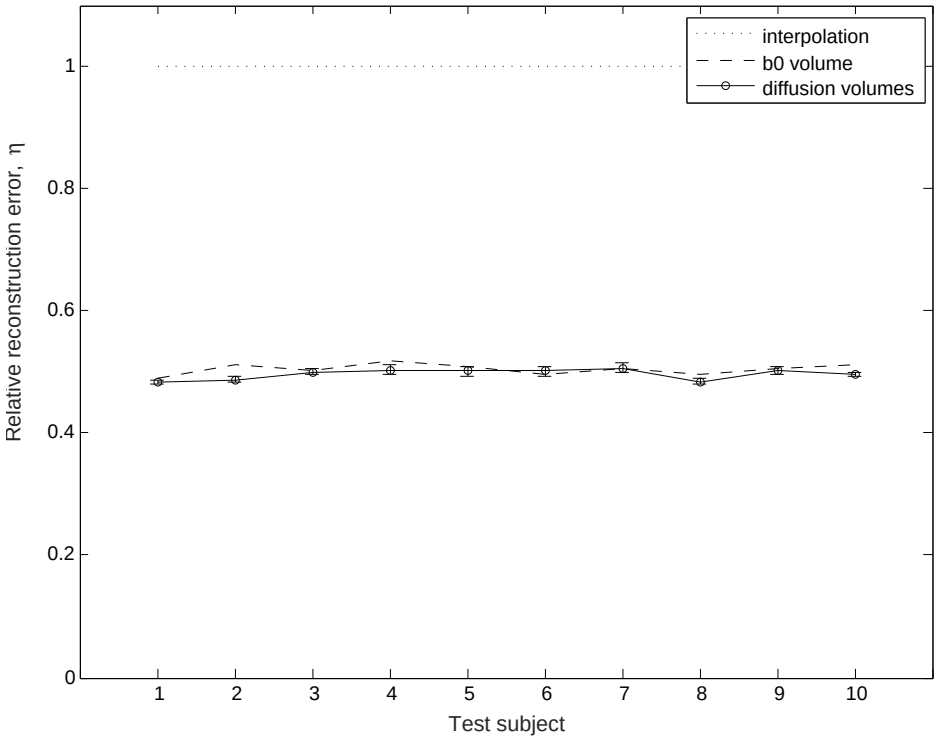

Figure 5.3: Relative reconstruction error η for Kirby dataset [51] test subjects using dictionaries trained on the remaining subjects in the same dataset. The vertical bars show the standard deviation of error across diffusion directions.

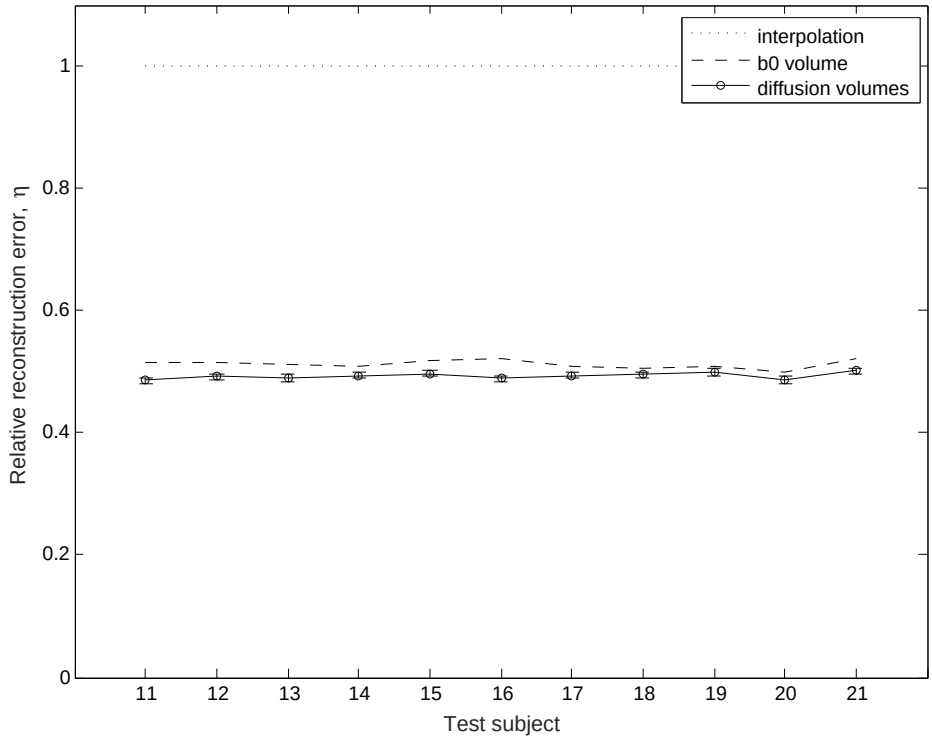

Figure 5.4: Relative reconstruction error η for Kirby dataset [51] test subjects, trained on the remaining subjects in the dataset. The selection of test versus training subjects in this figure is the opposite of Figure 5.3.

struction of diffusion images, outperforms interpolation methods. To the best of our knowledge, CLASR is the only existing super-resolution method developed for dwMRI that is independent of acquisition and the diffusion model employed, which are objectives of this thesis.

To quantify the improvement, we analyzed the consistency between measures of intra-subject SC and FC. We estimated SC using the fiber counts between brain region pairs, and FC using Pearson's correlation between parcel time courses. We chose to employ deterministic streamline tractography with the diffusion tensor model, which is by far the most popular tractography approach to date. However, we highlight that our super-resolution approach can be used with any diffusion model and any tractography method. Tractography was carried out using Dipy [29], with 750,000 seed points for all examined volumes.

For each subject, SC and FC are vectors of size $d(d-1)/2$ comprising the corresponding connectivity estimates between each region pair, where d is the number of brain regions. We then calculated Pearson's correlation between intra-subject SC and FC to quantify the consistency between the two connectivity estimates. Using this correlation measure, we compared the proposed super-resolution approach with trilinear and spline interpolation in addition to a CLASR.

Figure 5.5 shows the SC-FC correlation for each subject tested. Taking the average SC-FC correlation across the group when using the original data as a baseline, the improvement was 5.7% with spline interpolation, 13.6% with CLASR, and 27.1% with our proposed method. On the other hand, there was a 6.3% decrease in the correlation when trilinear interpolation was used. The difference in the performance of our method and every other method tested was found to be statistically significant at $p < 0.01$ based on the Wilcoxon signed-rank test, showing its potential for enhanced structural connectivity assessment. Our results thus suggest that low spatial resolu-

tion of dMRI data can partially account for the low SC-FC correlation, and statistically significant improvements can be achieved using super-resolved dwMRI data.

To investigate why trilinear interpolation resulted in a lower SC-FC correlation compared to the original data, we calculated the number of tracts reconstructed with each method. The local intra-parcel connections were excluded since they have no effect on SC-FC correlation. Figure 5.6 shows the number of inter-parcel tracts averaged across the group along with the corresponding standard deviations. As observed from this figure, performing tractography on volumes upsampled with trilinear interpolation resulted in a lower number of tracts compared to the original volumes, even though the same number of seed points were used to initiate tracking for all of the methods we compared. We speculate that the reason of this phenomenon is the additional partial volume effects introduced by the blurring of the data during trilinear interpolation, which hamper the tractography quality. Spline interpolation, however, is known to cause less blurring compared to trilinear interpolation, and our results suggest that upsampling dMRI data using spline interpolation can be beneficial for tractography. The overall trend of inter-parcel tract counts closely resembles to that of the SC-FC correlation, with our proposed method outperforming all other methods tested. This shows that dictionary based super-resolution is a viable post-processing solution for dwMRI that can help in mapping the white matter brain architecture more accurately.

5.3 Qualitative Results

We now present a qualitative comparison between the fiber tracts reconstructed from the original (2 mm) and super resolved (1 mm) dwMRI data. We employ the same tractography approach as in the SC-FC comparisons

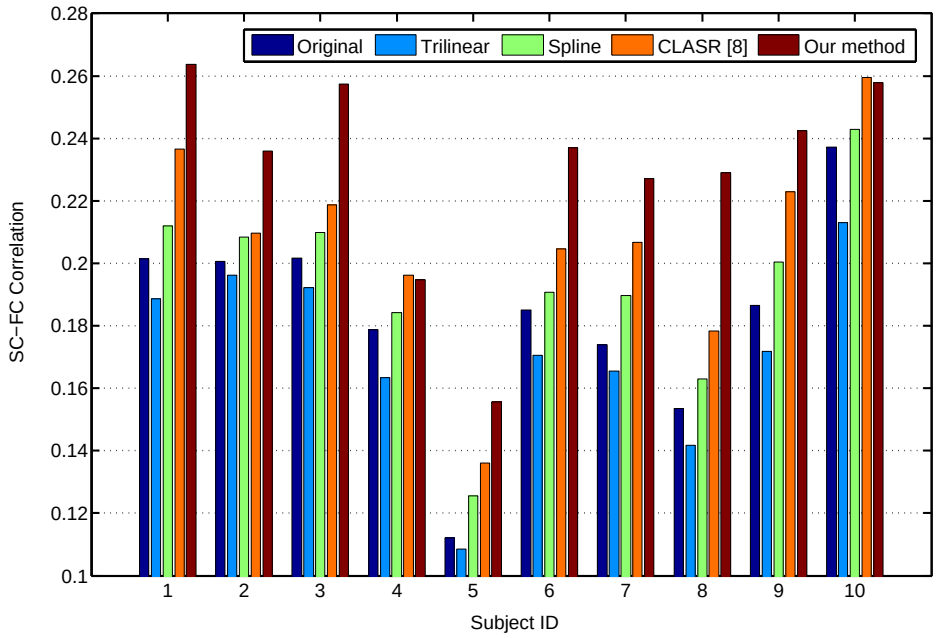

Figure 5.5: SC-FC correlation for 10 subjects with SC estimated from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR, and the proposed method. Our method outperforms all other methods tested for eight of the subjects, and performs comparable to CLASR for two subjects (subjects 4 and 10).

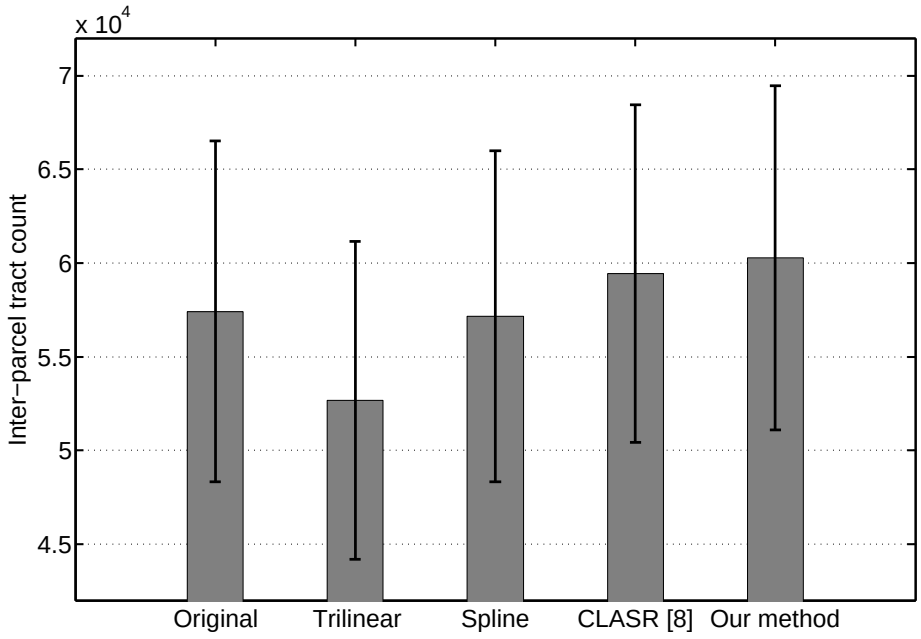

Figure 5.6: Number of inter-parcel tracts reconstructed from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method. Intra-parcel tracts are not included here since they do not contribute to SC-FC correlation. We emphasize that tractography is initiated with the same number of seeds for each method.

(deterministic streamline tractography using the diffusion tensor model with 750,000 seed points). We generated the tract-density maps by calculating the total number of fiber tracts present in each voxel. Figure 5.7 (a),(c) and (b),(d) show sample tract-density maps with the original and super-resolved dMRI data, respectively. As observed from these figures, the tract-density maps generated from the super-resolution data clearly show more spatial information. Figure 5.7 (e),(f) and (g),(h) show the corticospinal tracts extracted using a region of interest (ROI) placed on the brain stem for two representative subjects. It can be observed that fiber tracts reconstructed from the super-resolution data can capture the fan-shape configuration of the corticospinal tract more fully.

Next, we present a qualitative comparison between the raw dwMRI images obtained from original (2mm) and super resolved (1 mm) data. Figure 5.8 shows an example comparison. As observed from the figure, the proposed method has better resemblance to ground truth. The structural features in the image are also more pronounced, while appearing blurred for spline interpolation. This provides an example of the type of details obtained by super-resolving the dwMRI data.

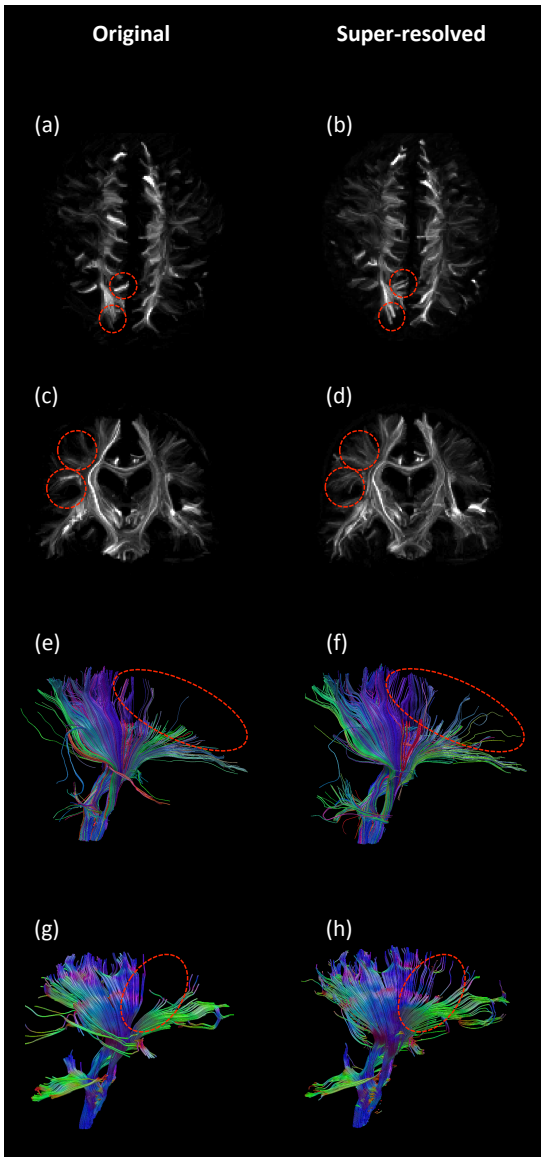

Figure 5.7: Qualitative comparison between the tract-density maps and fiber tracts reconstructed from the original (left) and super-resolved (right) dwMRI data. Original data from the Kirby set has 2 mm isotropic resolution which is super-resolved to 1 mm isotropic resolution. Each row corresponds to a different test subject. Tract-density maps of super-resolved data ((b) and (d)) show markedly improved spatial detail compared to those of original data ((a) and (c)). Corticospinal tracts reconstructed from super-resolved data ((f) and (h)) can capture the fan-shape configuration more accurately than those generated from original data ((e) and (g))

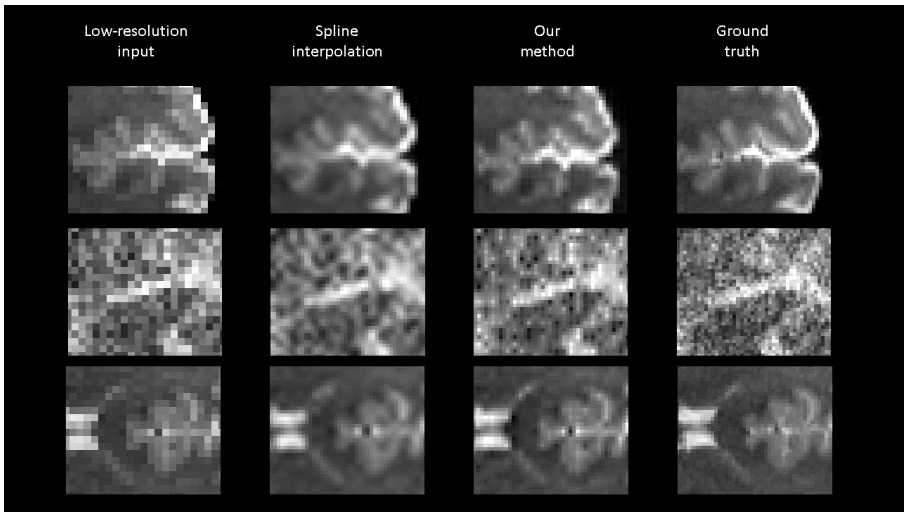

Figure 5.8: Qualitative of comparison of raw diffusion images from the HCP dataset [91]. Note the closer resemblance between the proposed method and ground truth.

Chapter 6

Conclusions

6.1 Discussion

Low spatial resolution is a known limitation of dwMRI, which often hinders the performance of subsequent analysis and determination of structural information. We proposed the use of a simple yet effective super-resolution processing pipeline on dwMRI to capture a more accurate portrayal of the white matter architecture. This approach does not require multiple dwMRI acquisitions and is applicable to legacy data. Quantitatively, we demonstrated that SC-FC consistency can be markedly increased with the use of our approach in estimating SC. We also qualitatively illustrated that the gain in spatial resolution remarkably improves the fiber tracts and tract-density maps generated. Taken collectively, our results suggest a super-resolution based framework holds great promise in enhancing the spatial resolution in dwMRI, without requiring additional scans or any modifications of the acquisition protocol.

We also presented a closer investigation of the validation strategy and presented a comparison of four SC metrics computed from tractography results with respect to their relationship to FC. Among the metrics considered, we showed that volume normalized fiber count has the highest correlation with FC for both resting-state and task data. On the other hand, our results showed that average FA has the lowest correlation with FC. We speculate that the reason of this low correlation is the non-specificity of FA,

with several inadvertent factors (such as axonal density, axonal ordering, and amount of myelination) modulating it along with the reliability of local diffusion information. In addition, we also demonstrated that total fiber length metric reduces the fiber length bias associated with shorter fibers. Our results therefore suggest that average FA may not be the best metric to quantify SC, and that the choice among other SC metrics warrants special attention depending on the question being addressed and the scale of the problem (e.g. whole-brain or local regional analysis).

6.2 Thesis Contributions and Future Work

We now describe the major contributions of this thesis. We proposed a dwMRI data processing pipeline (built on a dictionary learning approach) that enhances the resolution of dwMRI after the data has been acquired. The pipeline does not require modifications to the dwMRI acquisition process, and therefore is more practical in clinical conditions where simple acquisition methods are typically used

Due to the absence of acquisition requirements, the pipeline can be readily used with existing legacy databases of dwMRI data. This can be useful when trying to utilize or reuse datasets which has been acquired with older acquisition technologies. Absence of acquisition requirements also makes it usable with common types of dwMRI, such as DTI or HARDI (single, and multiple-shell). Some of the existing super-resolution approaches are specific to HARDI.

We also note that the proposed pipeline does not require repeated dwMRI acquisitions, unlike classical shifting-based super-resolution methods. Another contribution is that the pipeline performs a resolution enhancement directly on raw dwMRI data, resulting in a model-free pipeline.

This lack of required modeling is useful because it does not restrict the

end user to a particular model, which gives the freedom of using different diffusion models in different research or clinical situations, depending on the specific research conditions or objectives. Finally, we note that the model-independence also allows for *cumulative* enhancements, in which any other diffusion-specific or tractography-specific enhancements can be readily applied on top of the pipeline.

While we demonstrated the benefits of the described pipeline, it is important to acknowledge that the performance of the proposed method inherently depends on the training dataset, as in any machine learning method that involves training or prior information. The age span of the subjects we used in our experiments was 23-61, showing that the method can generalize to a large range of ages. However, how well abnormalities such as tumor and edema can be modeled with dictionary learning is currently unclear and warrants further research in future work.

Bibliography

[1] Iman Aganj, Christophe Lenglet, Guillermo Sapiro, Essa Yacoub, Kamil Ugurbil, and Noam Harel. *Multiple q-shell ODF reconstruction in q-ball imaging*. Springer, 2009.

[2] 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.

[3] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *Signal Processing, IEEE Transactions on*, 54(11):4311–4322, 2006.

[4] 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.

[5] DC Alexander, GJ Barker, and SR Arridge. Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine*, 48(2):331–340, 2002.

[6] Babak A Ardekani, Ali Tabesh, Serge Sevy, Delbert G Robinson, Robert M Bilder, and Philip R Szeszko. Diffusion tensor imaging reliably differentiates patients with schizophrenia from healthy volunteers. *Human brain mapping*, 32(1):1–9, 2011.

[7] Mohammad Bajammal, Burak Yoldemir, and Rafeef Abugharbieh. Comparison of structural connectivity metrics for multimodal brain image analysis. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pages 934–937. IEEE, 2015.

[8] Roland Bammer, Burak Acar, and Michael E Moseley. In vivo mr tractography using diffusion imaging. *European journal of radiology*, 45(3):223–234, 2003.

[9] Peter J Basser, James Mattiello, and Denis LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1):259, 1994.

[10] Sean Borman and Robert L Stevenson. Super-resolution from image sequences-a review. In *Midwest Symposium on Circuits and Systems*, pages 374–378. Notre Dame, Indiana, 1998.

[11] Qingjiu Cao, Ni Shu, Li An, Peng Wang, Li Sun, Ming-Rui Xia, Jin-Hui Wang, Gao-Lang Gong, Yu-Feng Zang, Yu-Feng Wang, et al. Probabilistic diffusion tractography and graph theory analysis reveal abnormal white matter structural connectivity networks in drug-naive boys with attention deficit/hyperactivity disorder. *The Journal of Neuroscience*, 33(26):10676–10687, 2013.

[12] Emmanuel Caruyer, Jian Cheng, Christophe Lenglet, Guillermo Sapiro, Tianzi Jiang, and Rachid Deriche. Optimal design of multiple q-shells experiments for diffusion mri. In *MICCAI Workshop on Computational Diffusion MRI-CDMRI'11*, 2011.

[13] Emmanuel Caruyer, Christophe Lenglet, Guillermo Sapiro, and Rachid Deriche. Design of multishell sampling schemes with uniform coverage in diffusion mri. *Magnetic Resonance in Medicine*, 69(6):1534–1540, 2013.

[14] Hanbo Chen, Kaiming Li, Dajiang Zhu, Xi Jiang, Yixuan Yuan, Peili Lv, Tuo Zhang, Lei Guo, Dinggang Shen, and Tianming Liu. Inferring group-wise consistent multimodal brain networks via multi-view spectral clustering. *Medical Imaging, IEEE Transactions on*, 32(9):1576–1586, 2013.

[15] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1998.

[16] Chen-Huan Chiu, Yu-Chun Lo, Hwa-Sheng Tang, I-Chao Liu, Wen-Yang Chiang, Fang-Cheng Yeh, Fu-Shan Jaw, and Wen-Yih Isaac Tseng. White matter abnormalities of fronto-striato-thalamic circuitry in obsessive–compulsive disorder: a study using diffusion spectrum imaging tractography. *Psychiatry Research: Neuroimaging*, 192(3):176–182, 2011.

[17] Daan Christiaens, Marco Reisert, Thijs Dhollander, Stefan Sunaert, Paul Suetens, and Frederik Maes. Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model. *NeuroImage*, 123:89–101, 2015.

[18] Z Chu, EA Wilde, JV Hunter, SR McCauley, ED Bigler, M Troyanskaya, R Yallampalli, JM Chia, and HS Levin. Voxel-based analysis of diffusion tensor imaging in mild traumatic brain injury in adolescents. *American journal of neuroradiology*, 31(2):340–346, 2010.

[19] Thomas E Conturo, Robert C McKinstry, Erbil Akbudak, and Bruce H Robinson. Encoding of anisotropic diffusion with tetrahedral gradients: a general mathematical diffusion formalism and experimental results. *Magnetic Resonance in Medicine*, 35(3):399–412, 1996.

[20] 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.

[21] 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.

[22] 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.

[23] Ian Ellison-Wright and Ed Bullmore. Meta-analysis of diffusion tensor imaging studies in schizophrenia. *Schizophrenia research*, 108(1):3–10, 2009.

[24] Kjersti Engan, Sven Ole Aase, and John Håkon Husøy. Multi-frame compression: Theory and design. *Signal Processing*, 80(10):2121–2140, 2000.

[25] A Feinstein, P O'Connor, N Akbar, L Moradzadeh, CJM Scott, and NJ Lobaugh. Diffusion tensor imaging abnormalities in depressed multiple sclerosis patients. *Multiple Sclerosis*, 2009.

[26] Mark Fiecas, Hernando Ombao, Dan van Lunen, Richard Baumgartner, Alexandre Coimbra, and Dai Feng. Quantifying temporal correlations: A test–retest evaluation of functional connectivity in resting-state fmri. *Neuroimage*, 65:231–241, 2013.

[27] Lawrence R Frank. Anisotropy in high angular resolution diffusion-weighted mri. *Magnetic Resonance in Medicine*, 45(6):935–939, 2001.

[28] Brendan J Frey and Delbert Dueck. Clustering by passing messages between data points. *science*, 315(5814):972–976, 2007.

[29] Eleftherios Garyfallidis, Matthew Brett, Bagrat Amirbekian, Ariel Rokem, Stefan Van Der Walt, Maxime Descoteaux, Ian Nimmo-Smith, and Dipy Contributors. Dipy, a library for the analysis of diffusion mri data. *Frontiers in neuroinformatics*, 8, 2014.

[30] Elizabeth R Gerstner and A Gregory Sorensen. Diffusion and diffusion tensor imaging in brain cancer. In *Seminars in radiation oncology*, volume 21, pages 141–146. Elsevier, 2011.

[31] Matthew F Glasser, Stamatios N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, et al. The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80:105–124, 2013.

[32] H Greenspan, G Oz, N Kiryati, and S Peled. Mri inter-slice reconstruction using super-resolution. *Magnetic resonance imaging*, 20(5):437–446, 2002.

[33] Donald W Gross. Diffusion tensor imaging in temporal lobe epilepsy. *Epilepsia*, 52(s4):32–34, 2011.

[34] Kanghui Guo and Demetrio Labate. Optimally sparse multidimensional representation using shearlets. *SIAM journal on mathematical analysis*, 39(1):298–318, 2007.

[35] Vikash Gupta, Nicholas Ayache, and Xavier Pennec. Improving dti resolution from a single clinical acquisition: a statistical approach using spatial prior. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, pages 477–484. Springer, 2013.

[36] Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J Honey, Van J Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biol*, 6(7):e159, 2008.

[37] Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J Honey, Van J Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biol*, 6(7):e159, 2008.

[38] Dennis M Healy, Harrie Hendriks, and Peter T Kim. Spherical deconvolution. *Journal of Multivariate Analysis*, 67(1):1–22, 1998.

[39] Christopher J Honey, Jean-Philippe Thivierge, and Olaf Sporns. Can structure predict function in the human brain? *Neuroimage*, 52(3):766–776, 2010.

[40] M Inglese and Maxim Bester. Diffusion imaging in multiple sclerosis: research and clinical implications. *NMR in Biomedicine*, 23(7):865–872, 2010.

[41] Michal Irani and Shmuel Peleg. Improving resolution by image registration. *CVGIP: Graphical models and image processing*, 53(3):231–239, 1991.

[42] Mark Jenkinson, Christian F Beckmann, Timothy EJ Behrens, Mark W Woolrich, and Stephen M Smith. Fsl. *Neuroimage*, 62(2):782–790, 2012.

[43] Derek K Jones. The effect of gradient sampling schemes on measures derived from diffusion tensor mri: a monte carlo studyâĂă. *Magnetic Resonance in Medicine*, 51(4):807–815, 2004.

[44] 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.

[45] Mark H Khachaturian, Jonathan J Wisco, and David S Tuch. Boosting the sampling efficiency of q-ball imaging using multiple wavevector fusion. *Magnetic Resonance in Medicine*, 57(2):289–296, 2007.

[46] Kwang In Kim and Younghee Kwon. Single-image super-resolution using sparse regression and natural image prior. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(6):1127–1133, 2010.

[47] Nicolas Kunz, Hui Zhang, Lana Vasung, Kieran R O'Brien, Yaniv Assaf, François Lazeyras, Daniel C Alexander, and Petra S Hüppi. Assessing white matter microstructure of the newborn with multi-shell diffusion mri and biophysical compartment models. *Neuroimage*, 96:288–299, 2014.

[48] Li-Wei Kuo, Jyh-Horng Chen, Van Jay Wedeen, and Wen-Yih Isaac Tseng. Optimization of diffusion spectrum imaging and q-ball imaging on clinical mri system. *Neuroimage*, 41(1):7–18, 2008.

[49] Gitta Kutyniok and Wang-Q Lim. Compactly supported shearlets are optimally sparse. *Journal of Approximation Theory*, 163(11):1564–1589, 2011.

[50] Demetrio Labate, Wang-Q Lim, Gitta Kutyniok, and Guido Weiss. Sparse multidimensional representation using shearlets. In *Optics & Photonics 2005*, pages 59140U–59140U. International Society for Optics and Photonics, 2005.

[51] Bennett A Landman, Alan J Huang, Aliya Gifford, Deepti S Vikram, Issel Anne L Lim, Jonathan AD Farrell, John A Bogovic, Jun Hua, Min Chen, Samson Jarso, et al. Multi-parametric neuroimaging reproducibility: a 3-t resource study. *Neuroimage*, 54(4):2854–2866, 2011.

[52] Denis Le Bihan. Looking into the functional architecture of the brain with diffusion mri. In *International Congress Series*, volume 1290, pages 1–24. Elsevier, 2006.

[53] Harvey S Levin, Elisabeth Wilde, Maya Troyanskaya, Nancy J Petersen, Randall Scheibel, Mary Newsome, Majdi Radaideh, Trevor Wu, Ragini Yallampalli, Zili Chu, et al. Diffusion tensor imaging of mild to moderate blast-related traumatic brain injury and its sequelae. *Journal of neurotrauma*, 27(4):683–694, 2010.

[54] Xinghua Li, Huanfeng Shen, Liangpei Zhang, Hongyan Zhang, Qiangqiang Yuan, and Gang Yang. Recovering quantitative remote sensing products contaminated by thick clouds and shadows using multitemporal dictionary learning. *Geoscience and Remote Sensing, IEEE Transactions on*, 52(11):7086–7098, 2014.

[55] Ching-Po Lin, Van Jay Wedeen, Jyh-Horng Chen, Ching Yao, and Wen-Yih Isaac Tseng. Validation of diffusion spectrum magnetic resonance imaging with manganese-enhanced rat optic tracts and ex vivo phantoms. *Neuroimage*, 19(3):482–495, 2003.

[56] Yawu Liu, Gabriela Spulber, Kimmo K Lehtimäki, Mervi Könönen, Ilona Hallikainen, Heidi Gröhn, Miia Kivipelto, Merja Hallikainen, Ritva Vanninen, and Hilkka Soininen. Diffusion tensor imaging and tract-based spatial statistics in alzheimer's disease and mild cognitive impairment. *Neurobiology of aging*, 32(9):1558–1571, 2011.

[57] Yu-Chun Lo, Wei-Tsuen Soong, Susan Shur-Fen Gau, Yu-Yu Wu, Meng-Chuan Lai, Fang-Cheng Yeh, Wen-Yang Chiang, Li-Wei Kuo, Fu-Shan Jaw, and Wen-Yih Isaac Tseng. The loss of asymmetry and reduced interhemispheric connectivity in adolescents with autism: a study us-

ing diffusion spectrum imaging tractography. *Psychiatry Research: Neuroimaging*, 192(1):60–66, 2011.

[58] Stephan E Maier, Yanping Sun, and Robert V Mulkern. Diffusion imaging of brain tumors. *NMR in biomedicine*, 23(7):849–864, 2010.

[59] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 689–696. ACM, 2009.

[60] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 689–696. ACM, 2009.

[61] Rogier B Mars, Saad Jbabdi, Jérôme Sallet, Jill X O'Reilly, Paula L Croxson, Etienne Olivier, MaryAnn P Noonan, Caroline Bergmann, Anna S Mitchell, Mark G Baxter, et al. Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity. *The Journal of Neuroscience*, 31(11):4087–4100, 2011.

[62] C Meca, S Chabert, and D Le Bihan. Diffusion mri at large b values: whatâĂŹs the limit. In *Proceedings of the 12th Annual Meeting of ISMRM, Kyoto, Japan*, page 1196, 2004.

[63] Vincent Michel, Alexandre Gramfort, Gaël Varoquaux, Evelyn Eger, Christine Keribin, and Bertrand Thirion. A supervised clustering approach for fmri-based inference of brain states. *Pattern Recognition*, 45(6):2041–2049, 2012.

[64] Vincent Michel, Alexandre Gramfort, Gaël Varoquaux, Evelyn Eger, Christine Keribin, and Bertrand Thirion. A supervised clustering ap-

proach for fmri-based inference of brain states. *Pattern Recognition*, 45(6):2041–2049, 2012.

[65] Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.

[66] Shahrum Nedjati-Gilani, Daniel C Alexander, and Geoff 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.

[67] P Neher, Bram Stieltjes, Ivo Wolf, H Meinzer, and K Maier-Hein. Analysis of tractography biases introduced by anisotropic voxels. In *Proc. Annual Meeting ISMRM*, 2013.

[68] Peter F Neher, Bram Stieltjes, Marco Reisert, Ignaz Reicht, Hans-Peter Meinzer, and Klaus H Fritzsche. Mitk global tractography. In *SPIE medical imaging*, pages 83144D–83144D. International Society for Optics and Photonics, 2012.

[69] Weihua Ou, Xinge You, Dacheng Tao, Pengyue Zhang, Yuanyan Tang, and Ziqi Zhu. Robust face recognition via occlusion dictionary learning. *Pattern Recognition*, 47(4):1559–1572, 2014.

[70] Slavche Pejoski, Venceslav Kafedziski, and Dusan Gleich. Compressed sensing mri using discrete nonseparable shearlet transform and fista. *Signal Processing Letters, IEEE*, 22(10):1566–1570, 2015.

[71] 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.

[72] Adolf Pfefferbaum, Edith V Sullivan, Maj Hedehus, Kelvin O Lim, Elfar Adalsteinsson, and Michael Moseley. Age-related decline in

brain white matter anisotropy measured with spatially corrected echo-planar diffusion tensor imaging. *Magnetic resonance in medicine*, 44(2):259–268, 2000.

[73] Dirk HJ Poot, Ben Jeurissen, Yannick Bastiaensen, Jelle Veraart, Wim Van Hecke, Paul M Parizel, and Jan Sijbers. Super-resolution for multislice diffusion tensor imaging. *Magnetic Resonance in Medicine*, 69(1):103–113, 2013.

[74] Vesna Prčkovska, Alard F Roebroeck, WLPM Pullens, Anna Vilanova, and Bart M ter Haar Romeny. Optimal acquisition schemes in high angular resolution diffusion weighted imaging. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2008*, pages 9–17. Springer, 2008.

[75] Klaus Scheffler. Superresolution in mri? *Magnetic resonance in medicine*, 48(2):408–408, 2002.

[76] 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.

[77] Claire E Sexton, Ukwuori G Kalu, Nicola Filippini, Clare E Mackay, and Klaus P Ebmeier. A meta-analysis of diffusion tensor imaging in mild cognitive impairment and alzheimer's disease. *Neurobiology of aging*, 32(12):2322–e5, 2011.

[78] Ling Shao, Ruomei Yan, Xuelong Li, and Yan Liu. From heuristic optimization to dictionary learning: a review and comprehensive comparison of image denoising algorithms. *Cybernetics, IEEE Transactions on*, 44(7):1001–1013, 2014.

[79] Richard Z Shilling, Marijn E Brummer, and Klaus Mewes. Merging multiple stacks mri into a single data volume. In *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on*, pages 1012–1015. IEEE, 2006.

[80] Richard Z Shilling, Senthil Ramamurthy, and Marijn E Brummer. Sampling strategies for super-resolution in multi-slice mri. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pages 2240–2243. IEEE, 2008.

[81] 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.

[82] Bertrand Thirion, Gaël Varoquaux, Elvis Dohmatob, and Jean-Baptiste Poline. Which fmri clustering gives good brain parcellations? *Frontiers in neuroscience*, 8, 2014.

[83] Benjamin Thyreau, Bertrand Thirion, Guillaume Flandin, and Jean-Baptiste Poline. Anatomo-functional description of the brain: a probabilistic approach. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, volume 5, pages V–V. IEEE, 2006.

[84] J-Donald Tournier, Fernando Calamante, and Alan Connelly. Robust determination of the fibre orientation distribution in diffusion mri: non-negativity constrained super-resolved spherical deconvolution. *NeuroImage*, 35(4):1459–1472, 2007.

[85] 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.

[86] J-Donald Tournier, Chun-Hung Yeh, Fernando Calamante, Kuan-Hung Cho, Alan Connelly, and Ching-Po Lin. Resolving crossing fibres using constrained spherical deconvolution: validation using diffusion-weighted imaging phantom data. *Neuroimage*, 42(2):617–625, 2008.

[87] David S Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[88] 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.

[89] 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.

[90] Hanoch Ur and Daniel Gross. Improved resolution from subpixel shifted pictures. *CVGIP: Graphical Models and Image Processing*, 54(2):181–186, 1992.

[91] David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, WU-Minn HCP Consortium, et al. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.

[92] Zhijiang Wang, Zhengjia Dai, Gaolang Gong, Changsong Zhou, and Yong He. Understanding structural-functional relationships in the human brain a large-scale network perspective. *The Neuroscientist*, 21(3):290–305, 2015.

[93] 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.

[94] 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.

[95] David J Werring, Ahmed T Toosy, Christopher A Clark, Geoffrey JM Parker, Gareth J Barker, David H Miller, and Alan J Thompson. Diffusion tensor imaging can detect and quantify corticospinal tract degeneration after stroke. *Journal of Neurology, Neurosurgery & Psychiatry*, 69(2):269–272, 2000.

[96] Gavin P Winston, Laura Mancini, Jason Stretton, Jonathan Ashmore, Mark R Symms, John S Duncan, and Tarek A Yousry. Diffusion tensor imaging tractography of the optic radiation for epilepsy surgical planning: a comparison of two methods. *Epilepsy research*, 97(1):124–132, 2011.

[97] Jianchao Yang, John Wright, Thomas S Huang, and Yi Ma. Image super-resolution via sparse representation. *Image Processing, IEEE Transactions on*, 19(11):2861–2873, 2010.

[98] Burak Yoldemir, Bernard Ng, and Rafeef Abugharbieh. Effects of tractography approach on consistency between anatomical and functional connectivity estimates. In *Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on*, pages 250–253. IEEE, 2014.

[99] Dajiang Zhu, Tuo Zhang, Xi Jiang, Xintao Hu, Hanbo Chen, Ning Yang, Jinglei Lv, Junwei Han, Lei Guo, and Tianming Liu. Fusing dti and fmri data: a survey of methods and applications. *Neuroimage*, 102:184–191, 2014.