PHYSICS-CONSTRAINED UNSUPERVISED DEEP LEARNING FOR ACCELERATED DIFFUSION MRI

A THESIS SUBMITTED TO

THE GRADUATE SCHOOL OF ENGINEERING AND SCIENCE

OF BILKENT UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR

THE DEGREE OF

MASTER OF SCIENCE

IN

ELECTRICAL AND ELECTRONICS ENGINEERING

By

Atakan Topcu

August 2025

PHYSICS-CONSTRAINED UNSUPERVISED DEEP LEARNING FOR ACCELERATED DIFFUSION MRI

By Atakan Topcu

August 2025

We certify that we have read this thesis and that in our opinion it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Science.

Emine Ülkü Sarıtaş Çukur (Advisor)

Tolga Çukur

Esin Öztürk Işık

Approved for the Graduate School of Engineering and Science:

Orhan Arıkan
Director of the Graduate School

ABSTRACT

PHYSICS-CONSTRAINED UNSUPERVISED DEEP LEARNING FOR ACCELERATED DIFFUSION MRI

Atakan Topcu
M.S. in Electrical and Electronics Engineering
Advisor: Emine Ülkü Sarıtaş Çukur
August 2025

Diffusion Magnetic Resonance Imaging (dMRI) is a noninvasive technique that probes the microscopic Brownian movement of water molecules within neural tissues, providing insights into the underlying microstructural architecture. In dMRI, the displacement of spins is encoded in a domain called q-space through the use of diffusion-sensitizing gradients. Classical dMRI models, such as diffusion tensor imaging (DTI), require only a few samples in q-space, but fall short in resolving crossing or diverging fiber bundles. To address these limitations, High Angular Resolution Diffusion Imaging (HARDI) was introduced to enhance fiber characterization by densely sampling the q-space across multiple spherical shells defined by different b-values, thereby detecting several fiber orientations within a single voxel. Building on this framework, advanced multi-shell techniques such as Multi-Shell Spherical Deconvolution (MSMT-CSD) and Neurite Orientation Dispersion and Density Imaging (NODDI) have been developed, offering refined insights into complex microstructural features. Nevertheless, the requirement for densely sampling q-space renders advanced dMRI techniques extremely time-consuming and impractical for clinical use. This thesis proposes a deep unsupervised **Q**-space **U**psampling via physics-**C**onstrained **C**oordinate-based **I**mplicit network (QUCCI) to accelerate multi-shell dMRI. QUCCI models the underlying volume as a continuous function in both spatial coordinates and q-space, enabling the sampling of q-space along arbitrary directions without the constraints of fixed sampling schemes. An encoder maps coordinates to a latent code, and an MLP predicts the signal, allowing arbitrary q-space sampling without large training datasets or vendor harmonization. Physics-based regularization stabilizes learning. Tested on 10 subjects at R = 10, 15, 22.5, and extended to joint q-space interpolation plus in-plane super-resolution for submillimeter whole-brain dMRI, QUCCI surpasses a recent deep-learning competitor,

a least-squares baseline, and raw undersampled data. Slice-, subject-, and metric-level evaluations, and downstream DTI, MSMT-CSD, and NODDI maps confirm its superior fidelity. QUCCI enables accelerated dMRI with minimal information loss, advancing the clinical feasibility of advanced multi-shell methods.

Keywords: Diffusion magnetic resonance imaging, multi-shell diffusion imaging, q-space, deep learning, neural fields, unsupervised learning.

ÖZET

HIZLANDIRILMIŞ DİFÜZYON MRG İÇİN FİZİK KISITLAMALI DENETİMSİZ DERİN ÖĞRENME

Atakan Topcu
Elektrik ve Elektronik Mühendisliği, Yüksek Lisans
Tez Danışmanı: Emine Ülkü Sarıtaş Çukur
Ağustos 2025

Difüzyon Manyetik Rezonans Görüntüleme (dMRG), sinir dokusundaki su moleküllerinin mikroskobik Brown hareketini inceleyerek mikroyapısal mimariyi ortaya koyan bir tekniktir. Spin yer değiştirmeleri, difüzyon duyarlı gradyanlarla q-uzayında kodlanır. Klasik Difüzyon Tensör Görüntüleme (DTG) az sayıda q-uzayı örneğiyle tek doğrultulu lifleri tanımlar, ancak liflerin kesiştiği bölgelerde yetersizdir. Bu eksikliği gidermek için, farklı b-değerleriyle tanımlanan küresel kabukları yoğun biçimde örnekleyen Yüksek Açısal Çözünürlüklü Difüzyon Görüntüleme (HARDI) geliştirilmiştir. HARDI'nin çok kabuklu uzantıları olan Çoklu-Doku Çoklu-Kabuk Kısıtlı Küresel Ters Evrişim (ÇDÇK-KKTE) ve Nörit Oryantasyon Dağılımı ve Yoğunluk Görüntüleme (NODDI) karmaşık lif mimarisini daha ayrıntılı betimler, ancak yoğun q-uzayı örneklemesi klinik uygulanabilirliği kısıtlayan uzun çekim süreleri gerektirir. Bu tez, çok kabuklu dMRG'yi hızlandırmak için "Q-Uzayı Yukarı Örneklemesi: Fiziksel Kısıtlamalı Koordinat Tabanlı Örtük Ağ" (QUCCI) adlı gözetimsiz, derin öğrenme tabanlı bir yaklaşım sunmaktadır. QUCCI, hacmi hem uzaysal hem de q-uzayı koordinatlarının sürekli bir fonksiyonu olarak modeller; böylece sabit şemalara bağlı kalmadan istenen yönde örnekleme yapılabilir. Bir kodlayıcı koordinatları yüksek boyutlu gizli koda çevirir, çok katmanlı algılayıcı (ÇKAA) ise sinyali tahmin eder. Yöntem yalnızca koordinat girişi kullandığından büyük eğitme kümelerine veya cihaz uyumlaştırmasına gerek duymaz; fizik temelli düzenleme kararlı yakınsamayı sağlar. QUCCI, 10 gönüllüde R = 10, 15, 22.5 hızlandırma seviyeleri ve q-uzayı enterpolasyonuna ek olarak düzlem-içi süper çözünürlükle milimetre altı çözünürlükte tam beyin dMRG'de sınanmıştır. Sonuç olarak QUCCI, güncel bir derin öğrenme yöntemine, en küçük kareler temelli klasik yaklaşıma ve doğrudan alt örnekleme ile elde edilen veriye göre üstünlük göstermiştir. Dilim, denek ve

metrik düzeyindeki değerlendirmeler ile DTG, ÇDÇK-KKTE ve NODDI haritaları da bu üstünlüğü doğrulamaktadır. QUCCI yaklaşımının bir uzantısında ise, q-uzayı içinde yapılan enterpolasyonun, düzlem içi süper çözünürlükle birleştirilerek tüm beyni kapsayan alt milimetre çözünürlüklü dMRG elde edilmesi hedeflenmiştir. Sonuç olarak QUCCI, bilgi kaybını en aza indirerek dMRG'yi klinik açıdan uygulanabilir hızlara taşımaktadır.

Anahtar sözcükler: Difüzyon manyetik rezonans görüntüleme, çoklu-kabuk difüzyon görüntüleme, q-uzayı, derin öğrenme, nöral alan, denetimsiz öğrenme.

Acknowledgement

Before I start, I would like to acknowledge that this thesis was supported Scientific and Technological Research Council of Turkey (TUBITAK Grant No: 122E162).

To say the least, I've never been one to relish writing as I always leaned on Prof. Saritas' feedback for that. However, writing this section has been a blast. It's my chance to thank the extraordinary people who've carried me through these past years.

I could not have asked for a more inspiring advisor than Prof. Emine Ulku Saritas. You are the kindest, most patient soul I know. I still remember my freshman self pacing outside your office, desperately hoping you'd take a chance on an eager undergrad. From day one, your door was always open even when your schedule was buried under a mountain of emails, and you never tired of my endless questions. You taught me to think critically, to persist when every experiment went haywire, and to find humor in the toughest moments. I hope to embody even a fraction of your grace and wisdom. Without your mentorship, I wouldn't be standing here today.

Working with Prof. Tolga Çukur has also been a honor. Your blend of mathematical rigor and engineering vision reshaped my approach to every problem. You were the one who first demystified deep learning for me, showing that beneath the jargon lies surprisingly intuitive principles. I'll forever admire your professionalism and vision.

A huge thank you to my SARITAS lab family, Abdallah, Sevil, Ege, Merve, and Zeynep for turning this lab into a home. Abdallah, your expert deep-learning tips (and timely jokes) kept me going and bolstered my confidence. Sevil, your evolution into a brilliant researcher has been truly inspiring. Ege, your relentless optimism grounded me, and Zeynep, your travel tales and punchlines made even the longest office days a joy. I can't wait to see where your journeys take you all.

I would also like to thank former lab members. To Mustafa and Aslı, thank

you for helping me beat my anxiety when I first joined the lab. I actually knew Mustafa for a long time since I joined the lab as an undergraduate. I was so scared to get bad results and I nearly quit research. Without his optimism and support, I don't think I could have continued it. Even to this day, he still checks on me. He is an amazing friend and an amazing mentor. After Mustafa, Aslı was the one that helped me get into master's life. Her patience with me still continues as she helps me get accustomed to life in UC Berkeley. I sincerely thank her for being patient with me and hope to see our friendship continue in Berkeley.

To my friends Hasan, Atakan, and Emirhan, thank you for keeping me grounded. Whether we were trying to fix a broken code or celebrating a breakthrough, your belief in me never wavered, and I've made lifelong friends I'll always cherish. Hasan, thank you for caring for me when I was sick or down, for sharing my burdens, and for your infectious optimism. Your presence reminded me I'm never alone. You've truly found your people in us, and you've got a friend in me. Atakan, I'm so glad you finally accepted to hang out with me last year (after twenty tries!). Calling you my brother feels right, your childlike wonder, sincerity, and eagerness to help have pushed me to be better, and I owe much of my progress to you. I can't wait for more of the Bay Area adventures together. Emirhan, you were my first friend here and helped me through the toughest early days of my master's. Your enthusiasm for research and comforting company kept me going despite anxiety and setbacks. You and Ecem have brought so much joy and laughter into my life. Of all, Elif is the one whose kindness I can never adequately thank. At every turn she encouraged me, offered thoughtful feedback on my research, lifted my spirits in difficult moments, and celebrated my successes with genuine joy. I'm grateful to have her in my life and look forward to sharing countless moments together. Here's to more memories, California dreaming together!

I would also like to thank my friends, Ata Yavuzyılmaz, Saitcan, Eda, Kerem, Ecem, Yağızalp, Mehmet Emin, Baykam, Cemre, Ata Korkusuz, Barış, Berke and Özgür. Thank you for supporting me thoughout years. Even though you were physically far away, it seemed so you were here with me all the time. Looking back, 10 years old Atakan would have never guessed he would have that many

friends. I am so lucky my life turned around out to be this way.

Mom and Dad, I wouldn't be here today without your sacrifices. You have always supported me, encouraged me and provided me with the means to pursue the future I want. I hope this achievement makes you as proud as I am proud to be your son.

I owe a piece of this work to each of you. From the bottom of my heart, thank you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Diffusion Magnetic Resonance Imaging (dMRI) is a non-invasive technique that tracks the Brownian motion of water molecules in living tissue [2, 3]. Within neural tissue, axonal membranes and surrounding microstructures constrain this motion, so water preferentially diffuses either *along* fiber tracts or through the spaces between them. Because diffusion behavior mirrors the local environment, dMRI can reveal the organization of white-matter axons and provide complementary insights into gray matter and cerebrospinal fluid. MRI pulse sequences are sensitized to diffusion by applying a pair of gradient pulses. The resulting phase dispersion attenuates the signal in proportion to random Brownian motion, thereby encoding microscopic displacements. By tuning these diffusion gradients, dMRI quantifies fine-scale details, such as fiber orientation, axonal density, and tissue integrity, at spatial scales well below conventional MRI image resolution [2]. Consequently, dMRI has become essential for mapping brain connectivity, monitoring developmental changes, and diagnosing neurodegenerative disorders [4, 5].

Early clinical practice has largely relied on diffusion tensor imaging (DTI), which represents diffusion in each voxel with a single second-order tensor estimated from at least six non-collinear diffusion-weighted directions (i.e., q-space directions) and one non-diffusion-weighted (b_0) volume [6, 7]. The eigenvalues and eigenvectors of the tensor yield biomarkers such as fractional anisotropy

(FA) and mean diffusivity (MD) that are sensitive to demyelination, edema, and axonal loss. Because DTI assumes one dominant fiber orientation per voxel, however, it collapses information in regions where multiple tracts cross, branch, or fan, thereby masking true pathway geometry and artificially lowering FA [8].

High-angular-resolution diffusion imaging (HARDI) overcomes this limitation by densely sampling the diffusion signal in q-space [2, 9]. Modern HARDI protocols acquire data on several concentric "shells," where each shell corresponds to a different b-value and therefore a different radius in q-space. Higher b-values are achieved by increasing amplitude, duration, or separation of the diffusion gradient pulses, which in turn amplifies sensitivity to restricted intra-axonal diffusion but lowers signal-to-noise ratio (SNR) [2]. Model-free reconstructions such as constrained multi-shell multi-tissue spherical deconvolution (MSMT-CSD) treat the measured signal as the convolution of a single-fiber response with an orientation distribution function (ODF) and estimate multiple fiber peaks by non-negative deconvolution, whereas model-based approaches such as neurite orientation dispersion and density imaging (NODDI) fit a three-compartment biophysical model to derive intracellular volume fraction (ICVF) and orientation dispersion indices (ODI) [10, 11]. Both MSMT-CSD and NODDI require the acquisition of dozens to hundreds of gradient directions across low- and high-b shells to achieve numerical stability and biological specificity. Furthermore, additional reverse phase-encoded acquisitions also need to be acquired for robust susceptibility correction. As a result, the total scan times can easily exceed thirty minutes, rendering comprehensive multi-shell HARDI impractical for routine clinical workflows [12, 13].

Sub-millimeter dMRI pushes these demands even further. Generalized Slice Dithered Enhanced Resolution (gSlider) acquires thick RF-encoded slabs and subsequently super-resolves them into thin slices, offering an attractive compromise between spatial resolution and SNR [14]. Yet, when gSlider is combined with dense q-space sampling and repeated averages, total acquisition times can approach one to two hours, precluding routine clinical use or studies involving motion-prone cohorts [15]. These extended protocols increase the risk of motion-related artifacts and patient discomfort, which can degrade data quality [16].

Consequently, there is a pressing need for acceleration methods that achieve sub-millimeter resolution while dramatically reducing scan duration.

This thesis seeks to accelerate dMRI using a physics-inspired unsupervised deep learning method. This thesis proposes a unified, self-supervised neural-field framework that combines implicit coordinate representations with explicit physical priors to reconstruct high-resolution, high-angular-resolution and high-radial-resolution dMRI data from highly accelerated dMRI acquisitions. Implicit neural representations can learn a continuous mapping from voxel coordinates to diffusion signals of a *single* subject, eliminating the need for large training datasets. The proposed method, QUCCI (Q-space Upsampling with physics-Constrained Coordinate-based Implicit network), introduces an unsupervised framework that jointly enforces angular and radial continuity, achieving high-fidelity interpolation from as few as 12 acquired gradient directions. Building on this idea, this thesis further proposes sq-QUCCI that cascades two neural-field modules, one operating along the RF-encoding dimension, the other along the q-space dimension, to super-resolve undersampled gSlider data. Whole-brain 500 µm dMRI can thus be acquired in approximately fifteen minutes with seven-fold acceleration. Comprehensive evaluations on in vivo brain data demonstrate substantial scan time reductions while maintaining the accuracy of DTI, MSMT-CSD, and NODDI metrics.

Table 1.1: Glossary of Acronyms

Acronym	**Meaning / Definition**
AFD	Apparent Fiber Density
ADC	Apparent Diffusion Coefficient (mm^2/s)
B_0	Main static magnetic field (Tesla)
B_1	Radio-frequency magnetic field
CAIPI	Controlled Aliasing in Parallel Imaging
CSD	Constrained Spherical Deconvolution
CSF	Cerebrospinal Fluid
dMRI	Diffusion Magnetic Resonance Imaging
DTI	Diffusion Tensor Imaging
EPI	Echo-Planar Imaging
FA	Fractional Anisotropy
fODF	fiber Orientation Distribution Function
g-factor	Geometry factor (noise amplification)
GRAPPA	Generalized Autocalibrating Partially Parallel Acquisitions
HARDI	High Angular Resolution Diffusion Imaging
ICVF	Intra-cellular (neurite) Volume Fraction
MB	Multiband (Simultaneous Multi-Slice)
MSMT-CSD	Multi-Shell Multi-Tissue CSD
NODDI	Neurite Orientation Dispersion and Density Imaging
nuFO	Number of Fiber Orientations
ODF	Orientation Distribution Function
ODI	Orientation Dispersion Index
q-space	Wave-vector space encoding diffusion attenuation
RF	Radio Frequency
SAR	Specific Absorption Rate
SENSE	Sensitivity Encoding (image-space parallel MRI)
SH	Spherical Harmonics
SMS	Simultaneous Multi-Slice imaging
TE	Echo Time (ms)
TR	Repetition Time (ms)
gSlider	generalized SLIce Dithered Enhanced Resolution

Chapter 2

Background

2.1 Diffusion MRI Fundamentals

Diffusion magnetic resonance imaging (dMRI), often referred to as diffusion-weighted imaging (DWI), is an *in vivo*, non-invasive method that measures the microscopic motion of water molecules inside biological tissue without perturbing that motion [6]. Because water displacement patterns change with cellular architecture and pathology, dMRI is able to probe tissue microstructure in both healthy and diseased states.

Clinically, the modality is invaluable for the diagnosis of numerous conditions, such as the early detection of acute ischemic stroke [17, 18, 19], lymph node staging [20], liver neoplasms [21], renal lesions [22, 23], prostate tumors [24], and colorectal cancer [25]. In addition to its clinical use, dMRI has also become popular in neuroscience for studying white matter microstructure and for reconstructing fiber pathways via tractography for brain connectivity and neurodevelopmental monitoring [26, 27, 28].

2.1.1 Diffusion in Tissue

In still water, and with no outside forces acting on it, each water molecule moves randomly due to thermal energy. If we look at an ensemble of molecules, their spread of positions follows a bell-shaped (Gaussian) curve with zero average movement [29]. In n dimensions, the average squared distance traveled after time t is

$$\langle r^2 \rangle \;=\; 2\,n\,D\,t, \tag{2.1}$$

where D is the diffusion coefficient (in mm^2/s), a constant that comes from Fick's laws of diffusion [30]. Equation (2.1) is the multi-dimensional form of the classic Einstein–Smoluchowski equation and describes free, isotropic (i.e., equal-in-all-directions) diffusion.

Most tissues do not let water molecules move freely and equally in all directions. In the brain, for instance, cell membranes, myelin, and other tiny structures change how water molecules can move: outside the cells movement is hindered, while inside it can be restricted in certain directions [31]. Because of these barriers, measured diffusion is anisotropic, i.e., it depends on direction. To describe this behavior, researchers use models that range from the simple diffusion tensor [6] to more detailed multi-compartment methods such as Neurite Orientation Dispersion and Density Imaging (NODDI) [11]. These models provide useful biomarkers that inform us about the tissue microstructure and health.

2.1.2 Pulse Sequence Sensitization

The classic way to "see" diffusion in MRI is to add a pair of gradient pulses to a spin echo sequence, an idea introduced in the 1950s and later refined by Stejskal and Tanner [32, 33, 34]. An illustration of a basic spin-echo diffusion sequence can be seen in Fig. 2.1. When the first gradient is on, each spin accrues up a position-dependent phase. If a spin stays perfectly still, the second gradient cancels that phase and a coherent signal reappears at the echo time (TE). Spins undergoing random movement, however, see a slightly different magnetic field during the

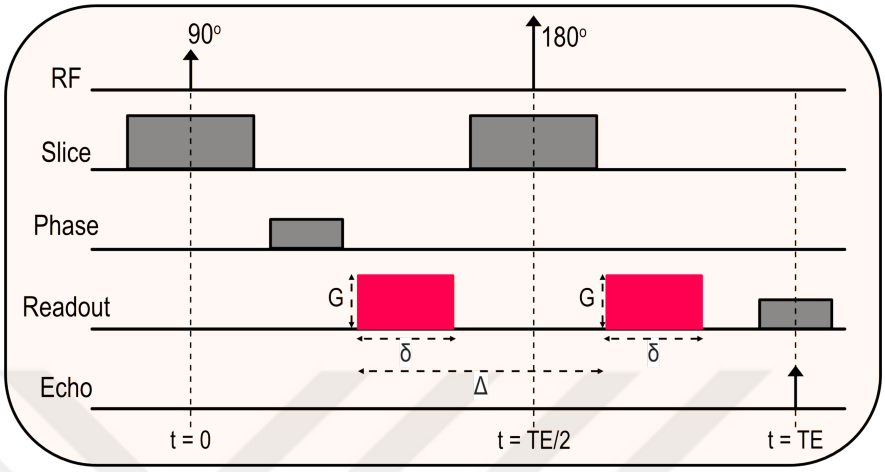

Figure 2.1: Basic spin-echo diffusion MRI sequence. Two strong diffusion-sensitizing gradients (red) with amplitude G and duration δ are applied on either side of the 180° RF pulse, separated by the diffusion time Δ. Changing G, δ, Δ or the direction of the diffusion-sensitizing gradients controls the amount and orientation of diffusion weighting. The echo is sampled at the echo time TE.

second gradient pulse, so that their phase is not fully canceled. Because millions of spins move randomly during diffusion, their phases add up randomly. While the average phase remains zero, the net MRI signal is reduced. This signal loss is the essence of diffusion contrast.

For a single spin following a path, $\mathbf{x}(t)$, the accumulated phase at the echo time TE can be expressed as:

$$\phi(TE) = \gamma\Big[\int_0^{TE/2} \mathbf{G}(t)\cdot\mathbf{x}(t)\,dt - \int_{TE/2}^{TE} \mathbf{G}(t)\cdot\mathbf{x}(t)\,dt\Big], \tag{2.2}$$

where γ is the gyromagnetic ratio and $\mathbf{G}(t)$ is the diffusion-sensitizing gradient. Because the spin displacements $\mathbf{x}(t)$ in Equation (2.2) are Gaussian distributed, the resulting phase $\phi(TE)$ also follows a Gaussian distribution with zero mean and variance σ_ϕ^2 [33] (see Subsection 2.1.1). To compute the diffusion-weighted signal relative to the unweighted signal, we average the complex phase factor $e^{i\phi(TE)}$ over all spins, i.e.,

$$S = S_0 \int_{-\infty}^{\infty} f_\phi e^{i\phi}\,d\phi = S_0\; e^{-\frac{\sigma_\phi^2}{2}} = S_0\; e^{-bD}, \tag{2.3}$$

where f_ϕ is the probability density function (PDF) of ϕ, S_0 is the signal in the absence of diffusion-sensitizing gradients, b refers to the so-called "b-value" (in

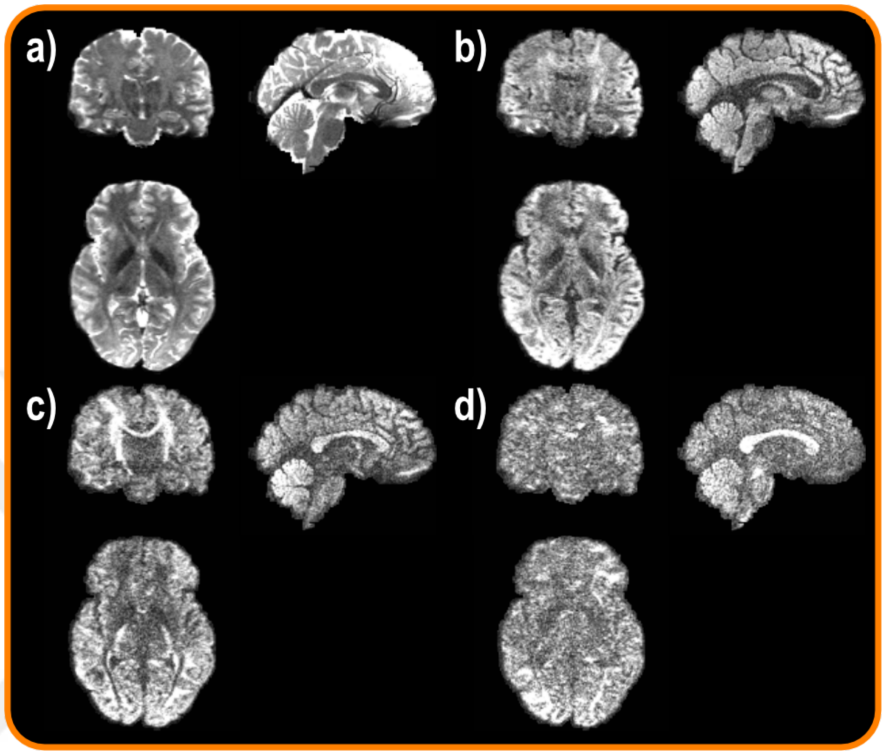

Figure 2.2: Example brain dMRI images at four different b-values, shown in coronal (top left), sagittal (top right), and axial (bottom left) orientations for: (a) $b = 0$ s/mm^2, (b) $b = 1000$ s/mm^2, (c) $b = 2000$ s/mm^2, (d) $b = 3000$ s/mm^2. Images are from the preprocessed dMRI data of the Human Connectome Project [1].

s/mm^2), and D is the apparent diffusion coefficient (ADC) (in mm^2/s). Because the phase shifts are distributed in a Gaussian fashion, they interfere destructively, producing an exponential drop in signal amplitude while leaving the average phase unchanged [35].

In Equation (2.3), the b-value summarizes how "strong" the diffusion encoding is. For rectangular gradient lobes (see Fig. 2.1), it can be expressed as:

$$b = q^2\left(\Delta - \tfrac{\delta}{3}\right), \tag{2.4}$$

where

$$q = \gamma \Delta G \tag{2.5}$$

Here, q is referred to as the wavevector. By varying the direction and amplitude of the gradient pulse, we sample different points in the so called q-space, the domain in which diffusion-induced signal attenuation is encoded and measured.

According to Equation (2.4), stronger gradients, longer gradient pulses, or a longer separation of gradient pulses all increase diffusion weighting. Images acquired with $b = 0$ s/mm^2 give an S_0 reference; images with one or more non-zero b-values and gradient directions encode diffusion along the selected axes. Combining these measurements with suitable models (e.g., ADC, DTI, NODDI) allows us to extract rich information about tissue microstructure.

Figure 2.2 shows how increasing the b-value progressively emphasizes restricted diffusion in brain white matter, illustrating the practical effect of the b-value on image contrast.

2.1.3 Diffusion MRI Acquisition Strategies

Echo Planar Imaging: Diffusion gradients introduce sensitivity to motion at the micrometer scale. When global rigid motion occurs during these diffusion-sensitizing gradients, it introduces additional phase errors and shifts in k-space, adversely affecting the image quality. The sensitivity to subject motion diminishes the effectiveness of standard imaging sequences such as line-by-line acquisitions [36]. Examples of such global movements include subject head motion or breathing-related motion [37]. Additionally, diffusion microstructural modeling often requires multiple scans at varying b-values and/or diffusion directions. These demands necessitate sequences that are robust to motion, maintain high signal-to-noise ratio (SNR) efficiency, and achieve short scanning durations suitable for clinical environments [36]. The most widely adopted sequence meeting these criteria for diffusion MRI is echo planar imaging (EPI) [38]. Unlike traditional line-by-line acquisitions, EPI traverses k-space in a zigzag trajectory, rapidly capturing the entire k-space data.

As shown in Fig. 2.3(a), single-shot EPI (ss-EPI) acquires the entire k-space data after a single RF excitation. This characteristic makes ss-EPI particularly robust against motion during the diffusion-sensitizing gradients [39]. Although widely used for dMRI, ss-EPI is prone to specific artifacts stemming from its strong readout gradients and relatively low bandwidth in the phase-encoding

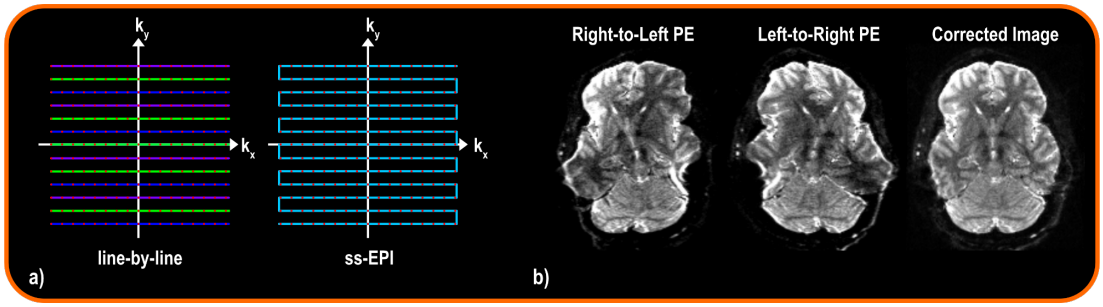

Figure 2.3: (a) In a conventional acquisition, each line of k-space is collected sequentially, with a separate RF excitation for each line. In single-shot EPI, the entire k-space data is acquired in one continuous readout following a single RF excitation. (b) Example images in the axial plane, showing geometric distortions in opposing directions for EPI images acquired with reversed PE directions (right-to-left and left-to-right in this example). These two EPI images are processed by FSL's TOPUP to estimate a displacement field and apply distortion correction. Images are from the preprocessed dMRI data of the Human Connectome Project [1].

(PE) direction. Two of the most disruptive artifacts for ss-EPI are susceptibility-induced geometric distortion and eddy-current–induced geometric distortion.

Susceptibility artifacts occur due to the prolonged readout duration in EPI sequences that reduce the readout bandwidth along the PE direction, combined with inhomogeneities in the main magnetic field, B_0. These field inhomogeneities stem from variations in magnetic susceptibilities among different tissues, as well as the presence of air or metal [40].

Eddy-current–induced distortions arise from the rapid switching of gradients, which generates eddy currents within conductive components of the scanner [41]. Diffusion-weighted sequences experience particularly pronounced eddy currents due to their strong diffusion-sensitizing gradients. These eddy currents counteract the intended gradients, creating unintended effective gradients. Additionally, eddy currents generated by diffusion-sensitizing gradients may persist over an extended period, influencing the subsequent data readout. This persistence results in further geometric distortions of the image, as well as spin dephasing that reduces the SNR.

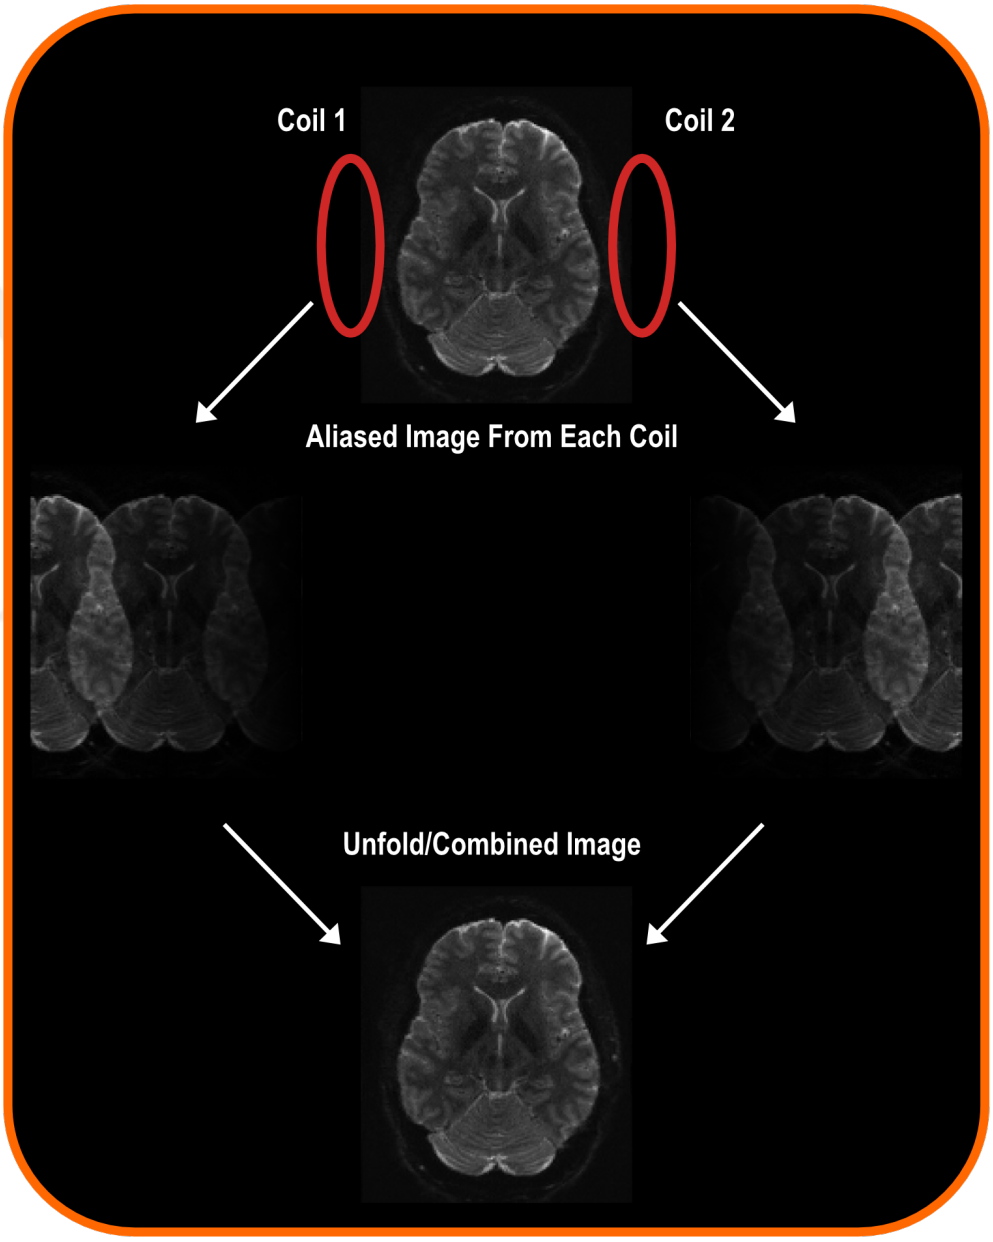

Figure 2.4: SENSE is a parallel imaging technique that operates in image domain, where the k-space is undersampled and the resulting aliased images from each coil are unfolded by explicitly using the coil sensitivity profiles. Images generated using the preprocessed dMRI data of the Human Connectome Project [1].

Postprocessing methods for correcting such artifacts most often estimate the susceptibility-induced off-resonance field directly from pairs of images acquired with opposite PE directions, as these data-driven approaches consistently outperform both registration-based corrections and techniques using separately measured field maps [42]. As shown in Fig. 2.3(b), the estimated field is then used to "unwarp" the blip-up/blip-down image pair, and the degree to which the corrected images coincide serves as an accuracy measure for the estimated field. Well-known implementations of this reversed-PE correction strategy include FSL's TOPUP [43, 44] and SPM's HySCO toolbox [45, 46]. Eddy-current–induced distortions are subsequently addressed with dedicated eddy correction, most commonly performed with FSL's `eddy` tool, which registers each diffusion-weighted volume to a model prediction generated by a Gaussian-process representation of the diffusion signal [47]. By exploiting the assumptions that neighboring gradient directions yield similar signals and that antiparallel directions are identical, the `eddy` tool can disentangle genuine diffusion contrast from eddy-current effects.

Parallel Imaging: Parallel imaging undersamples k-space data by skipping selected PE lines, while exploiting the distinct spatial sensitivity patterns of a multi-channel receive coil array. Two of the most popular parallel imaging techniques in MRI are Sensitivity Encoding (SENSE) and Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA). SENSE reconstructs images in the spatial domain by solving a linear system of equations that uses measured or estimated coil sensitivity maps [48]. Figure 2.4 shows example aliased images from individual coils and the SENSE reconstructed image. GRAPPA works instead in k-space, interpolating missing k-space lines from neighboring acquired lines in k-space, after estimating interpolation coefficients using a fully sampled central calibration region in k-space [49]. In the case of EPI, typical readout reduction factors of 2–4 using parallel imaging roughly halve the echo-train length, reducing geometric distortions and T_2 blurring, although higher factors amplify noise through the so-called noise-amplification or g-factor penalty [50].

Simultaneous Multi-Slice (SMS) Imaging: Simultaneous Multi-Slice (SMS) Imaging, also called multiband (MB) Imaging, excites multiple slices at once with an RF pulse composed of multiple frequency-shifted bands (see Fig. 2.5). During

readout, Controlled Aliasing in Parallel Imaging (CAIPI) introduces small slice-dependent *blips* in the PE direction ("blipped-CAIPI") so that each slice shifts to a different in-plane location before aliasing. The combined use of CAIPI shifts and coil sensitivity differences allows robust slice unmixing with reduced noise. Practically, SMS factors of $N_{\mathrm{MB}} = 2$–6 shorten whole-brain scan time by the same factor; pairing SMS with in-plane GRAPPA yields 2 mm isotropic, 64-direction diffusion-encoded data in under five minutes on a 3 T MRI scanner [51]. Accelerating beyond $N_{\mathrm{MB}} \approx 8$ brings diminishing returns: residual slice-leakage artifacts grow, heating due to RF specific absorption rate (SAR) rises, and performance becomes sensitive to coil geometry and RF-field (B_1) uniformity [50].

2.2 Classical Diffusion Models

The earliest approach to modeling the dMRI signal treats each voxel as a homogeneous, isotropic medium characterized by a single diffusion coefficient. This classical signal model describes the signal attenuation purely as a function of the diffusion-weighting factor b, without accounting for any directional dependence. Building on this idea, a more advanced approach inherits the mono-exponential form but extends it to capture anisotropic diffusion by encoding direction-dependent attenuation through a diffusion tensor formulation. Despite their widespread clinical use and computational efficiency, these classical models are inherently limited in regions of complex microstructure, motivating the development of high-angular-resolution and multi-compartment methods described in the later sections.

2.2.1 Apparent Diffusion Coefficient

Before introducing orientation-dependent models, it is instructive to consider the simplest diffusion-weighted signal model, which assumes a single, isotropic diffusion coefficient per voxel. In this mono-exponential framework, the measured

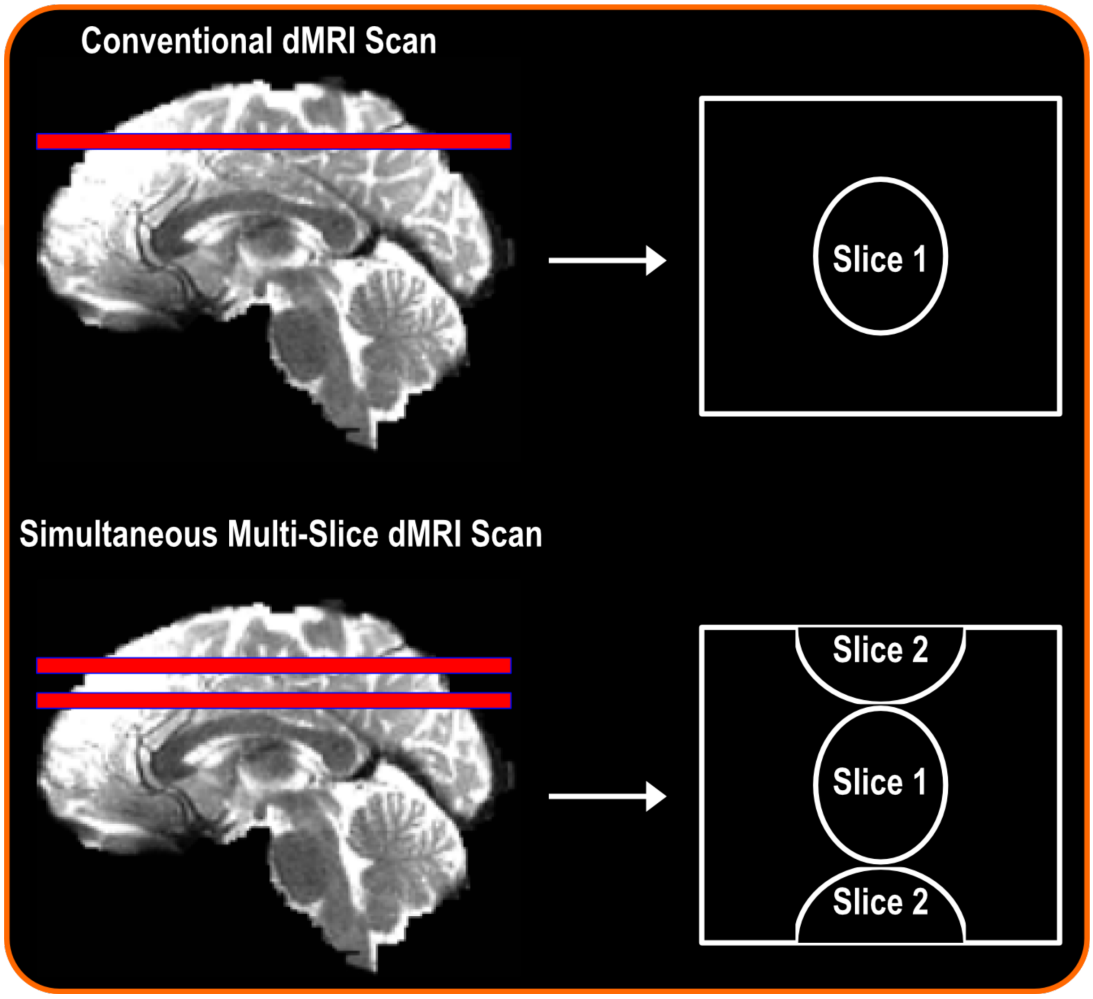

Figure 2.5: Schematic comparison of conventional single-slice and simultaneous multi-slice (SMS) diffusion MRI acquisitions. In the conventional scan (top), each slice is excited and diffusion-encoded (red line) during its own slice-encoding period. In the SMS scan (bottom), two slices are excited and encoded simultaneously using a multiband RF pulse, via controlled aliasing methods that shift images from each slice to different in-plane locations. The acquired signal from both slices is then separated into individual slice images by coil-sensitivity–based SMS reconstruction. Images are from preprocessed dMRI data of the Human Connectome Project [1].

signal attenuation $S(b)$ at diffusion-weighting b is described by

$$S(b) \;=\; S_0\, e^{-b\, D_{\mathrm{ADC}}}, \tag{2.6}$$

where S_0 is the non-diffusion-weighted (i.e., $b = 0$) signal and D_{ADC} is the ADC [34, 35]. By acquiring just two images, one at $b = 0$ and one at a single nonzero b, D_{ADC} can be estimated via linear regression of $\ln[S(b)/S_0]$ versus b using Eq. 2.6. This model provides a rapid, quantitative measure of water mobility that has proven invaluable in clinical settings, for example in acute stroke assessment [18]. However, because it treats diffusion as isotropic and neglects any directional dependence, the mono-exponential ADC is unable to resolve microstructural anisotropy or multiple fiber populations.

2.2.2 Diffusion Tensor Imaging

Diffusion Tensor Imaging (DTI) is one of the most widely used methods for characterizing diffusion anisotropy. Under the assumption of symmetric diffusion, the process is described by a positive-definite, symmetric tensor [9]:

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix}. \tag{2.7}$$

where D_{ij} is the covariance of diffusion displacements along the i-th and j-th axis [52].

Decomposing $\mathbf{D}$ yields eigenvectors $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ with corresponding eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$, which together define an ellipsoid whose axes align with the eigenvectors and whose axis lengths scale with the corresponding eigenvalues. As shown in Fig. (2.6)(a), the eigenvector $\mathbf{e}_1$ associated with the largest eigenvalue λ_1 indicates the principal diffusion direction, denoting the dominant fiber orientation within that voxel. Extending Eq. (2.3), the resulting signal model is [9]:

$$S(\mathbf{g}, b) \;=\; S_0\, e^{-b\, \mathbf{g}^T\, \mathbf{D}\, \mathbf{g}}, \tag{2.8}$$

where $\mathbf{g} = \mathbf{q}/\|\mathbf{q}\|$ is the unit-length gradient direction. Because the tensor $\mathbf{D}$ is symmetric and contains six unknowns, resolving the $\mathbf{D}$ together with S_0 requires six diffusion-weighted acquisitions in non-collinear directions plus a single non-diffusion-weighted (i.e., $b = 0$) acquisition.

A few important metrics derived from the diffusion tensor are widely utilized in clinical research on neurodegenerative diseases. One of the key metrics is fractional anisotropy (FA), which reflects the degree of directional dependence of water diffusion:

$$\mathrm{FA} = \frac{\sqrt{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_3 - \lambda_1)^2}}{\sqrt{2\left(\lambda_1^2 + \lambda_2^2 + \lambda_3^2\right)}}, \quad 0 \leq \mathrm{FA} \leq 1. \tag{2.9}$$

Another essential metric is mean diffusivity (MD), which quantifies the average magnitude of diffusion by averaging the three eigenvalues:

$$\mathrm{MD} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} = \frac{\mathrm{Tr}(\mathbf{D})}{3}. \tag{2.10}$$

Here, $\mathrm{Tr}(\mathbf{D})$ denotes the trace (i.e., the sum of the diagonal entries) of $\mathbf{D}$.

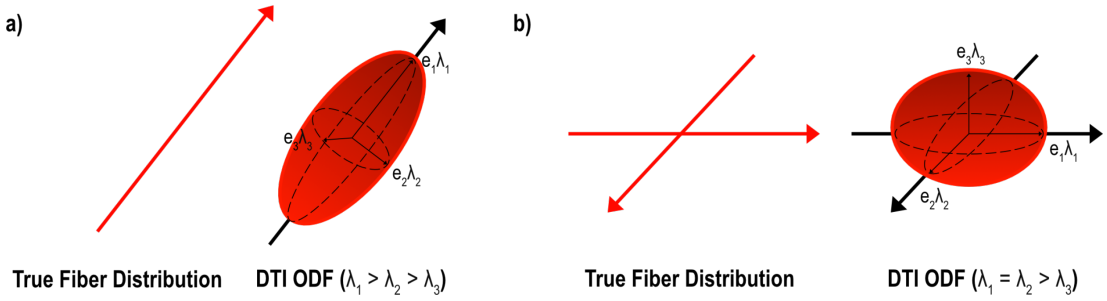

Figure 2.6: (a) In regions with a single dominant fiber population ($\lambda_1 > \lambda_2 > \lambda_3$), DTI represents diffusion as an elongated ellipsoid whose longest axis aligns with the principal fiber direction. (b) When two fiber bundles cross at equal strength ($\lambda_1 = \lambda_2 > \lambda_3$), the DTI analysis yields an oblate spheroid and fails to distinguish the individual pathways.

2.2.3 Limitations of DTI

DTI analysis is limited to a single dominant fiber direction per voxel, implicitly assuming simple, unidirectional fiber structures. For example, in the case of two crossing fibers as shown in Fig. 2.6(b), DTI analysis will yield an oblate diffusion ellipsoid with two equal axes (i.e., $\lambda_1 = \lambda_2$), failing to characterize the presence of a secondary fiber orientation. In practice, many voxels contain multiple fiber populations that intersect, fan, or run in close proximity, especially in complex white-matter regions [53]. Accurately disentangling these overlapping pathways requires more sophisticated diffusion MRI approaches.

2.3 High-Angular and Multi-Shell Diffusion Models

Modern dMRI protocols acquire data with dozens to hundreds of gradient directions across multiple b-values, enabling detailed characterization of tissue microstructure beyond what DTI can offer. High Angular Resolution Diffusion Imaging (HARDI) builds on dMRI by sampling the q-space more completely, using many gradient directions and b-values [9]. From these measurements, the Orientation Distribution Function (ODF) is estimated, which describes the probability of diffusion in each direction.

In these models, first the signal attenuation is described as follows [9]:

$$E(\mathbf{q},t) \;=\; \frac{S(\mathbf{q},t)}{S_0} \;=\; \int_{\mathbb{R}^3} p(\mathbf{r},t)\, e^{-2\pi i\, \mathbf{q}^T \mathbf{r}}\, \mathrm{d}\mathbf{r} \;=\; \mathcal{F}\{p(\mathbf{r},t)\}, \tag{2.11}$$

where $p(\mathbf{r},t)$ gives the probability that a water molecule moves to position $\mathbf{r}$ after time t, and $\mathcal{F}(\cdot)$ denotes Fourier transform. In the simple (isotropic) case, $p(\mathbf{r},t)$ is a Gaussian function [9]. If we measure the signal attenuation $E(\mathbf{q},t)$ for the entire q-space, we can recover $p(\mathbf{r},t)$ by taking its inverse Fourier transform [9]:

$$p(\mathbf{r},t) \;=\; \mathcal{F}^{-1}\{E(\mathbf{q},t)\}. \tag{2.12}$$

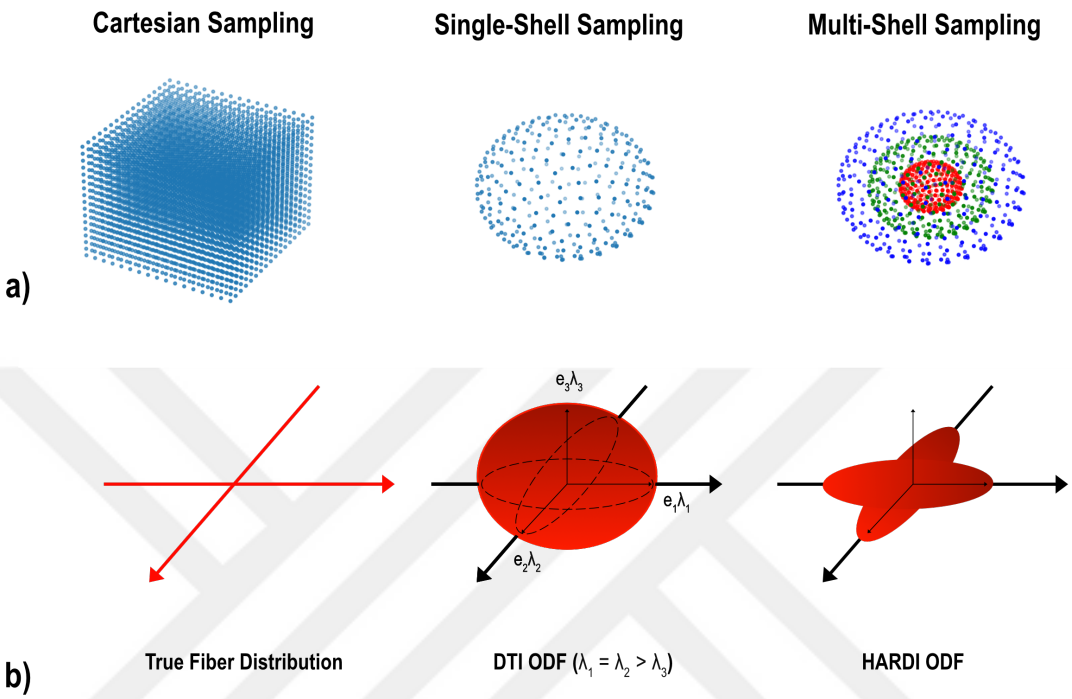

Figure 2.7: (a) HARDI acquisition schemes differ in how they sample q-space. DSI densely covers q-space using a Cartesian grid, usually with more than 200 samples, providing broad coverage. In contrast, single-shell and multi-shell HARDI acquire data on spherical surfaces whose radii are set by the b-value, achieving high angular resolution with far fewer samples than DSI. (b) The richer angular information in HARDI supports more advanced diffusion signal models. Unlike DTI, HARDI can recover ODF coefficient sets whose peaks correspond to the actual fiber orientations in the case of crossing fibers.

Since fully sampling the q-space is not feasible, we approximate $p(\mathbf{r}, t)$ by its ensemble average, called the Ensemble Average Propagator (EAP). The ODF is then defined from the EAP and captures all directional information [2]:

$$\Psi(\theta, \phi) \;=\; \int_0^\infty p(r, \theta, \phi)\; r^2 \, \mathrm{d}r, \tag{2.13}$$

where $p(r, \theta, \phi)$ is $p(\mathbf{r}, t)$ in spherical coordinates, with $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$. The peaks of $\Psi(\theta, \phi)$ line up with the main fiber directions, which helps in further analysis and visualization [9].

There are two common approaches to acquire HARDI data as shown in Fig. 2.7(a):

- **Diffusion Spectrum Imaging (DSI):** The q-space is sampled on a dense, Cartesian grid, often requiring more than 200 samples [9, 54].
- **Single-Shell HARDI:** The q-space is sampled on a single spherical shell at a fixed b-value [9].

Throughout this thesis, single-shell HARDI acquisition is considered rather than DSI, as single-shell HARDI requires far fewer gradient directions (around 60 compared to 200), making data acquisition faster and more efficient.

Multi-shell diffusion MRI generalizes single-shell HARDI by acquiring diffusion-weighted signals on several concentric spheres in q-space, each defined by a different b-value [2]. This strategy combines the high angular resolution provided by large b-values (e.g., $b \geq 2000$ s/mm^2) with the improved SNR of lower b-values (e.g., $b \leq 1000$ s/mm^2). By densely sampling multiple shells, multi-shell protocols enhance the angular fidelity of orientation distribution estimates, while preserving adequate SNR for robust fitting as shown in Fig 2.7(b). Multi-shell dMRI enables:

- **Improved ODF estimation:** The additional radial information provided by multiple b-values helps disambiguate fiber populations and reduces potential bias in the estimated ODF.
- **Model-based compartment analysis:** Fitting biophysical models across multiple b-values provides estimates of neurite density, orientation dispersion, and free-water fraction.

By sampling the full radial–angular space, multi-shell imaging strikes a balance between angular detail and quantitative specificity, forming the foundation for advanced unsupervised and physics-guided reconstruction methods proposed in this thesis.

2.3.1 Spherical Harmonic Decomposition

On each shell (i.e., fixed b_i), the diffusion signal I_{ik} in voxel k along the i^{th} gradient direction (denoted in spherical coordinates as (θ_i, ϕ_i)) admits an expansion in the spherical-harmonic (SH) basis as follows [9]:

$$I_{ik} \;=\; \sum_{l=0}^{L} \sum_{m=-l}^{l} a_{lm,k}\, Y_l^m(\theta_i, \phi_i), \tag{2.14}$$

where Y_l^m are the real-symmetric spherical harmonics up to order L, and $a_{lm,k}$ are the corresponding coefficients for voxel k. Because I_{ik} is real and antipodally symmetric, only even l is typically employed and each complex basis function into its real and imaginary parts are converted as follows:

$$Y_l^m(\theta, \phi) \;\rightarrow\; \begin{cases} \sqrt{2}\, Re(Y_l^m(\theta, \phi)), & m > 0, \\ Y_l^0(\theta, \phi), & m = 0, \\ \sqrt{2}\, Im(Y_l^{|m|}(\theta, \phi)), & m < 0. \end{cases} \tag{2.15}$$

This decomposition yields a smooth, continuous representation of the angular signal in each voxel using the spherical harmonics illustrated in Fig. 2.8.

2.3.2 Fiber Orientation Distribution Estimation

In single-shell HARDI, the measured signal $I(\theta, \phi)$ at each orientation (θ, ϕ) can be modeled as a spherical convolution between the true fiber orientation distribution function (fODF), $F(\theta, \phi)$, and a common single-fiber response function $R(\theta)$ [55]:

$$I(\theta, \phi) = F(\theta, \phi) \otimes R(\theta). \tag{2.16}$$

Fiber Response and fODF Representation: Under the assumption that all coherent white-matter fiber bundles share the same microscopic diffusion behavior, $R(\theta)$ can be measured from voxels with highly aligned fibers (e.g., FA>0.7).

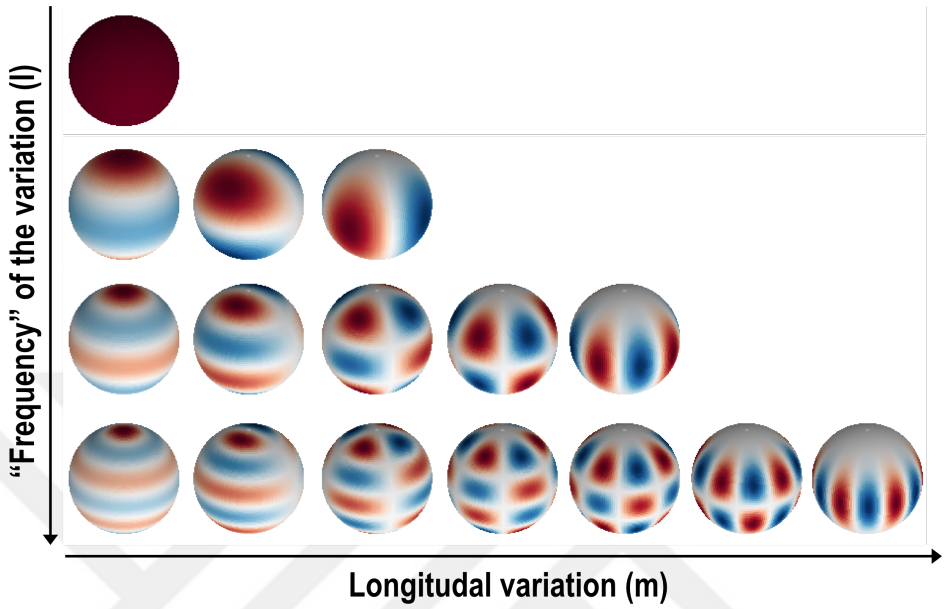

Figure 2.8: An illustration of spherical harmonics Y_l^m for even degrees $l = 0, \ldots, 6$ and orders $m = 0, \ldots, l$, rendered on unit spheres with an orthographic projection at a 30° elevation and 30° azimuth. The diverging "seismic" colormap highlights positive lobes in red and negative lobes in blue.

The fODF itself is modeled as a weighted sum of Dirac functions [55]:

$$F(\theta, \phi) = \sum_{i=1}^{N} v_i \, \delta\big(\theta - \theta_i, \phi - \phi_i\big), \tag{2.17}$$

where each volume fraction v_i indicates the proportion of fibers oriented along (θ_i, ϕ_i).

Spherical Harmonics Formulation: To solve the convolution efficiently, we expand I, F, and R in a SH basis up to order $l_{\max}$. Let $\mathbf{a}_l$, and $\mathbf{f}_l$ be the vectors of SH coefficients by the l^{th} order for $I(\theta, \phi)$ and $F(\theta, \phi)$ respectively. The convolution reduces to a matrix multiplication as follows [56]:

$$\mathbf{a}_l \;=\; \mathbf{R}_l \, \mathbf{f}_l, \tag{2.18}$$

where $\mathbf{R}_l$ is a matrix whose entries are given by the l^{th} order SH coefficients of $R(\theta)$. Then, given $\mathbf{a}_l$ and $\mathbf{R}_l$, one can find $\mathbf{f}_l$ using linear least squares estimation.

Constrained Spherical Deconvolution (CSD): Direct inversion for $\mathbf{f}_l$ is unstable, especially at higher l (i.e., the components of the signal capturing higher

angular variations are more easily corrupted) [56]. Constrained spherical deconvolution (CSD) imposes a non-negativity constraint on the reconstructed fODF coefficients:

$$\hat{\mathbf{f}} = \arg\min_{\mathbf{f}} \ \frac{1}{2}\left\|\mathbf{C}\,\mathbf{f} - \mathbf{I}\right\|^2 \quad \text{subject to} \quad \mathbf{L}\,\mathbf{f} \ \geq \ 0, \tag{2.19}$$

where $\mathbf{I}$ is the measured single-shell HARDI signal intensities for all q-space directions, $\mathbf{f}$ is the unknown fODF to be estimated, $\mathbf{C}$ is the spherical convolution kernel, and $\mathbf{L}$ is the constraint matrix enforcing non-negativity. This constrained formulation yields physically meaningful, non-negative fODFs [56].

Multi-Shell Multi-Tissue CSD (MSMT-CSD): To account for multiple tissue compartments (e.g., white matter, gray matter, and CSF) and take advantage of multi-shell data, MSMT-CSD extends Eq. (2.19) to jointly deconvolve signals from m shells and n tissue types [10]:

$$\begin{aligned}\{\hat{\mathbf{f}}_j\}_{j=1}^{n} = \arg\min_{\{\mathbf{f}_j\}} \ \tfrac{1}{2}\Bigg\| \underbrace{\begin{bmatrix} \mathbf{C}_{1,1} & \cdots & \mathbf{C}_{1,n} \\ \vdots & & \vdots \\ \mathbf{C}_{m,1} & \cdots & \mathbf{C}_{m,n} \end{bmatrix}}_{\mathbf{C}_{\text{all}}} \begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_n \end{bmatrix} - \begin{bmatrix} \mathbf{I}_1 \\ \vdots \\ \mathbf{I}_m \end{bmatrix} \Bigg\|^2, \\ \text{subject to} \quad \underbrace{\begin{bmatrix} \mathbf{L}_1 & & \\ & \ddots & \\ & & \mathbf{L}_n \end{bmatrix}}_{\mathbf{L}_{\text{all}}} \begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_n \end{bmatrix} \ \geq \ 0,\end{aligned} \tag{2.20}$$

where $\mathbf{I}_i$ is the vector of the HARDI signal intensities on the i^{th} shell, $\mathbf{C}_{i,j}$ is the convolution kernel for the j^{th} tissue on the i^{th} shell, $\mathbf{f}_j$ are the fODF coefficients for the j^{th} tissue, and $\mathbf{L}_j$ is the constraint matrix enforcing non-negativity for the j^{th} tissue.

White matter is modeled with full SH order (i.e., anisotropic response), while gray matter and CSF are often restricted to order zero (i.e., isotropic response). This joint estimation improves the separation of tissue-specific signals and yields simultaneous maps of fiber orientations and tissue volume fractions [10].

Metrics Derived from CSD/MSMT-CSD: The following metrics are commonly derived for tissue microstructure analysis using CSD/MSMT-CSD [57]:

- **Apparent Fiber Density (AFD):** The magnitude of fODF, which denotes the fiber density per voxel. This metic can also be used to evaluate the performance of the reconstruction.
- **Number of Fiber Orientations (nuFO):** The count of distinct peaks in the fODF above a certain threshold, reflecting the local complexity of fiber crossings.

2.3.3 Neurite Orientation Dispersion Density Imaging

NODDI is a multi-shell HARDI method that models the normalized diffusion signal A as a sum of three compartments, restricted (intra-neurite), hindered (extra-neurite), and isotropic (CSF), each weighted by its volume fraction. The combined signal is expressed as [11]:

$$A \;=\; \big(1-\nu_{\text{iso}}\big)\Big[\nu_{\text{ic}}\,A_{\text{ic}} + \big(1-\nu_{\text{ic}}\big)\,A_{\text{ec}}\Big] \;+\; \nu_{\text{iso}}\,A_{\text{iso}}, \tag{2.21}$$

where A_{ic}, A_{ec}, and A_{iso} are the normalized signals from the intra-neurite, extra-neurite, and CSF compartments, respectively. In addition, ν_{ic} is the intra-neurite volume fraction and ν_{iso} the CSF (isotropic) volume fraction.

Under the dispersed-cylinder model, the intra-neurite signal is given by

$$A_{\text{ic}} \;=\; \int_{\mathbb{S}^2} f(\mathbf{n}) \exp\big[-\,b\,d_{\parallel}\,(\mathbf{q}\cdot\mathbf{n})^2\big]\,\mathrm{d}\mathbf{n}, \tag{2.22}$$

where $d_{\parallel}$ is the intrinsic diffusivity along each neurite direction $\mathbf{n}$, and $f(\mathbf{n})$ is the orientation distribution (modeled as a Watson distribution).

The extra-neurite compartment follows a Gaussian model given as

$$\log A_{\text{ec}} \;=\; -\,b\,\mathbf{q}^T\Big(\int_{\mathbb{S}^2} f(\mathbf{n})\,D(\mathbf{n})\,\mathrm{d}\mathbf{n}\Big)\mathbf{q}, \tag{2.23}$$

with $D(\mathbf{n})$ a cylindrically symmetric diffusion tensor whose principal axis is $\mathbf{n}$. In addition, the perpendicular diffusivity $d_{\perp}$ is linked to $d_{\parallel}$ by

$$d_{\perp} \;=\; d_{\parallel}\,(1-\nu_{\text{ic}}). \tag{2.24}$$

Finally, the CSF compartment is modeled as a simple isotropic diffusion with its own Gaussian profile.

The primary aim in NODDI analysis is to recover the neurite density, ν_{ic}, and the orientation dispersion encoded in $f(\mathbf{n})$, rather than the full high-order ODF [11]. The primary metric of the NODDI analysis is the Orientation Dispersion Index (ODI), defined as

$$\mathrm{ODI} \;=\; \frac{2}{\pi}\,\arctan\!\big(\kappa^{-1}\big), \tag{2.25}$$

where κ is the concentration parameter of the Watson distribution in the intra-neurite ("stick") model. Lower κ denotes more dispersed orientations and yields higher ODI, indicating greater angular spread around the mean direction. Another important NODDI metric is the Intracellular Volume Fraction (ICVF), which represents the overall neurite density and is equal to the estimated ν_{ic} .

NODDI fits the multi-compartment signal model in Eq. (2.21) to the HARDI data using a Gauss–Newton scheme to compute maximum-likelihood estimates of ν_{ic}, ν_{iso}, and κ. A set of diffusivities is held at fixed a priori values, and at least two distinct b-value shells are required for robust fitting [11].

2.4 Diffusion MRI Metrics

Diffusion MRI enables high-resolution, high-SNR mapping of neurite microstructure, allowing us to extract intra-cellular volume fraction, diffusion anisotropy, fiber orientations, and dispersion simultaneously. These metrics provide critical insights into tissue integrity and are widely used in the diagnosis of neurodegenerative diseases and focal brain disorders. As such, the scalar maps derived from dMRI data form the principal outputs of the methods described in previous sections.

The metrics presented above fall into two categories: anisotropy indices and fiber-density indices. In this thesis, we consider six different diffusion measures:

FA and MD from DTI, AFD and nuFO from MSMT-CSD, and ODI and ICVF from NODDI.

As introduced in Section 2.2.2, FA reflects diffusion anisotropy while MD represents mean diffusivity. Although these DTI-based measures capture important aspects of white-matter integrity, HARDI-derived metrics offer greater specificity and sensitivity. For instance, the NODDI metrics ICVF and ODI (see Section 2.3.3) separate the two principal microstructural influences that FA conflates: neurite density and orientation dispersion. This distinction enables a more detailed analysis of anisotropy within fiber populations [11]. Similarly, as described in Section 2.3.2, MSMT-CSD analysis yields AFD, which is directly proportional to the intra-neurite volume fraction and provides an accurate estimate of fiber density per unit volume, especially in voxels with crossing fibers [10]. In addition, the nuFO metric quantifies local white-matter complexity. Together, these NODDI and MSMT-CSD metrics allow us to isolate and investigate the individual microstructural contributions within the diffusion signal.

Chapter 3

QUCCI for Unsupervised q-Space Upsampling

This chapter is based in parts on the publication titled:

> **A. Topcu**, A. Z. Alkilani, T. Çukur and E. U. Saritas, "Unsupervised q-Space Interpolation Using Physics-Constrained Coordinate-Based Implicit Network" *Proceedings of the 32nd Annual Meeting of ISMRM*, Singapore, May 2024.

3.1 Introduction

Unsupervised q-Space Upsampling Using Physics-Constrained Coordinate-Based Implicit Networks (QUCCI) presents an unsupervised framework that uses implicit neural representations (INRs) to reconstruct dMRI volumes by outputting spherical polar Fourier (SPF) basis coefficients, thereby treating the signal as a continuous function of spatial coordinates. This design affords uninterrupted interpolation across both angular and radial directions in q-space, permitting flexible sampling of q-space without adherence to fixed acquisition schemes. The core

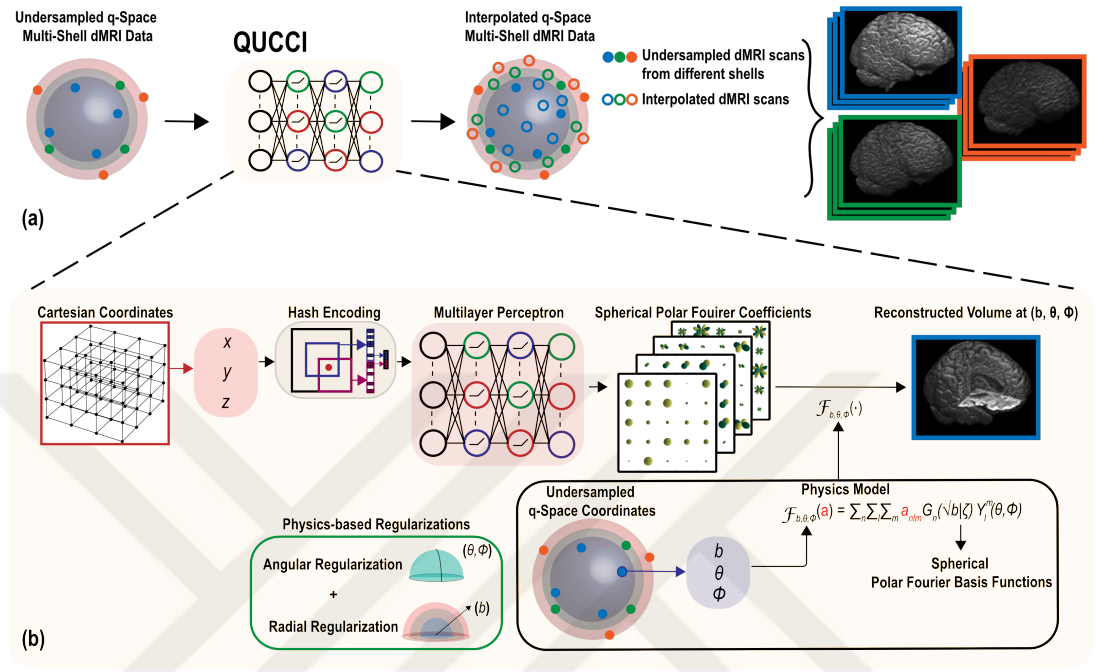

Figure 3.1: (a) QUCCI is capable of smoothly interpolating throughout q-space in both angular and radial dimensions, while embedding the complete q-space signal representation within its parameters. (b) Schematic of the QUCCI architecture. During training, samples were drawn from both Cartesian coordinates (x, y, z) and q-space parameters (b, θ, ϕ). For the Cartesian points, random sampling was applied to mimic artifact-induced variations. These sampled Cartesian coordinates are then encoded via a multi-resolution hash grid structure that supports look-up and interpolation. The resulting coordinate embeddings are input to an implicit neural network, which predicts the SPF coefficients. Finally, these coefficients are transformed into the corresponding pixel intensity at the specified q-space location using a physics-based forward model, and compared against the ground-truth pixel value.

innovation of QUCCI is its subject-specific, task-agnostic nature, which achieves state-of-the-art reconstruction accuracy at acceleration factors of $R = 10$–22.5 for multi-shell dMRI, all the while requiring minimal storage once training is complete. This performance stems from the integration of multi-resolution hash grid encoding, modeling outputs directly as SPF coefficients, and a combined regularization strategy that fuses MLP-based implicit priors with physics-driven explicit constraints. Collectively, these features enable QUCCI to address the demands of high-angular and high-radial resolution dMRI reconstruction.

3.2 Theory

3.2.1 Problem Formulation

3.2.1.1 Formulating q-Space Interpolation as an Inverse Problem

An inverse formulation for reconstructing diffusion MRI data $\mathbf{I}$ in q-space is given by

$$\mathbf{I} = \mathbf{M}\,\mathbf{a} + \boldsymbol{\epsilon}, \tag{3.1}$$

where $\mathbf{M}$ denotes the measurement operator encoding the imaging physics, $\mathbf{a}$ contains the acquired q-space samples, and $\boldsymbol{\epsilon}$ accounts for model and measurement errors. Depending on the method, $\mathbf{M}$ may act in image, k-space, or other transform domains, and $\mathbf{a}$ accordingly represents measurements in those domains. Because this problem is ill-posed, a regularized reconstruction can be formulated as follows:

$$\min_{\mathbf{a}} \ \tfrac{1}{2}\left\|\mathbf{I} - \mathbf{M}\,\mathbf{a}\right\|_2^2 \ + \ \lambda\, R(\mathbf{a}), \tag{3.2}$$

where the data-fidelity term enforces consistency with the measurements and $R(\mathbf{a})$ injects prior knowledge (e.g., sparsity or smoothness) into the solution.

3.2.1.2 Classical INR–Based q-Space Interpolation

In a standard INR framework, the network directly predicts the voxel intensity for an unseen q-space coordinate, so that $\mathbf{Ma}$ is realized as the INR output itself. This direct regression requires ground-truth dMRI intensities for every new q-space sample, making it unsuitable for unsupervised training. To enable self-supervision, explicit forward operator $\mathbf{M}$ can be used. For example, the signal can be expanded in a spherical-harmonic (SH) basis [58]. The resulting reconstruction problem then becomes

$$\min_{\{\mathbf{a}_k\}} \ \frac{1}{B}\sum_{k=1}^{B}\left[\mathrm{L1}_{\mathrm{smooth}}\left(\mathbf{I}_k - \mathbf{M}\,\mathbf{a}_k\right) \ + \ \lambda\left\|\mathbf{a}_k\right\|_1\right], \tag{3.3}$$

where B is the batch size, $\mathbf{I}_k$ is the vector of measured dMRI signals at the observed q-space indices for voxel k, and $\mathbf{a}_k$ contains its corresponding SH coefficients.

3.2.2 QUCCI Formulation

3.2.2.1 Signal Characterization

In QUCCI, the dMRI signal is modeled using an orthonormal SPF basis, assuming that sufficiently large radial and angular truncation orders N and L allow any I_{ik} to be accurately represented [59]. The full signal can be denoted by $\mathbf{I} \in \mathbb{R}^{N_q \times N_x}$, where N_q is the number of q-space samples and N_x is the voxel count. Then, I_{ik} corresponds to the signal at the k-th voxel for the i-th q-space point, and can be expressed in terms of the SPF basis functions $B^{SPF}_{nlm}(b_i, \theta_i, \phi_i \mid \zeta)$ as follows [59]:

$$\begin{aligned} I_{ik} &= \mathcal{F}_{b_i,\theta_i,\phi_i}(\boldsymbol{a_k}) \\ &= \sum_{n=0}^{N} \sum_{l=0}^{L} \sum_{m=-l}^{l} a_{nlm} B^{SPF}_{nlm}(b_i, \theta_i, \phi_i | \zeta) \end{aligned} \tag{3.4}$$

where

$$B^{SPF}_{nlm}(b_i, \theta_i, \phi_i | \zeta) = G_n(\sqrt{b_i} | \zeta)\, Y_l^m(\theta_i, \phi_i). \tag{3.5}$$

Here, $\boldsymbol{a_k}$ is the vector of SPF coefficients for voxel k, and (b_i, θ_i, ϕ_i) are the spherical coordinates of the i-th q-space sample. The SPF basis B^{SPF}_{nlm} itself combines two components:

1. A Gaussian–Laguerre radial function $G_n(\sqrt{b_i} \mid \zeta)$ of order n. The parameter ζ controls the radial scale of the basis [59].
2. A real-valued, symmetric spherical-harmonic $Y_l^m(\theta_i, \phi_i)$ of even order l and degree m.

In this formulation, the angular variation of I_{ik} is captured by spherical-harmonic functions $Y_l^m(\theta_i, \phi_i)$ [2]:

$$Y_l^m(\theta_i, \phi_i) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-|m|)!}{(l+|m|)!}} P_l^m(\cos(\theta_i)) e^{im\phi_i}, \tag{3.6}$$

where P_l^m denotes the associated Legendre polynomial. The set $\{Y_l^m\}$ forms an orthonormal basis on the unit sphere, corresponding to the angular solutions of Laplace's equation in spherical coordinates [60]. Because the measured signal I_{ik} is real-valued and symmetric, QUCCI uses the real-valued SH basis defined as follows [61]:

$$Y_l^m = \begin{cases} \sqrt{2}\,\mathrm{Re}\big(Y_l^m\big), & 0 < m \le l, \\ Y_l^0, & m = 0, \\ \sqrt{2}\,\mathrm{Im}\big(Y_l^{|m|}\big), & -l \le m < 0. \end{cases} \tag{3.7}$$

Given the typically limited number of q-space shells, the radial profile of I_{ik} tends to be sparse. Thus, the radial component is reconstructed using Gaussian–Laguerre functions [59]:

$$G_n(\sqrt{b_i} \mid \zeta) = k_n(\zeta)\, \exp\!\Big(-\tfrac{b_i}{2\zeta}\Big)\, L_n^{\frac{1}{2}}\big(\tfrac{b_i}{\zeta}\big), \tag{3.8}$$

with the normalization constant

$$k_n(\zeta) = \Big[\tfrac{2}{\zeta^{3/2}}\, \tfrac{n!}{\Gamma(n+\frac{3}{2})}\Big]^{1/2}. \tag{3.9}$$

Here, $L_n^{1/2}(\cdot)$ is the generalized Laguerre polynomial and $\Gamma(z) = \int_0^\infty t^{z-1}e^{-t}\,dt$ is the gamma function.

3.2.2.2 Unsupervised Reconstruction

To enable unsupervised learning in QUCCI, this thesis adapts Eq. 3.1 by introducing the SPF-based measurement operator $\mathbf{M}_{\mathrm{SPF}} \in \mathbb{R}^{N_q \times \frac{(N+1)(L+1)(L+2)}{2}}$, which encodes the dMRI signal via SPF basis functions. The matrix $\boldsymbol{\epsilon} \in \mathbb{R}^{N_q \times N_x}$ denotes the residual error between $\mathbf{I}$ and $\mathbf{M}_{\mathrm{SPF}}\mathbf{a}$ caused by truncating the radial and angular SPF orders. The matrix $\boldsymbol{a} \in \mathbb{R}^{\frac{(N+1)(L+1)(L+2)}{2} \times N_x}$ is the network's output matrix of SPF coefficients.

3.2.2.3 Hash Grid Encoding

Given a volumetric signal, a traditional INR seeks to learn a continuous mapping from spatial coordinates to signal values using a neural network, typically an MLP. Formally, the INR defines

$$f(\mathbf{x}) = \mathrm{MLP}\big(\pi(\mathbf{x})\big), \tag{3.10}$$

where $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$ are the voxel coordinates and $\pi(\cdot)$ is a feature-mapping that lifts $\mathbf{x}$ into a higher-dimensional embedding. Early INRs used multi-scale sine and cosine encodings for $\pi(\cdot)$ to counteract the MLP's tendency to favor low-frequency functions [62]. However, this approach places heavy demands on the MLP's parameters and can slow convergence.

To alleviate these drawbacks, QUCCI adopts the hash encoding strategy presented in [63]. Instead of a fixed Fourier embedding, trainable parameters are arranged across multiple 3D grids. For each level $l \in \{1, \ldots, L_{level}\}$, one can define

$$\Pi_l \in \mathbb{R}^{N_l \times N_l \times N_l \times F}. \tag{3.11}$$

Here, N_l interpolates geometrically from $N_{\min}$ to $N_{\max}$ via

$$N_l = N_{\min} \cdot v_{\text{base}}^{l-1}, \tag{3.12}$$

where

$$v_{\text{base}} = \exp\!\big(\tfrac{\ln N_{\max} - \ln N_{\min}}{L_{level} - 1}\big). \tag{3.13}$$

Each coordinate $\mathbf{x}$ is encoded by trilinear interpolation across the eight nearest vertices of Π_l, yielding a feature vector at that level. To avoid the cubic memory growth of dense grids, each Π_l is implemented as a compact hash table $\Pi_l^{\text{hash}} \in \mathbb{R}^{T \times F}$. A spatial hash function

$$h(\mathbf{x}) = \big(\bigoplus_{i=1}^{d} \pi_i \, x_{\pi(i)}\big) \bmod T \tag{3.14}$$

maps grid indices (i, j, k) to table entries, and

$$\Pi_l(i, j, k) = \Pi_l^{\text{hash}}\big[h(i, j, k)\big]. \tag{3.15}$$

The full encoding $\pi(\mathbf{x})$ is the concatenation of all level-wise embeddings:

$$\pi(\mathbf{x}) = [\pi_1(\mathbf{x}), \ldots, \pi_{L_{level}}(\mathbf{x})]. \tag{3.16}$$

This hierarchical hash-based representation is both memory-efficient and end-to-end trainable, accelerating convergence and capturing detail across scales.

3.2.2.4 Implicit Network

The overall architecture of QUCCI is illustrated in Fig. 3.1(a). Starting from a coordinate $\mathbf{x}$, QUCCI first samples from an isotropic Gaussian $\mathcal{N}(\mathbf{x}, \Sigma)$ to introduce spatial smoothing and model residual noise $\boldsymbol{\epsilon}$. These perturbed points are then passed through the multi-level hash encoder $\pi(\cdot)$ and fed into the MLP to predict SPF coefficients:

$$\mathrm{MLP}\big(\pi\big(\mathcal{N}(\mathbf{x}, \Sigma)\big)\big) = \mathbf{a} + \boldsymbol{\epsilon}. \tag{3.17}$$

Finally, QUCCI converts the MLP output back to a voxel intensity at $\mathbf{x}$ using the forward SPF model:

$$\hat{I}_{ik} = \mathcal{F}_{b_i,\theta_i,\phi_i}\big(\mathrm{MLP}(\pi(\mathcal{N}(\mathbf{x}, \Sigma)))\big), \tag{3.18}$$

which effectively applies the operator $\mathbf{M}_{\mathrm{SPF}}$ at the specific q-space location (b_i, θ_i, ϕ_i).

3.2.2.5 Loss Functions

During training of QUCCI, the network parameters are optimized by minimizing

$$\arg\min_{\Theta} \mathcal{L}_{QUCCI}(\Theta), \tag{3.19}$$

where

$$\mathcal{L}_{QUCCI}(\Theta) = \mathcal{L}_{MSE}(\Theta) \; + \; \lambda\, \mathcal{R}_{Reg}(\Theta). \tag{3.20}$$

Here, $\mathcal{L}_{MSE}$ ensures data fidelity via the mean squared error, $\mathcal{R}_{Reg}$ encodes the chosen prior, λ is the regularization weight, and Θ comprises all trainable weights of the MLP and the hash grids.

Since the true SPF coefficients are unavailable, first the predicted coefficients are converted into voxel intensities $\hat{I}_{ik}$ using Eq. 3.18. Then, the MSE is computed against the measured signals I_{ik}. For the i-th acquired q-space sample within a batch of size N_{batch}, the data-consistency loss is defined as

$$\mathcal{L}_{MSE} = \frac{1}{N_{\text{batch}}} \sum_{k=1}^{N_{\text{batch}}} \left(I_{ik} - \hat{I}_{ik}\right)^2. \tag{3.21}$$

This process is repeated sequentially over each acquired q-space measurement, enforcing voxel-wise agreement between predictions and observations.

Because q-space interpolation is inherently underdetermined, an explicit regularizer $\mathcal{R}_{Reg}$ is incorporated to enhance image quality and suppress noise. Beyond the implicit regularization offered by the network architecture [64, 65], QUCCI features a physics-informed smoothness prior on the sphere. Specifically, it employs the Laplace–Beltrami operator as a unit-sphere smoothness term [61], which for a single shell (radial order $N = 0$) simplifies to the divergence-of-gradient metric on the spherical harmonic representation.

$$\begin{aligned} E(f) &= \int_{\Omega} (\Delta_b f)^2 d\Omega \\ &= \int_{\Omega} (\Delta_b \sum_{n=0,l,m} a_{nlm} Y_l^m(\theta,\phi))(\Delta_b \sum_{n=0,l',m'} a_{nl'm'} Y_{l'}^{m'}(\theta,\phi)) d\Omega \\ &= \sum_{i=1}^{(L+1)(L+2)/2} a_{0l_im_i}^2 l_i^2 (l_i+1)^2 = \mathbf{a^T L a} \end{aligned} \tag{3.22}$$

Here, Δ_b denotes the Laplace–Beltrami operator, which extends the classical Laplacian to functions defined on a surface and quantifies their smoothness over the unit sphere.

Because any spherical function f expanded in SH basis satisfies $\Delta_b Y_l^m = -l(l+1)Y_l^m$, the operator simplifies to the following [2, 66]

$$\Delta_b f = \frac{\partial^2 f}{\partial \theta^2} + \frac{1}{\tan\theta} \frac{\partial f}{\partial \theta} + \frac{1}{\sin^2\theta} \frac{\partial^2 f}{\partial \phi^2} = -l(l+1)\, f, \tag{3.23}$$

where $f(\theta, \phi) = \sum_{n,l,m} a_{nlm}\, Y_l^m(\theta, \phi)$. Next, the SH orders are collected into a diagonal matrix $\mathbf{L} \in \mathbb{Z}^{\frac{(L+1)(L+2)}{2} \times \frac{(L+1)(L+2)}{2}}$, whose i-th diagonal entry is $l_i^2(l_i+1)^2$. This yields a closed-form regularizer in matrix notation.

To additionally regularize along the radial dimension and encourage mono-exponential signal decay across shells [59], each diagonal term was augmented by $n_i^2(n_i+1)^2$, resulting in

$$\Lambda_{nlm} = \lambda_l\, l_i^2(l_i+1)^2 \;+\; \lambda_n\, n_i^2(n_i+1)^2, \tag{3.24}$$

where $\mathbf{\Lambda} \in \mathbb{Z}^{\frac{(N+1)(L+1)(L+2)}{2} \times \frac{(N+1)(L+1)(L+2)}{2}}$. In block form:

$$\Lambda = \begin{bmatrix} \Lambda_{000} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \Lambda_{NLL} \end{bmatrix}. \tag{3.25}$$

Finally, the Frobenius norm was applied to $\mathbf{a}^\top \Lambda\, \mathbf{a}$ to form the regularization penalty, $\mathcal{R}_{Reg}$.

3.3 Methods

3.3.1 Dataset

Demonstrations were conducted using the preprocessed dMRI data from the Human Connectome Project (HCP) 1200 Subjects Data Release [1]. From this collection, the preprocessed dMRI magnitude images of ten subjects were chosen. Data were acquired on a 3 T Siemens Skyra "Connectom" scanner employing a multi-band diffusion sequence with single-shot echo-planar imaging (ss-EPI) readouts and reversed phase-encoding (PE) directions (RL and LR). Imaging parameters included a 210×180 mm^2 FOV, 168×144 acquisition matrix, 1.25 mm isotropic resolution, multiband acceleration factor of 3, 6/8 partial-Fourier acquisition, echo spacing of 0.78 ms, bandwidth of 1488 Hz/Px, TR/TE of 5520/89.50 ms, and a flip angle of 78°. A total of 18 images at $b = 0$ s/mm^2 were collected. Three diffusion shells with b-values of 1000, 2000, and 3000 s/mm^2 were sampled at 90

q-space points each, yielding 270 q-space samples [67]. The dMRI preprocessing pipeline comprised intensity normalization across runs, EPI distortion correction via the TOPUP algorithm, eddy current and motion correction with the EDDY algorithm, gradient nonlinearity correction, computation of gradient deviation, and the registration of diffusion data and gradient orientations to the 1.25 mm structural space. Finally, brain masks generated through FreeSurfer segmentation were applied to the scans [67, 68, 47, 69].

3.3.2 q-Space Undersampling

QUCCI was evaluated across multiple acceleration factors to assess its effectiveness on reconstructed dMRI images and subsequent analyses. The acceleration factor (R) is defined as the ratio between the total number of q-space samples in the fully sampled dataset (i.e., 270 in this case) and that in the undersampled dataset. For each R, retrospective q-space undersampling was performed by first generating equidistributed points on the surface of each diffusion shell [70, 71]. Then, from the original q-space points, the ones nearest to the equidistributed points were selected.

Four distinct undersampling protocols were employed: $R = 22.5$ (2-shell), 15 (3-shell), 15 (2-shell), and 10 (3-shell). The term "3-shell" indicates that the undersampled q-space points were uniformly allocated across all three diffusion shells. In contrast, for the "2-shell" protocols, the middle q-space shell was omitted, and the undersampled points were evenly split between the remaining two shells (i.e., the inner and outer shells). These 2-shell scenarios specifically probe QUCCI's ability to perform radial interpolation.

3.3.3 QUCCI Implementation

The MLP comprises six fully connected layers, each with 256 neurons. Skip connections are inserted at the second and fourth layers, concatenating the hash-encoding output embeddings with the intermediate activations. Hash-encoding hyperparameters were set to $N_1 = 16$, $L_{level} = 16$, and $v_{base} = 1.39$. Parameters in the hash grid were initialized from a uniform distribution $\mathcal{U}(-10^4, 10^4)$.

For estimating the SPF coefficients of the dMRI signal, radial and angular truncation orders were fixed at $N = 2$ and $L = 7$, respectively. Spatial smoothing during training (see Eq. 3.17) used a Gaussian with covariance $\Sigma = \text{diag}(4 \times 10^{-4}, 4 \times 10^{-4}, 4 \times 10^{-4})$. The overall loss (see Eq. 3.20) employed a regularization weight $\lambda = 2 \times 10^{-13}$, with angular-to-radial weighting $\lambda_l = 1$ and $\lambda_n = 2$ in Eq. 3.24.

Training used the Adam optimizer, starting at a learning rate of 5×10^{-3} and halving it every five epochs [72]. QUCCI was trained for 100 epochs with a batch size of 21025. Because the signal representation includes the radial component, a single QUCCI model was trained per acceleration factor, covering all diffusion shells.

At inference, the Gaussian covariance was set to zero and fed coordinates to the MLP; the forward model (see Eq. 3.18) was then used to convert the predicted SPF coefficients into voxel intensities at q-space locations not seen during training. All computations ran in PyTorch on an NVIDIA RTX 3090 GPU.

3.3.4 Quantitative Assessments

The qualities of the predicted dMRI images and the downstream diffusion metrics were assessed quantitatively via Peak SNR (PSNR) and Structural Similarity Index Measure (SSIM) metrics. Three baselines were adopted as comparison techniques:

1) Undersampled: In this baseline, q-space interpolation was not performed and the available data were directly used for the downstream tasks.

2) SPF: In this baseline, SPF coefficients were estimated via linear least squares estimation [59], while keeping the same physics-driven regularization used in QUCCI (i.e., $\mathbf{a^T \Lambda a}$ for radial and angular q-space regularization). For hyper-parameter tuning, one subject was randomly picked from the dataset and the hyperparameters were adjusted to minimize the mean squared error between the previously unseen reconstructed volumes for each shell and the ground truths for all acceleration rates. Accordingly, the radial truncation order was chosen as $N = 1$ and the angular truncation order as $L = 8$ with $\lambda_{reg} = 10^{-7}$ and $\lambda_l = \lambda_n = 1$.

3) NeSH: NeSH is a previously proposed learning-based method that takes voxel-space coordinates as input and employs coordinate-based networks to predict SH coefficients for each voxel. However, NeSH supports only single-shell q-space upsampling [58]. Accordingly, separate NeSH models were trained for each shell for a given acceleration rate, then the outputs from these individual shells were merged to assemble a multi-shell data. Hyperparameters were adopted from the original work [58]: an angular truncation order of $L = 2$ and a regularization weight of $\lambda_{reg} = 10^{-5}$.

Table 3.1: Quantitative assessment of image quality for interpolated dMRI scans. Mean (± standard deviation) PSNR and SSIM values computed over all slices from 10 subjects.

Acceleration Rate	Method	1st Shell		2nd Shell		3rd Shell	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
$R = 10$ (3-shell)	SPF	35.42 (4.2)	96.86 (2.1)	**34.41 (4.5)**	95.93 (2.8)	32.23 (4.3)	93.64 (4.4)
	NeSH	32.09 (4.1)	93.57 (4.1)	30.16 (3.6)	91.69 (5.2)	28.10 (3.9)	88.99 (7.1)
	QUCCI	**35.49 (4.1)**	**97.09 (1.9)**	34.25 (4.1)	**96.05 (2.7)**	**32.68 (3.9)**	**94.31 (3.9)**
$R = 15$ (3-shell)	SPF	**33.63 (4.4)**	95.33 (3.2)	**31.99 (4.7)**	93.20 (4.8)	30.18 (4.4)	90.53 (6.5)
	NeSH	28.97 (4.2)	88.27 (7.4)	28.09 (3.7)	86.52 (8.4)	26.87 (3.8)	84.77 (9.6)
	QUCCI	33.26 (4.2)	**95.34 (3.1)**	31.94 (4.3)	**93.56 (4.4)**	**30.56 (4.1)**	**91.42 (5.8)**
$R = 15$ (2-shell)	SPF	**35.00 (4.3)**	96.62 (2.3)	32.15 (4.3)	94.56 (3.7)	29.85 (4.1)	91.39 (5.8)
	NeSH	32.09 (4.1)	93.57 (4.1)	-	-	28.10 (3.9)	88.99 (7.1)
	QUCCI	34.91 (4.1)	**96.70 (2.2)**	**32.27 (4.1)**	**94.81 (3.5)**	**30.32 (3.9)**	**92.34 (5.1)**
$R = 22.5$ (2-shell)	SPF	**33.23 (4.5)**	**94.98 (3.4)**	30.78 (4.4)	92.04 (5.5)	28.81 (4.2)	88.61 (7.7)
	NeSH	28.97 (4.2)	88.27 (7.4)	-	-	26.87 (3.8)	84.77 (9.6)
	QUCCI	32.77 (4.1)	94.90 (3.4)	**30.82 (4.1)**	**92.55 (5.0)**	**29.14 (3.8)**	**89.75 (6.8)**

Table 3.2: Quantitative performance comparison of different methods for downstream NODDI, MSMT-CSD, and DTI metrics. Mean (± standard deviation) PSNR and SSIM values computed over all slices from 10 subjects.

		NODDI				**MSMT-CSD**				**DTI**			
Acceleration Rate	**Method**	**ODI**		**ICVF**		**AFD**		**nuFO**		**FA**		**MD**	
		PSNR	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**
$R = 10$ (3-shell)	Undersampled	21.24 (4.4)	88.52 (6.9)	20.83 (4.5)	87.70 (6.5)	26.37 (3.1)	92.23 (4.4)	18.95 (3.7)	82.75 (9.4)	28.90 (3.6)	94.71 (3.5)	26.60 (4.5)	96.40 (2.4)
	SPF	24.06 (4.4)	90.89 (5.5)	**28.57 (4.7)**	94.89 (3.2)	28.93 (3.5)	93.77 (3.7)	**19.17 (3.7)**	83.23 (9.2)	27.88 (4.1)	94.93 (3.1)	41.02 (3.9)	98.80 (0.8)
	NeSH	22.29 (4.7)	87.09 (7.9)	26.02 (4.6)	92.65 (4.4)	26.13 (4.0)	89.17 (6.5)	17.61 (3.9)	79.81 (11.2)	26.05 (4.3)	90.85 (5.7)	37.95 (4.0)	98.38 (1.0)
	QUCCI	**25.05 (4.4)**	**92.21 (4.7)**	28.24 (4.6)	**95.06 (3.1)**	**29.55 (3.5)**	**94.29 (3.5)**	19.07 (3.7)	**83.26 (9.2)**	**29.30 (3.9)**	**95.73 (2.7)**	**41.26 (4.2)**	**99.02 (0.7)**
$R = 15$ (3-shell)	Undersampled	17.14 (4.7)	82.59 (10.4)	19.70 (4.2)	84.86 (8.0)	24.30 (3.4)	88.76 (6.5)	**18.47 (3.7)**	80.05 (10.9)	18.54 (3.0)	81.65 (9.5)	25.98 (4.6)	95.51 (3.0)
	SPF	21.61 (4.6)	86.45 (7.9)	26.92 (4.7)	93.30 (4.1)	25.62 (4.5)	88.89 (6.3)	17.91 (3.9)	79.70 (11.2)	23.99 (4.7)	89.86 (5.7)	38.76 (3.8)	98.09 (1.3)
	NeSH	19.34 (4.8)	81.07 (11.3)	24.20 (4.6)	89.29 (6.2)	22.90 (5.3)	80.14 (13.8)	17.21 (4.1)	76.38 (13.3)	22.18 (4.7)	83.60 (9.3)	35.55 (3.6)	97.58 (1.4)
	QUCCI	**22.70 (4.5)**	**88.31 (6.7)**	**26.97 (4.6)**	**93.81 (3.7)**	**26.53 (4.2)**	**90.38 (5.5)**	18.20 (3.7)	**80.49 (10.7)**	**25.38 (4.4)**	**91.74 (4.7)**	**39.54 (3.8)**	**98.47 (1.0)**
$R = 15$ (2-shell)	Undersampled	19.51 (4.5)	85.48 (8.7)	19.98 (4.3)	85.95 (7.4)	25.31 (3.1)	90.60 (5.3)	18.54 (3.8)	81.11 (10.4)	26.52 (3.5)	92.65 (4.6)	24.76 (4.4)	95.18 (3.1)
	SPF	22.57 (4.4)	88.57 (6.8)	26.10 (4.8)	93.56 (3.9)	27.90 (3.8)	92.30 (4.5)	**18.71 (3.7)**	81.63 (10.1)	26.34 (4.1)	93.08 (4.1)	37.01 (3.9)	98.15 (1.2)
	NeSH	21.44 (4.6)	85.78 (8.5)	24.90 (4.6)	92.33 (4.5)	25.95 (3.9)	88.59 (6.8)	17.44 (3.9)	79.05 (11.7)	25.70 (4.2)	89.94 (6.2)	36.92 (3.8)	98.27 (1.1)
	QUCCI	**23.61 (4.5)**	**90.64 (5.7)**	**27.01 (4.6)**	**94.33 (3.4)**	**28.30 (3.7)**	**92.68 (4.4)**	18.62 (3.7)	**81.72 (10.1)**	**28.34 (3.8)**	**94.43 (3.5)**	**38.01 (3.9)**	**98.62 (0.9)**
$R = 22.5$ (2-shell)	Undersampled	15.87 (4.9)	79.55 (12.2)	19.06 (4.2)	83.52 (8.7)	21.83 (4.5)	81.51 (14.1)	**18.21 (3.7)**	78.86 (11.6)	15.48 (3.0)	77.03 (11.8)	23.72 (4.6)	92.89 (4.8)
	SPF	20.34 (4.7)	84.63 (9.0)	25.84 (4.5)	92.39 (4.5)	23.40 (5.0)	83.67 (10.6)	17.37 (3.9)	78.39 (12.0)	23.05 (4.7)	88.01 (6.7)	36.64 (3.8)	97.44 (1.7)
	NeSH	18.96 (4.9)	79.89 (11.9)	23.98 (4.5)	89.63 (6.1)	21.54 (5.2)	77.00 (14.9)	17.18 (4.0)	75.82 (13.6)	22.32 (4.7)	83.27 (9.6)	35.22 (3.6)	97.58 (1.4)
	QUCCI	**21.72 (4.7)**	**87.32 (7.4)**	**26.89 (4.5)**	**93.54 (3.8)**	**25.07 (4.6)**	**87.69 (7.0)**	17.61 (3.9)	**79.41 (11.3)**	**24.80 (4.5)**	**90.60 (5.3)**	**37.99 (3.6)**	**98.12 (1.2)**

3.4 Results

Table 3.1 presents quantitative image quality metrics for interpolated dMRI scans compared to the reference scans. To avoid bias from noise, the reference dMRI scans were denoised using SPF with the full q-space data. In the highest shell (i.e., the third shell), QUCCI achieves a significantly higher performance than all other baselines ($p < 0.05$, paired Wilcoxon signed-rank test). This improvement is particularly pronounced at the extreme acceleration rates of $R = 22.5$ (2-shell) and $R = 15$ (2-shell), indicating that even for the case of severe q-space undersampling, QUCCI can effectively learn the implicit fiber representation and more accurately reconstruct dMRI scans for unseen q-space locations. In the second shell, QUCCI also significantly outperforms all baselines at the highest acceleration rates ($p < 0.05$). For the first shell, QUCCI and SPF yield comparable results. It is important to note that SPF's performance is expected to increase when SNR is higher, as is the case in lower shell dMRI scans. Nonetheless, downstream tasks typically rely on multi-shell data to accurately estimate volume fractions of different tissue types and to model fiber orientation distributions [10]. Accordingly, QUCCI's ability to interpolate higher-shell dMRI scans suggests potential gains in these downstream applications.

A representative qualitative comparison is presented in Fig. 3.2, displaying interpolated dMRI scans across all shells at $R = 10$ (3-shell). The superiority of QUCCI is most pronounced in the higher shells. In every case, QUCCI faithfully preserves the implicit diffusion contrast of the dMRI scans without overfitting the noise inherent to undersampled data.

Table. 3.2 reports the outcomes for downstream diffusion metrics, illustrating how the interpolated shells and directions influence the overall performance. In the NODDI-derived metrics of ODI and ICVF, QUCCI outperforms all other methods in terms of both PSNR and SSIM at every acceleration rate ($p < 0.05$), with the lone exception of PSNR for ICVF at $R = 10$. Because the NODDI algorithm requires a densely sampled q-space for accurate reconstruction, PSNR values remain relatively low across all ODI measurements [11]. Nonetheless, QUCCI

consistently achieves high SSIM—even when PSNR is comparatively reduced. This superior performance is illustrated in the example results in Fig. 3.3(a) for $R = 10$ (3-shell), where QUCCI more faithfully reproduces the reference ODI and ICVF maps under extreme undersampling, and preserves fine structural details.

For the MSMT-CSD metrics, QUCCI delivers superior AFD reconstruction in term of both PSNR and SSIM across all acceleration rates ($p < 0.05$). In the case of nuFO, QUCCI's PSNR is marginally lower than that of the undersampled data and the SPF method at every acceleration rate. However, since nuFO assumes only six discrete levels (between 0–5), even a one-pixel difference can cause substantial PSNR fluctuations, making SSIM a more reliable indicator of model performance for this case. Accordingly, QUCCI significantly outperforms all other methods in SSIM for every acceleration rate ($p < 0.05$). The example results in Fig. 3.3(b) illustrate QUCCI's advantage over competing approaches for the MSMT-CSD metrics. Because AFD encodes higher-order frequency components, interpolating unseen directions allows QUCCI to more accurately reconstruct fiber-orientation peaks. For nuFO, QUCCI's benefits are most apparent in regions with a high number of fiber orientations.

For the DTI metrics, QUCCI consistently outperforms all other methods in terms of both PSNR and SSIM across every tested acceleration rate ($p < 0.05$). Although the DTI model predates MSMT-CSD and NODDI and typically requires fewer dMRI acquisitions, it struggles to fit anisotropy values—especially FA—under extreme undersampling conditions such as $R = 22.5$ (2-shell) and $R = 15$ (3-shell). This observation highlights that, despite its lower sampling demands, DTI remains sensitive to the angular distribution of available q-space data in order to optimally estimate anisotropy. In contrast, QUCCI's ability to sample dMRI data at arbitrary radial and angular positions enables it to excel in classical downstream analyses like DTI, as demonstrated in the example results in Fig. 3.3(c). Even at a moderate acceleration rate of $R = 10$, QUCCI delivers superior quality maps and captures high-frequency features that are lost in NeSH reconstructions.

To illustrate the impact of imputed shells, Fig. 3.4 compares $R = 15$ with two

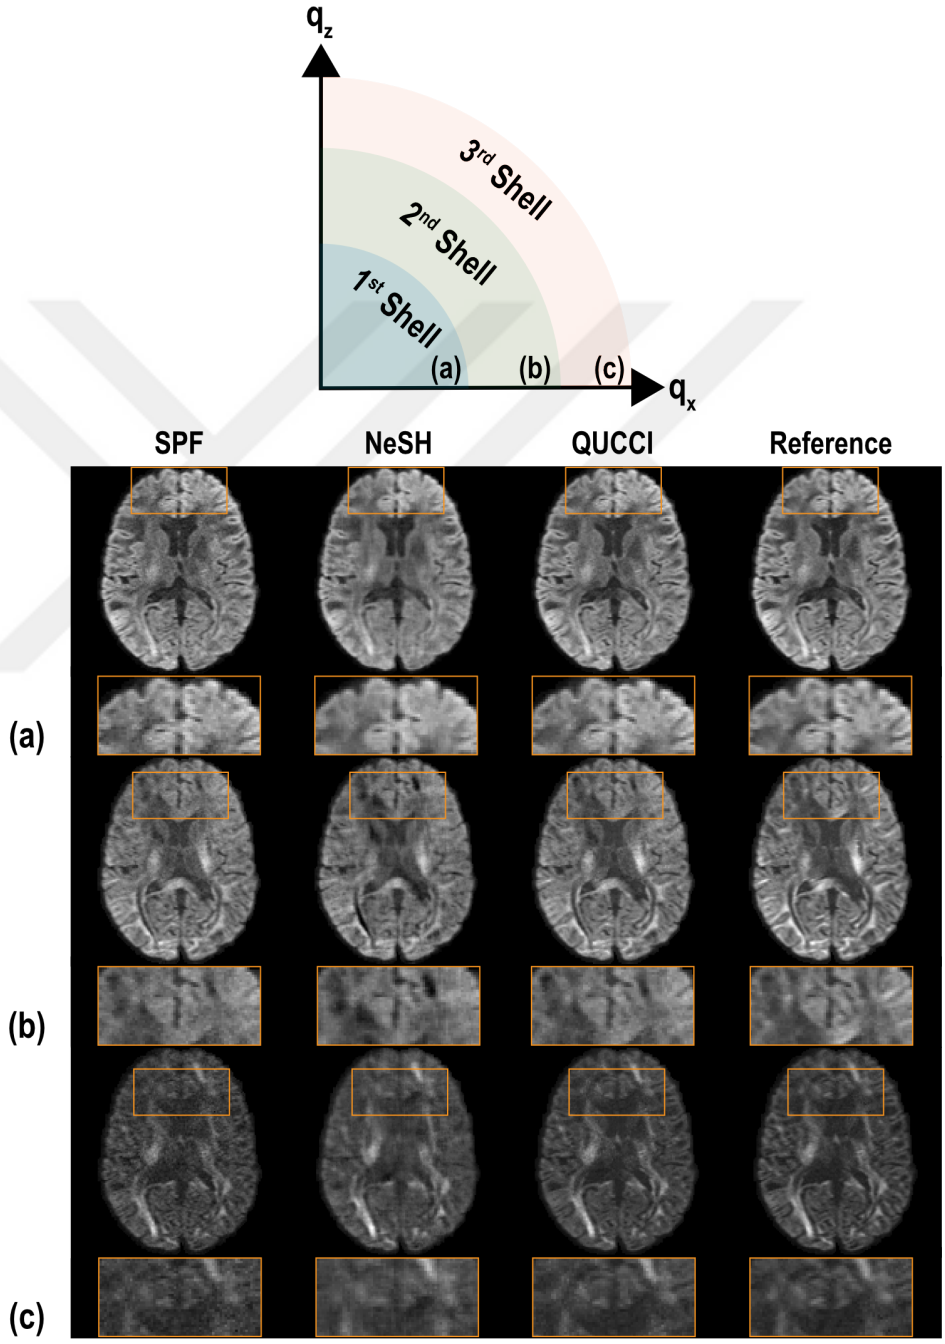

Figure 3.2: Illustrative results for the interpolated dMRI scans and the denoised reference for**(a)** the first shell, **(b)** the second shell, and **(c)** the third shell, at an acceleration rate of $R = 10$ (3-shell). QUCCI exhibits enhanced performance, particularly in the higher shells.

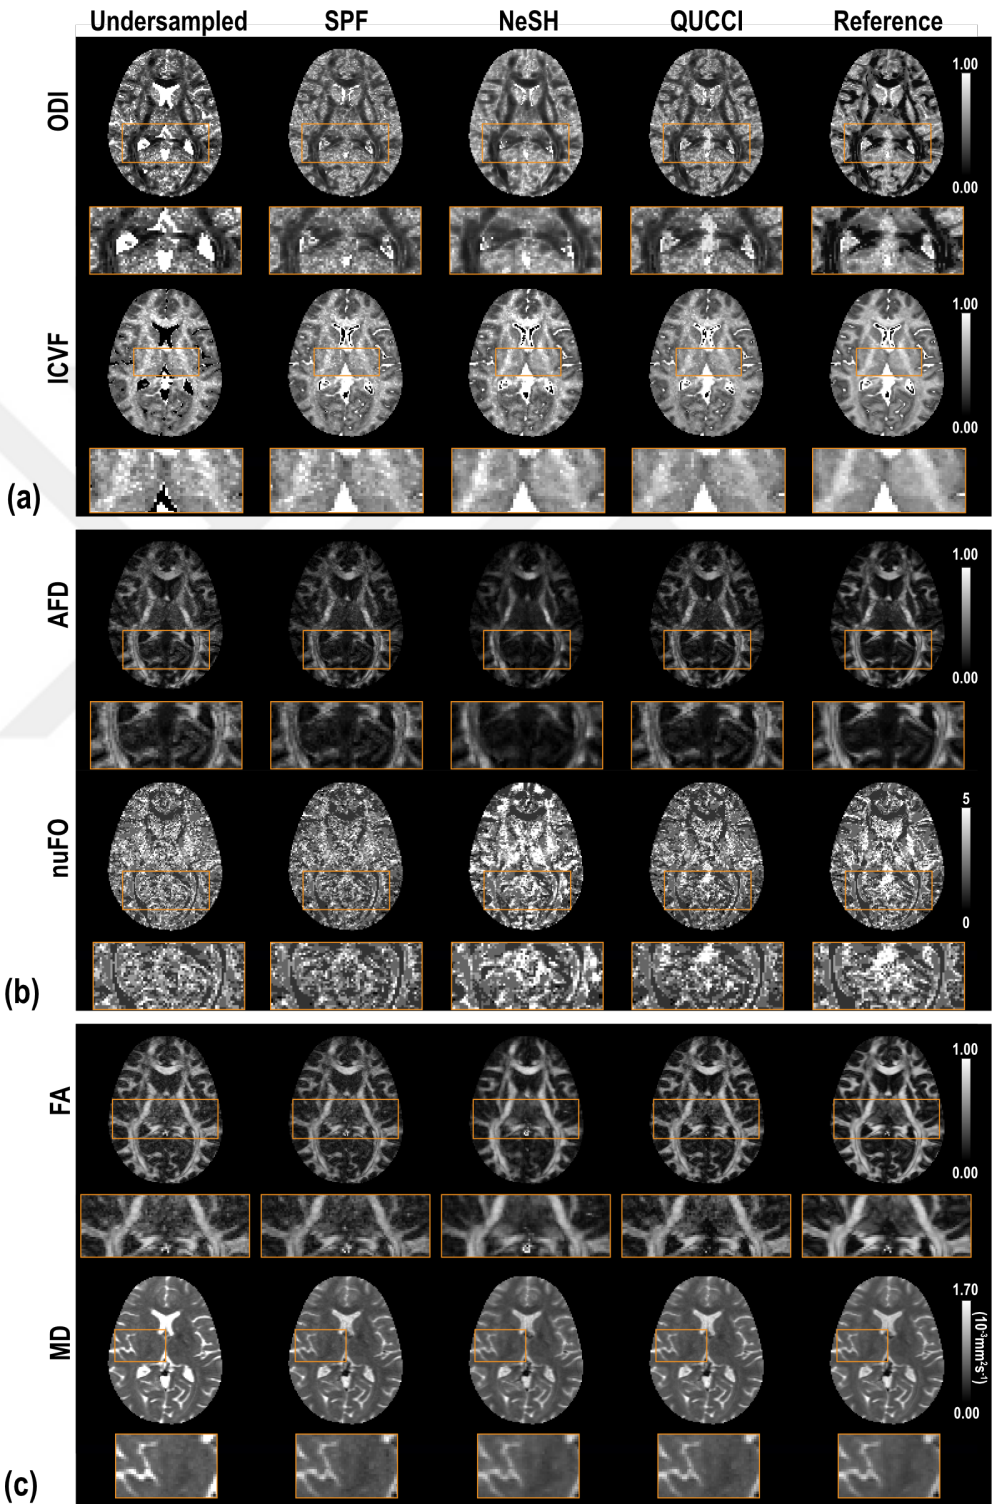

Figure 3.3: Illustrative comparison of different methods on dMRI metrics for **(a)** NODDI, **(b)** MSMT-CSD, and **(c)** DTI at an acceleration rate of $R = 10$ (3-shell).

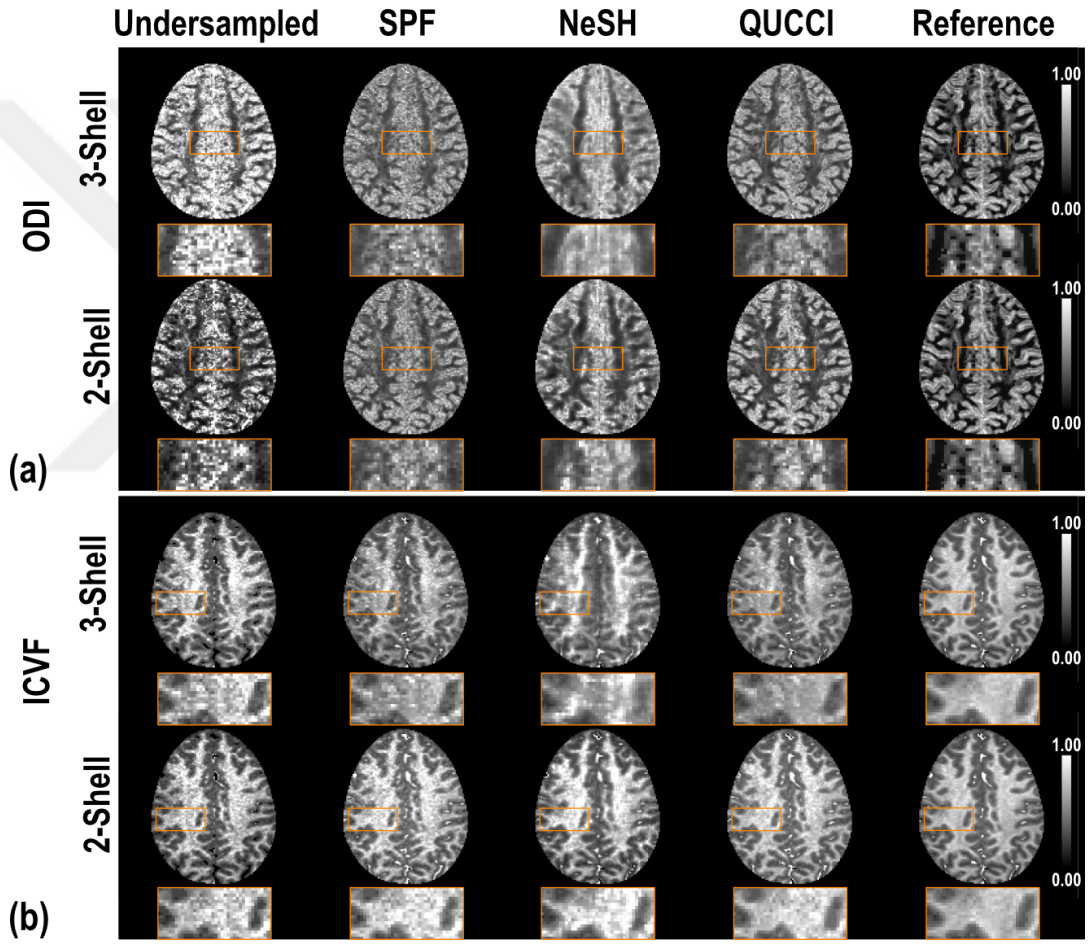

Figure 3.4: Illustrative comparison showing how varying the number of shells and the number of selected directions per shell during training influences the quality of dMRI interpolation at $R = 15$ for **(a)** the ODI metric and **(b)** the ICVF metric.

shells versus three shells for NODDI-derived metrics. For both ODI (Fig. 3.4(a)) and ICVF (Fig. 3.4(b)), QUCCI achieves the best results. However, across all approaches, training with $R = 15$ (2-shell) yields the highest accuracy, suggesting that allocating more angular samples per shell—even with fewer shells—is more effective for QUCCI and other methods than increasing the total number of shells at the cost of angular resolution.

3.4.1 Ablation Study

To quantify the impact of each subcomponent of QUCCI, an ablation study was conducted on five subjects from the HCP dataset at an acceleration factor of $R = 15$ (3-shell). Individually, random sampling, the angular regularization term, the radial regularization term, and the hash encoding were removed. The resulting PSNR and SSIM scores reported in Table 3.3 and the qualitative comparisons in Fig. 3.5 reveal that the full model consistently outperforms its ablated variants, retaining critical details that are lost in other variants.

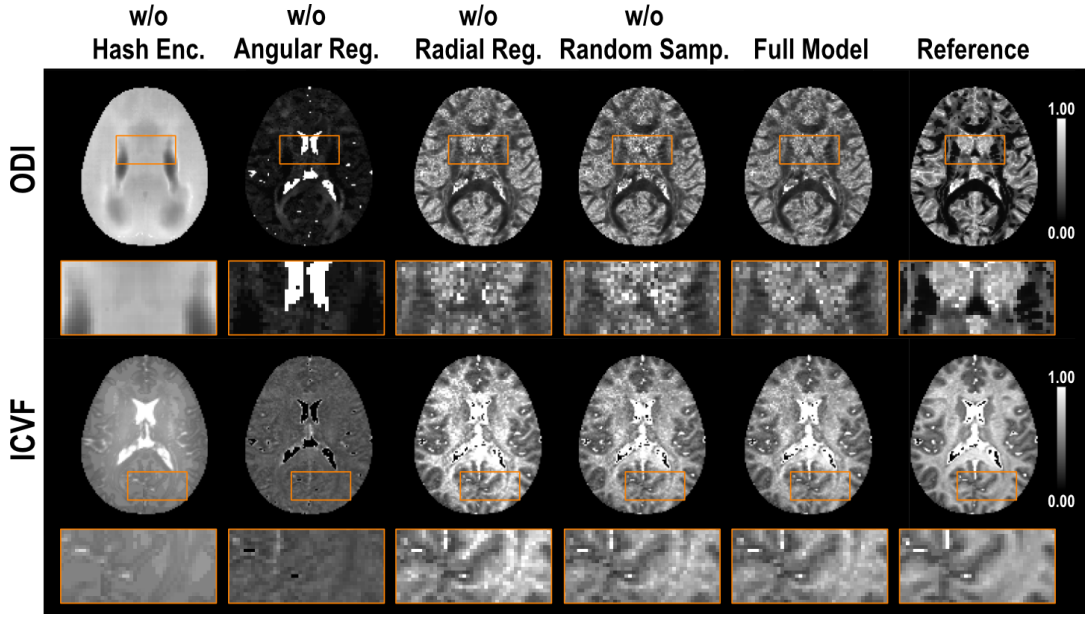

Figure 3.5: Representative comparison of how ablated models change the ODI and ICVF metrics for $R = 15$ (3-shell).

Table 3.3: Comparison of different ablated variants of the QUCCI model for NODDI, MSMT-CSD, and DTI metrics at $R = 15$ (3-shell). Mean ($\pm$ standard deviation) PSNR and SSIM values are computed over all slices from 5 subjects.

	NODDI				**MSMT-CSD**				**DTI**			
	ODI		**ICVF**		**AFD**		**nuFO**		**FA**		**MD**	
Methods	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**	**PSNR**	**SSIM**
full model	**22.47 (4.4)**	87.75 (7.1)	**26.76 (4.5)**	**93.56 (3.9)**	**26.38 (4.1)**	**90.06 (5.7)**	**18.02 (3.8)**	**79.97 (11.1)**	25.18 (4.5)	91.38 (4.9)	**39.52 (3.7)**	**98.40 (1.0)**
w/o random sampling	22.28 (4.5)	86.92 (7.6)	26.52 (4.6)	93.15 (4.2)	26.18 (4.0)	89.63 (6.0)	17.93 (3.8)	79.49 (11.4)	25.22 (4.4)	90.98 (5.2)	38.74 (3.9)	98.17 (1.2)
w/o radial regularization	22.35 (4.6)	**88.30 (6.9)**	26.19 (4.5)	92.79 (4.3)	26.17 (4.2)	89.62 (6.1)	17.86 (3.8)	79.85 (11.2)	**25.52 (4.5)**	**91.76 (4.8)**	38.11 (4.1)	97.92 (1.4)
w/o angular regularization	16.01 (4.4)	67.13 (17.9)	16.46 (4.2)	77.84 (11.9)	16.81 (3.9)	65.24 (17.9)	4.21 (7.0)	14.24 (29.2)	13.54 (3.5)	72.33 (14.3)	8.70 (3.8)	78.87 (10.6)
w/o hash encoding	14.86 (5.3)	72.55 (16.0)	22.19 (4.7)	80.74 (11.6)	21.07 (4.8)	72.03 (15.9)	16.40 (4.1)	69.57 (17.1)	16.99 (4.8)	69.17 (16.6)	33.20 (3.4)	96.82 (1.6)

Ablation experiments demonstrate that both hash encoding and angular regularization are particularly indispensable for QUCCI. To assess the impact of hash encoding, it was substituted with the positional encoding from NeSH, using a frequency expansion parameter of 12 to capture high-frequency dMRI details. In this configuration, the model fails to recover high-frequency components within 100 epochs and requires substantially longer training to reach convergence.

Eliminating angular regularization prevents the estimated SPF coefficients from exhibiting angular smoothness, rendering them vulnerable to noise and causing the model to overfit the undersampled volumes used during training. Consequently, metrics derived from interpolated scans become unreliable.

When radial regularization is omitted, radial smoothness in the SPF coefficients is lost, leading to noise between shells and an inability to faithfully model the mono-exponential decay of the dMRI signal as shell count increases.

Finally, without random sampling, the model's capacity to remove outliers diminishes, which serves as a critical factor for obtaining accurate metrics in MSMT-CSD and NODDI. MSMT-CSD depends on consistent, smooth fiber orientations across neighboring voxels for tractography, while NODDI requires coherent axonal alignment within each voxel to assess neurite integrity.

Chapter 4

sq-QUCCI for Super-Resolution in RF Encoding and q-Space Dimensions

This chapter is based in parts on the publication titled:

> **A. Topcu**, C. Liao, T. Çukur, K. Setsompop and E. U. Saritas, "Super-resolution across RF encoding and q-space dimensions via physics-driven neural fields for accelerated gSlider diffusion MRI," *Proceedings of the 33rd Annual Meeting of ISMRM*, Honolulu, May 2025.

4.1 Theory

Generalized Slice Dithered Enhanced Resolution (gSlider) is a previously proposed RF encoding method that addresses the dilemma between SNR and spatial resolution in diffusion MRI. In a conventional scan, achieving very thin slices

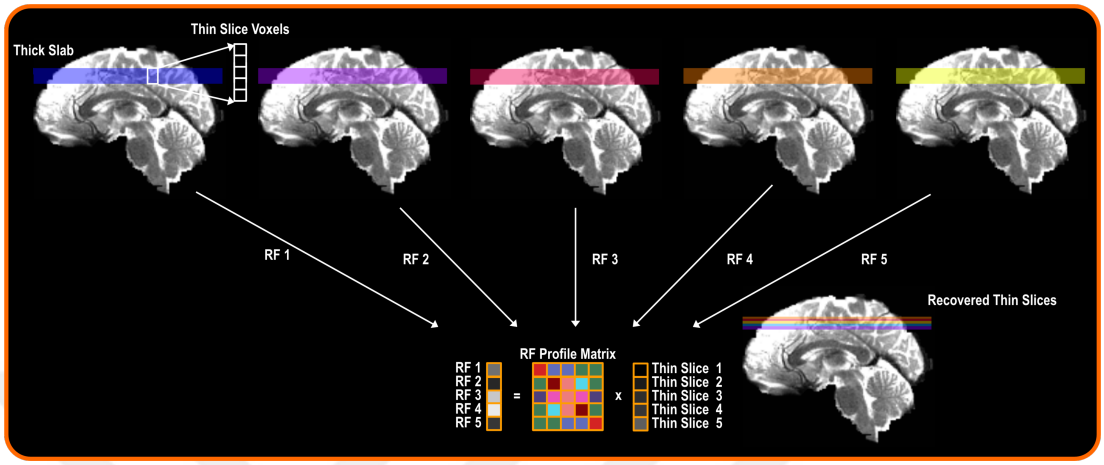

Figure 4.1: Schematic of the gSlider RF-encoding process. A single thick slab (e.g., 5 mm) is excited sequentially with five distinct RF phase encodings (labeled as RF 1–RF 5), each producing a different weighted combination of sub-slab "thin-slice" signals (colored strips). These slab responses populate the rows of the RF-profile encoding matrix, which when inverted recovers the individual thin-slice signals at sub-millimetric resolution, yielding the final high-resolution thin slices. Images generated using the preprocessed dMRI data from the Human Connectome Project [1].

(e.g., 0.7 mm) means exciting and reading out each thin slice, which greatly reduces SNR. Instead, gSlider excites a much thicker slab (e.g., 5 mm) but does so multiple times, each time using a slightly different RF pulse that "tags" the slab with a unique phase pattern. These phase-encoded acquisitions effectively mix together signals from a set of virtual thin slices inside the slab. After collecting this series of RF-encoded signals, gSlider performs a small matrix inversion (i.e., "unmixing") step: because the RF pulses are chosen to form an orthogonal basis along the slab direction, one can solve a linear system to recover each of the individual thin slice signals. The overall framework for gSlider can be seen in Fig. 4.1. The result is equivalent to directly acquiring sub-millimeter slices, but with higher SNR [14].

To speed things up, gSlider combines slab encoding with simultaneous multi-slice (SMS) acquisition: multiple slabs are excited and phase-encoded in parallel by a multiband RF pulse, and their signals are untangled both by the RF unmixing and by distinct coil sensitivity profiles. On high performance Connectom

scanners, this hybrid gSlider+SMS approach yields 660–760 µm isotropic voxels in about 25 minutes instead of the 60+ minutes that standard thin slice scans would require. Because gSlider relies on precise RF phase patterns and coil sensitivities, it is most effective on systems with uniform B_1 fields, stable B_0, and high gradient amplitudes [14].

As an example, images from 5 RF-encoded thick slabs with 2.5 mm slab thickness can be used to reconstruct images of 5 thin slices with 0.5 mm slice thickness. Formal equation to for this reconstruction problem is as follows [14, 15]:

$$\min_{\mathbf{S}} \ \frac{1}{2}\sum_{k=1}^{K}\left\|\mathbf{Y}_k^{\dagger}-\mathbf{D}_k\mathbf{S}\right\|_2^2 \ + \ \lambda_{\mathrm{Tik}}\left\|\mathbf{S}\right\|_2^2, \qquad K=5. \tag{4.1}$$

where $\mathbf{Y}_k^{\dagger}$ is the phase-corrected real gSlider data acquired with the k-th RF encoding basis. $\mathbf{D}_k$ is the complex gSlider down-sampling (forward) operator that maps an isotropic thin-slice volume to the k-th RF-encoded thick slice. Cross-talk between adjacent slabs is implicitly captured because each profile is simulated with the full Bloch model. $\mathbf{S}$ is the unobserved isotropic, high-resolution thin-slice dMRI volume to be estimated. λ_{Tik} is the Tikhonov regularization weight controlling the energy (smoothness) of $\mathbf{S}$ [14].

Because Eq. (4.1) decouples across (x, y) locations, it is solved independently at every in-plane voxel via the normal equations $\left(\sum_k \mathbf{D}_k^{\dagger}\mathbf{D}_k + \lambda_{\mathrm{Tik}}\mathbf{I}\right)\mathbf{S} = \sum_k \mathbf{D}_k^{\dagger}\mathbf{Y}_k^{\dagger}$, The RF bases are designed to be nearly orthogonal, yet the encoding matrix $\sum_k \mathbf{D}_k^{\dagger}\mathbf{D}_k$ remains mildly ill-conditioned because adjacent slab profiles overlap and because off-resonance effects perturb their ideal shapes. The Tikhonov term stabilizes the inversion.

4.2 Methods

4.2.1 Dataset

A single subject was scanned on a 3 T GE "Ultra-High-Performance" scanner equipped with a 32-channel Nova head coil, using the gSlider-BUDA circular-EPI sequence from Setsompop Lab at Stanford University [73].

Whole-brain 500 μm in-plane resolution gSlider-SMS data parameters were: 32 slabs (2.5 mm thick) with the same five RF encodings, matrix 440×440 (FOV = $220 \times 220\,\mathrm{mm}^2$), TR/TE = 3500/65 ms, multiband R_{MB}=2 and in-plane R_{PE}=4 (dual 5/8 Partial Fourier). Each volume contained ten $b = 0$ and 50 diffusion directions at $b = 1000\,\mathrm{s/mm}^2$. The protocol was repeated three times to boost SNR (total scan time of 105 min).

A FOV-matched, low-resolution B_1^+ map was also acquired with a 1-min Bloch–Siegert sequence at $3.4 \times 3.4 \times 5\,\mathrm{mm}^3$ resolution to correct B_1^+ inhomogeneity induced slab profile imperfections in gSlider reconstruction [74].

For sq-QUCCI, the gSlider data was retrospectively undersampled: only two of the five RF encodings were retained (R_{RF}=2) and 14 of the 50 uniformly spaced diffusion directions were kept ($R_q \approx 3.5$), giving an overall acceleration of $R \approx 7$. The resulting acquisition time for the undersampled data corresponds to a ~ 15 minute scan.

4.2.2 sq-QUCCI Implementation

While gSlider method for dMRI boosts SNR efficiency relative to single-band acquisitions, it introduces two fundamental bottlenecks:

1. **RF encoding:** A full gSlider protocol typically uses 5 encodings per slab. At each q-space point those encodings must be repeated, inflating scan time

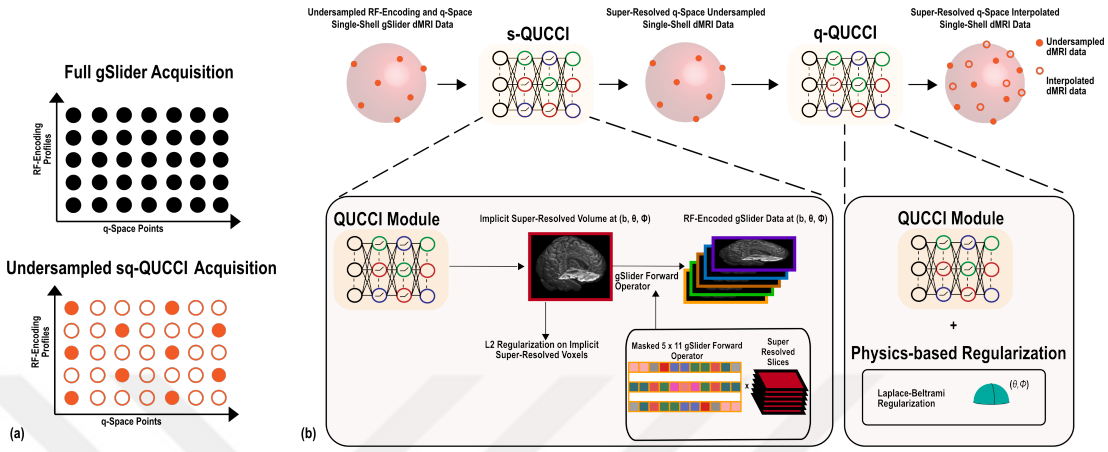

Figure 4.2: **(a)** Conventional gSlider versus undersampled sq-QUCCI acquisitions. Traditional gSlider repeats every RF encoding at every q-space location. sq-QUCCI instead uses a jointly undersampled pattern that alternates RF encodings across q-space, cutting the total scan time by a factor of 7. **(b)** Self-supervised reconstruction pipeline for sq-QUCCI. s-QUCCI first recovers missing RF encodings to generate a thin-slice volume. q-QUCCI then interpolates the sparse q-space grid to yield a fully sampled, super-resolved volume.

linearly with the number of RF encodings. Attempting to undersample the RF axis degrades slice unaliasing and makes the reconstruction ill-posed.

2. **q-space:** Micro-structural models (DTI, NODDI, CSD, etc.) require dozens of diffusion directions and multiple b-value shells. Exhaustive sampling therefore multiplies the already long RF-encoding loop, pushing whole-brain sub-millimeter protocols beyond 90–120 min and making them highly sensitive to motion and system drifts.

Because the two loops (RF encodings and q-space directions) are nested, accelerating only one loop leaves the other as a hard limit on scan efficiency. In practice, clinicians must trade spatial resolution for angular/radial resolution or accept prohibitively long sessions.

sq-QUCCI tackles these intertwined bottlenecks with two cascaded neural-field modules:

- **s-QUCCI** performs super-resolution along the RF encoding (slice) axis,

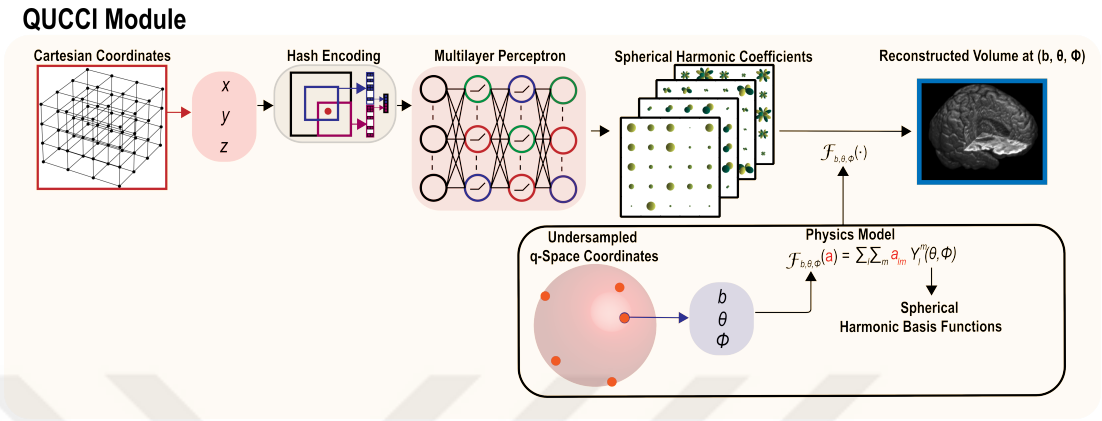

Figure 4.3: Neural field module shared by s-QUCCI and q-QUCCI. Cartesian voxel coordinates are hash-encoded at multiple resolutions, passed through an MLP, and converted into SPF (or SH) coefficients. A physics-based forward model projects these coefficients back to measurement space, enabling self-supervised consistency losses.

allowing only a subset of the RF encodings to be acquired.

- **q-QUCCI** performs super-resolution in the angular q-space domain, enabling sparse diffusion sampling and interpolation of unacquired directions.

The overall training scheme is given in Fig. 4.2. For **s-QUCCI** the model is composed using the gSlider RF encoding operator, creating a fully self-supervised loss that compares the synthetic thick-slice signal with the undersampled thick-slab measurements for which no ground-truth thin-slice data are required. After training, the s-QUCCI model is used as query for the training of **q-QUCCI**, which has the same training scheme and regularizations as the original QUCCI described in Chapter 3.

As shown in Fig. 4.3, both modules takes Cartesian voxel coordinates and output spherical-polar Fourier (SPF) coefficients via a multi-resolution hash-grid encoder followed by a multilayer perceptron (MLP). Implicit regularization arises from the spectral bias of the MLP, whereas explicit physics-driven constraints enforce [62]:

1. *Slice consistency*: an L_2 penalty between neighboring super-resolved voxels

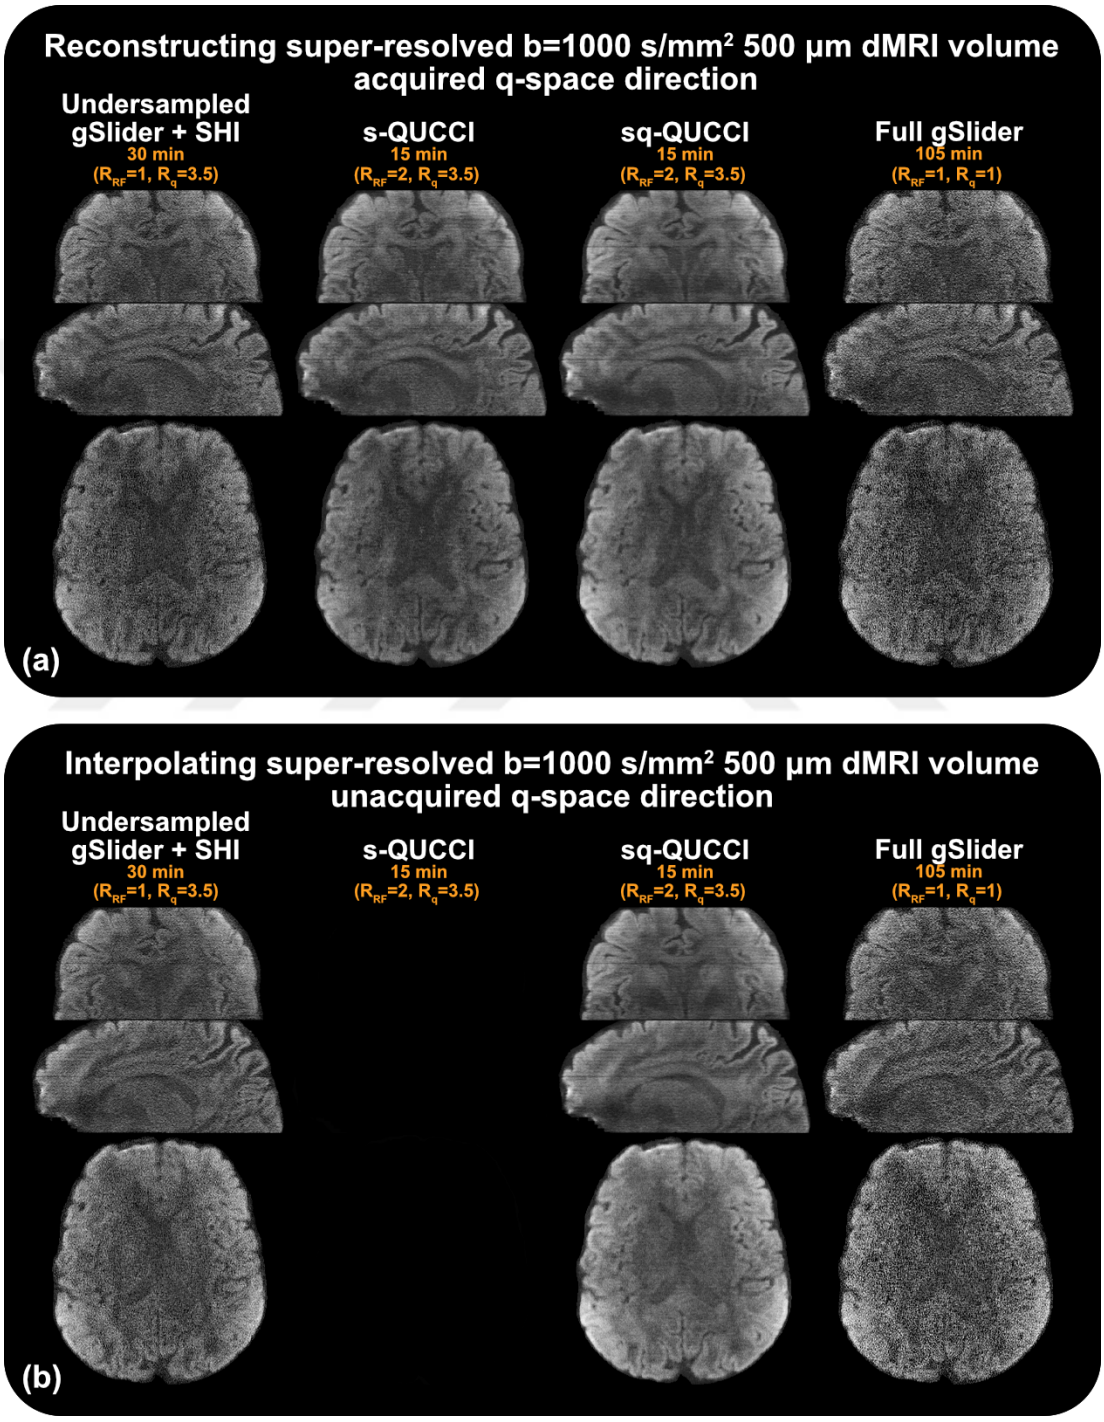

Figure 4.4: (a) Coronal, sagittal and axial slices at an *acquired* q-space location. (b) The same slices at an *unacquired* q-space location. sq-QUCCI preserves fine cortical details and white-matter contrast, while the baseline methods exhibit residual blurring and noise.

mitigates slab-boundary artifacts in **s-QUCCI**.

2. *Angular smoothness*: a single-shell Laplace–Beltrami penalty on SPF coefficients promotes directionally smooth signals in **q-QUCCI**.

s-QUCCI and q-QUCCI were trained sequentially for 14 epochs each in PyTorch with Adam. Losses compared forward-modelled predictions with acquired RF-encoded slabs (s-QUCCI) or sparse q-space samples (q-QUCCI), providing fully self-supervised optimization. For the estimation of SPF coefficients of the dMRI signal, the radial and angular truncation orders were chosen as $N = 0$ and $L = 8$, respectively for both s-QUCCI and q-QUCCI. The weight for the L_2 penalty in s-QUCCI is chosen as $\lambda = 3 \times 10^{-5}$. The weight for the Laplace-Beltrami penalty in q-QUCCI is chosen as $\lambda = 2 \times 10^{-13}$. The Gaussian sampling is turned off for s-QUCCI. The rest of the hyperparameters are directly taken as the same as in the original QUCCI model in Chapter 3.

sq-QUCCI was benchmarked against:

1. Standard gSlider reconstruction on q-space-undersampled data ($R_{\mathrm{RF}} = 1,\ R_q = 3.5$);
2. gSlider followed by spherical-harmonic interpolation (SHI) in q-space;
3. s-QUCCI alone on the jointly undersampled data ($R_{\mathrm{RF}} = 2,\ R_q = 3.5$).

All outputs were post-processed with FSL (`eddy`), AMICO and MRtrix3 to derive DTI, NODDI and CSD metrics [43, 6, 7, 11, 56].

4.3 Results

sq-QUCCI faithfully reconstructed high-resolution dMRI volumes despite a sevenfold acceleration, recovering fine anatomical details while suppressing noise amplification. Qualitatively, it best matched fully sampled gSlider results across FA,

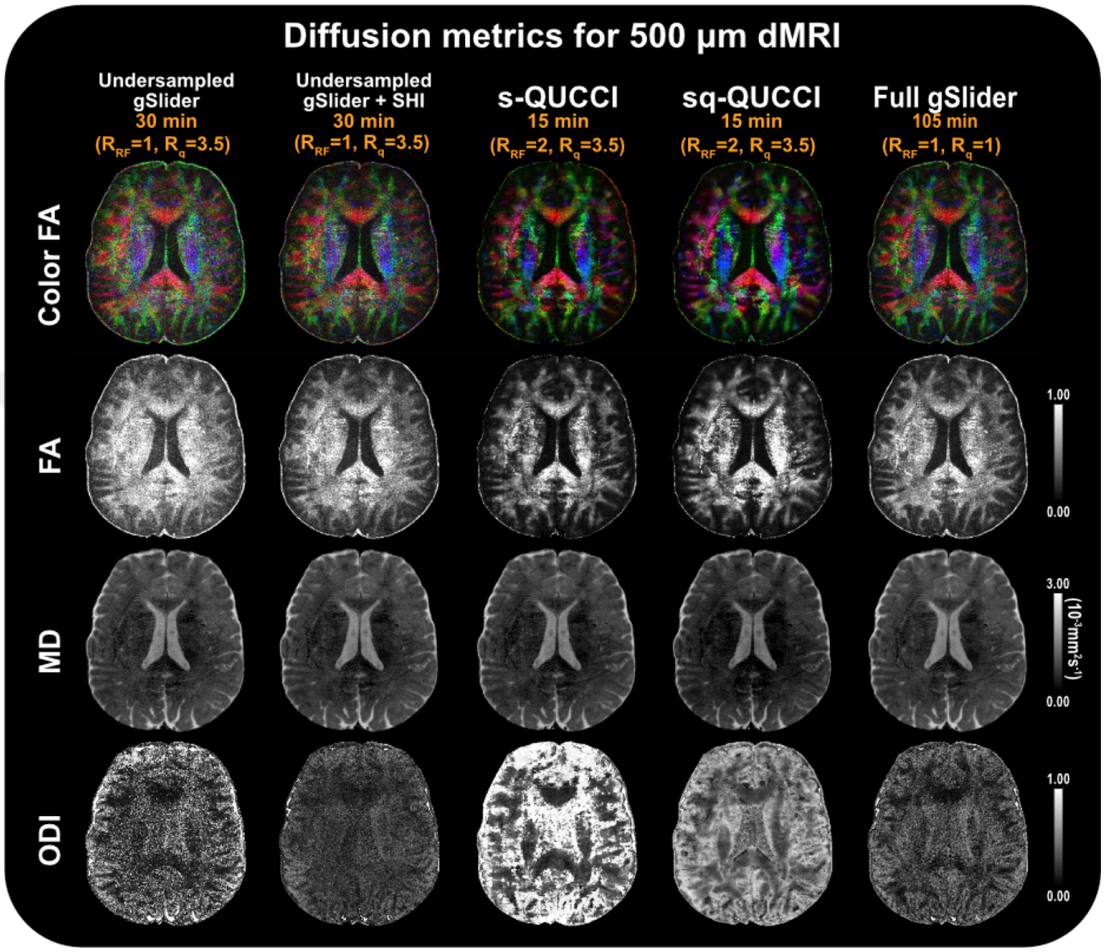

Figure 4.5: Diffusion metrics (color FA, FA, MD, ODI) for undersampled gSlider, undersampled gSlider+SHI, s-QUCCI, sq-QUCCI and fully sampled gSlider. sq-QUCCI preserves the accuracy of dMRI metrics despite aggressive acceleration.

ODI, and AFD metrics, outperforming SHI and s-QUCCI alone. Consequently, these results show that with sq-QUCCI, whole-brain 500 μm, 50-direction scans can be acquired in $\sim$ 15 min instead of 105 min.

Representative coronal, sagittal and axial slices reconstructed at an *acquired* and an *unacquired* q-space locations are shown in Fig. 4.4. Undersampled gSlider combined with SH interpolation (gSlider+SHI) retains anatomy but its SNR is similar to the noisy reference. s-QUCCI, which super-resolves only the RF-encoding direction, mitigates the noise but cannot recover angular details in

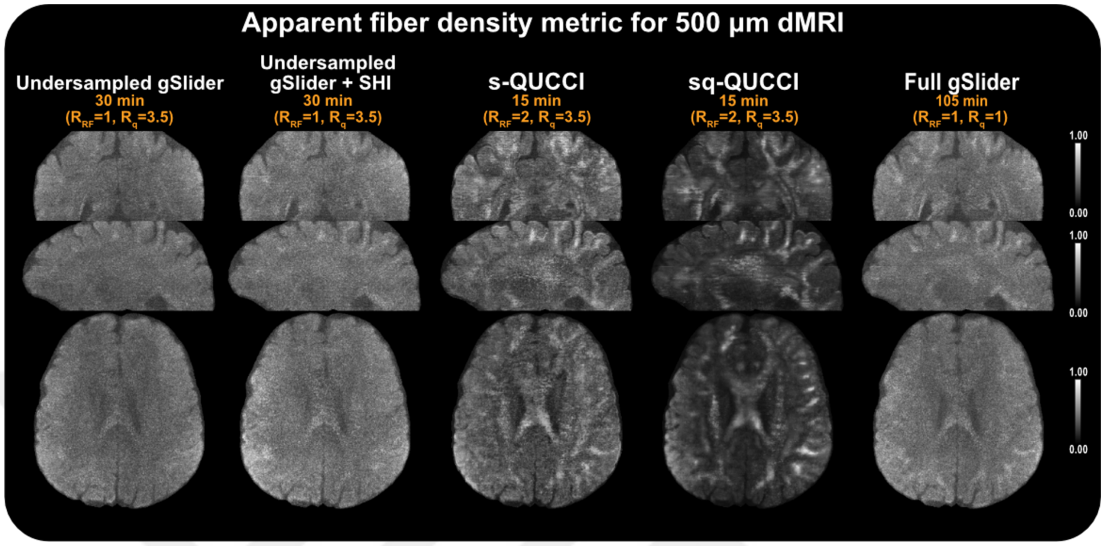

Figure 4.6: Apparent fiber density (AFD) maps. sq-QUCCI delivers the cleanest depiction of major fiber bundles, approaching the reference while using one-seventh of the scan time.

an unacquired q-space location. By contrast, sq-QUCCI preserves gray/white-matter boundaries and fine cortical folding in both *acquired* and an *unacquired* q-space locations, closely matching the fully sampled reference while maintaining low noise levels.

Voxel-wise diffusion metrics derived from DTI and NODDI are summarized in Fig. 4.5. For color FA and FA maps, s-QUCCI and sq-QUCCI reproduces the anisotropy patterns of the reference with high SNR, whereas gSlider+SHI underestimates or overestimates FA and suffers from low SNR. As for the q-space undersampled gSlider, SNR is qualitatively worse than the reference full gSlider. MD remains robust across all methods, reflecting its lower sensitivity to angular sparsity. ODI, however, is substantially problematic in all baselines, whereas sq-QUCCI recovers the dispersion difference between white matter and gray matter.

AFD maps in Fig. 4.6 further highlight sq-QUCCI's advantage. Compared with baselines, sq-QUCCI suppresses noise amplification and recovers the finer details, even surpassing the reference AFD map in terms of visual quality.

sq-QUCCI shows that physics-constrained neural fields can jointly super-resolve in RF-encoding and q-space dimensions, achieving state-of-the-art reconstructions at clinically practical scan times. Subject-specific training removes the need for large datasets, and the INR model yields compact models (a few MB) that can be archived with each exam.

Chapter 5

Discussion and Conclusion

In this thesis, QUCCI has been introduced as a scalable, physics-aware dMRI reconstruction framework based on coordinate-wise implicit neural representations. The method expresses the dMRI signal with spherical-polar Fourier (SPF) basis functions and enforces explicit angular and radial smoothness constraints derived from physical principles. A probabilistic noise model is added to suppress outliers during reconstruction. Extensive experiments on the HCP dataset demonstrate that QUCCI produces high-quality images that remain robust under aggressive q-space acceleration and against artifacts, while its off-the-grid MLP weights store the signal compactly and thus offer significant scalability over conventional approaches. QUCCI offers a particularly promising approach for accelerating neonatal and fetal dMRI scans, where scan time is limited by continual head motion [75, 76, 77].

For this thesis, DTI, MSMT-CSD, and NODDI were used to estimate diffusion orientations and to compute FA, MD, AFD, nuFO, ODI, and ICVF metrics to evaluate the performance of QUCCI. QUCCI is assessed to speed up multi-shell dMRI protocols by acceleration rates of R=10 to R=22.5. As for the gSlider single-shell dMRI protocol, the sq-QUCCI extension of QUCCI provided an acceleration rate of R=7, reducing the total acquisition time from 105 minutes to 15 minutes. In all cases, QUCCI and sq-QUCCI were both able to outperform

their baselines qualitatively and quantitatively for all acceleration rates.

Clinical and Practical Implications

The acceleration factors achieved by QUCCI translate directly into shorter scan times. For typical multi-shell protocols, the reduction from ∼55 min to 5–12 min (R = 22.5–10) makes advanced diffusion acquisitions feasible in routine clinical slots. Similarly, the sq-QUCCI application to gSlider data cuts a 105–minute high-resolution exam down to 15 minutes, a change that is especially valuable in settings involving neonates, fetuses, or uncooperative patients where motion intolerance is paramount. Because the accelerated data preserve key microstructural biomarkers (FA, MD, AFD, nuFO, ICVF, ODI) within acceptable error margins, clinicians can adopt these faster protocols without compromising downstream analyses or longitudinal comparability.

Table 5.1: Storage needed per subject for QUCCI and SPF on the HCP dataset. Because QUCCI is off-the-grid, its disk usage stays constant across resolutions in both spatial and q-space domains. The first row reports the grid dimensions, representing both the Cartesian image grid and the q-space sampling grid. Here, 270 is the total number of q-space samples, reduced to 18 for an acceleration rate of R = 15. The last column shows 2-fold spatial upsampling in all axes, yielding 8 times larger image matrix at each q-space location. Vol. = Volume.

	dMRI storage space		
Methods	*Vol.*×18	*Vol.*×270	*(8×Vol.)*×270
QUCCI	200MB	200MB	200MB
SPF	424MB	6.36GB	50.88GB

Scalability and Data Management

Table 5.1 highlights a secondary benefit of QUCCI: storage efficiency. The network weights act as a compact surrogate for the entire dMRI volume, decoupling disk usage from image resolution and q-space density. This compression is advantageous for biobanks or large patient registries that need to archive raw diffusion

data for future re-analysis. Moreover, a bank of pre-trained QUCCI weight sets could serve as a generative prior, enabling synthetic data augmentation or zero-shot reconstruction of novel protocols.

Limitations

Despite these advantages, several constraints remain:

- **Training latency:** Subject-specific training currently requires training on a high-end GPU up to a day. Integrating meta-learning or hyper-networks may cut this overhead.
- **Pathology coverage:** Validation was performed on 10 HCP subjects. The performance of QUCCI on lesions, tumors, or demyelinating disease must be assessed in future work.
- **Parallel-imaging physics:** For simplicity, the present forward model assumes a single sensitivity profile. Extending QUCCI to full multi-coil encoding could push acceleration rates further but will demand additional regularization.

This thesis demonstrates that a physics-constrained implicit neural representation can deliver rapid, storage-efficient, and artifact-resilient dMRI. Across both multi-shell and gSlider protocols, QUCCI and its variant sq-QUCCI consistently outperform classical baselines in reconstruction fidelity and downstream microstructural metrics, even under aggressive acceleration. These properties position QUCCI as a promising foundation for next-generation diffusion pipelines, where scan time, storage, and motion robustness are critical constraints.

Bibliography

[1] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, and K. Ugurbil, "The wu-minn human connectome project: An overview," *NeuroImage*, vol. 80, p. 62–79, Oct 2013.

[2] M. Descoteaux, *High Angular Resolution Diffusion MRI: from Local Estimation to Segmentation and Tractography*. Theses, Université Nice Sophia Antipolis, Feb. 2008.

[3] D. S. Novikov, E. Fieremans, and H.-H. Lee, "Chapter 5: Physics of diffusion imaging: Fundamentals," in *Handbook of Diffusion MR Tractography* (F. Dell'Acqua, M. Descoteaux, and A. Leemans, eds.), London, UK: Academic Press, 2024.

[4] Y. Shi and A. Toga, "Connectome imaging for mapping human brain pathways," *Molecular Psychiatry*, vol. 22, pp. 1230 – 1240, 2017.

[5] A. Bizzi, J. Y.-M. Yang, J. Aliaga-Arias, F. Dell'Acqua, J. P. Lavrador, and F. Vergani, "Chapter 31: Neurosurgical applications of clinical tractography," in *Handbook of Diffusion MR Tractography* (F. Dell'Acqua, M. Descoteaux, and A. Leemans, eds.), ch. 31, London, UK: Academic Press, 2024.

[6] D. L. Bihan, J. F. Mangin, C. Poupon, C. Clark, S. Pappatà, N. Molko, and H. Chabriat, "Diffusion tensor imaging: Concepts and applications," *Journal of Magnetic Resonance Imaging*, vol. 13, 2001.

[7] E. Özarslan and T. H. Mareci, "Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging," *Magnetic Resonance in Medicine*, vol. 50, no. 5, p. 955–965, 2003.

[8] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen, "High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity," *Magnetic Resonance in Medicine*, vol. 48, no. 4, pp. 577–582, 2002.

[9] M. Descoteaux, *High Angular Resolution Diffusion Imaging (HARDI)*, pp. 1–25. John Wiley & Sons, Ltd, 2015.

[10] B. Jeurissen, J.-D. Tournier, T. Dhollander, A. Connelly, and J. Sijbers, "Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion mri data," *NeuroImage*, vol. 103, p. 411–426, 2014.

[11] H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander, "Noddi: Practical in vivo neurite orientation dispersion and density imaging of the human brain," *NeuroImage*, vol. 61, no. 4, p. 1000–1016, 2012.

[12] E. Caruyer, C. Lenglet, G. Sapiro, and R. Deriche, "Design of multishell sampling schemes with uniform coverage in diffusion mri," *Magnetic Resonance in Medicine*, vol. 69, 2013.

[13] W. Consagra, A. Venkataraman, and Z. Zhang, "Optimized diffusion imaging for brain structural connectome analysis," *IEEE Transactions on Medical Imaging*, vol. 41, pp. 2118–2129, 2021.

[14] K. Setsompop, Q. Fan, J. Stockmann, B. Bilgic, S. Y. Huang, S. F. Cauley, A. Nummenmaa, F. Wang, Y. Rathi, T. Witzel, and L. L. Wald, "High-resolution in vivo diffusion imaging of the human brain with generalized slice dithered enhanced resolution: Simultaneous multislice (gslider-sms)," *Magnetic Resonance in Medicine*, vol. 79, no. 1, pp. 141–151, 2018.

[15] G. Ramos-Llordén, L. Ning, C. Liao, R. Mukhometzianov, O. Michailovich, K. Setsompop, and Y. Rathi, "High-fidelity, accelerated whole-brain submillimeter in vivo diffusion mri using gslider-spherical ridgelets (gslider-sr)," *Magnetic Resonance in Medicine*, vol. 84, no. 4, pp. 1781–1795, 2020.

[16] H. Chen, K. Dai, S. Zhong, J. Zheng, X. Zhang, S. Yang, T. Cao, C. Wang, E. Karasan, L. Frydman, and Z. Zhang, "High-resolution multishot diffusion-weighted mri combining markerless prospective motion correction and locally low-rank constrained reconstruction," *Magnetic Resonance in Medicine*, vol. 89, no. 2, pp. 605–619, 2023.

[17] M. Drake-Pérez, J. Boto, A. Fitsiori, K. Lovblad, and M. I. Vargas, "Clinical applications of diffusion weighted imaging in neuroradiology," *Insights Imaging*, vol. 9, no. 4, pp. 535–547, 2018.

[18] H. J. Kim, C. G. Choi, D. H. Lee, J. H. Lee, S. J. Kim, and D. C. Suh, "High-b-Value Diffusion-Weighted MR Imaging of Hyperacute Ischemic Stroke at 1.5T," *American Journal of Neuroradiology*, vol. 26, no. 2, pp. 208–215, 2005.

[19] R. G. González, P. W. Schaefer, F. S. Buonanno, L. H. Schwamm, R. F. Budzik, G. Rordorf, B. Wang, A. G. Sorensen, and W. J. Koroshetz, "Diffusion-weighted MR imaging: diagnostic accuracy in patients imaged within 6 hours of stroke symptom onset," *Radiology*, vol. 210, no. 1, pp. 155–162, 1999.

[20] I. Caglic and T. Barrett, "Diffusion-weighted imaging (DWI) in lymph node staging for prostate cancer," *Translational Andrology and Urology*, vol. 7, no. 5, pp. 814–823, 2018.

[21] B. Taouli, V. Vilgrain, E. Dumont, J.-L. Daire, B. Fan, and Y. Menu, "Evaluation of Liver Diffusion Isotropy and Characterization of Focal Hepatic Lesions with Two Single-Shot Echo-planar MR Imaging Sequences: Prospective Study in 66 Patients," *Radiology*, vol. 226, no. 1, pp. 71–78, 2003.

[22] M. Cova, E. Squillaci, F. Stacul, G. Manenti, S. Gava, G. Simonetti, and R. Pozzi-Mucelli, "Diffusion-weighted MRI in the evaluation of renal lesions:

preliminary results," *The British Journal of Radiology*, vol. 77, no. 922, pp. 851–857, 2004.

[23] E. Squillaci, G. Manenti, M. Cova, M. Di Roma, R. Miano, G. Palmieri, and G. Simonetti, "Correlation of Diffusion-weighted MR Imaging with Cellularity of Renal Tumours," *Anticancer Research*, vol. 24, no. 6, pp. 4175–4180, 2004.

[24] M. Eiber, K. Holzapfel, C. Ganter, K. Epple, S. Metz, H. Geinitz, H. Kübler, J. Gaa, E. J. Rummeny, and A. J. Beer, "Whole-body MRI including diffusion-weighted imaging (DWI) for patients with recurring prostate cancer: Technical feasibility and assessment of lesion conspicuity in DWI," *Journal of Magnetic Resonance Imaging*, vol. 33, no. 5, pp. 1160–1170, 2011.

[25] K. Nasu, Y. Kuroki, S. Kuroki, K. Murakami, S. Nawano, and N. Moriyama, "Diffusion-weighted Single Shot Echo Planar Imaging of Colorectal Cancer Using a Sensitivity-encoding Technique," *Japanese Journal of Clinical Oncology*, vol. 34, no. 10, pp. 620–626, 2004.

[26] M. Afzali, T. Pieciak, S. Newman, E. Garyfallidis, E. Özarslan, H. Cheng, and D. K. Jones, "The sensitivity of diffusion MRI to microstructural properties and experimental factors," *Journal of Neuroscience Methods*, vol. 347, p. 108951, 2021.

[27] K. Abhinav, F.-C. Yeh, A. Mansouri, G. Zadeh, and J. C. Fernandez-Miranda, "High-definition fiber tractography for the evaluation of perilesional white matter tracts in high-grade glioma surgery," *Neuro-Oncology*, vol. 17, no. 9, pp. 1199–1209, 2015.

[28] K. Abhinav, F.-C. Yeh, S. Pathak, V. Suski, D. Lacomis, R. M. Friedlander, and J. C. Fernandez-Miranda, "Advanced diffusion MRI fiber tracking in neurosurgical and neurodegenerative disorders and neuroanatomical studies: A review," *Biochimica et Biophysica Acta (BBA) - Molecular Basis of Disease*, vol. 1842, no. 11, pp. 2286–2297, 2014.

[29] W. Jost, *Diffusion in Solids, Liquids, Gases*. Academic Press, 1952.

[30] C. G. Koay and E. Özarslan, "Conceptual foundations of diffusion in magnetic resonance," *Concepts in Magnetic Resonance Part A*, vol. 42, no. 4, pp. 116–129, 2013.

[31] T. Schneider and C. A. M. Wheeler-Kingshott, "Chapter 3.2 - Q-Space Imaging: A Model-Free Approach," in *Quantitative MRI of the Spinal Cord*, pp. 146–155, Academic Press, 2014.

[32] E. L. Hahn, "Spin Echoes," *Physical Review*, vol. 80, pp. 580–594, 1950.

[33] H. Y. Carr and E. M. Purcell, "Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments," *Physical Review*, vol. 94, no. 3, pp. 630–638, 1954.

[34] E. O. Stejskal and J. E. Tanner, "Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient," *The Journal of Chemical Physics*, vol. 42, no. 1, pp. 288–292, 1965.

[35] D. Le Bihan, R. Turner, C. T. W. Moonen, and J. Pekar, "Imaging of diffusion and microcirculation with gradient sensitization: Design, strategy, and significance," *Journal of Magnetic Resonance Imaging*, vol. 1, no. 1, pp. 7–28, 1991.

[36] S. T. Skare and R. Bammer, "EPI-Based Pulse Sequences for Diffusion Tensor MRI," in *Diffusion MRI: Theory, Methods, and Applications*, Oxford University Press, 2010.

[37] A. W. Anderson and J. C. Gore, "Analysis and correction of motion artifacts in diffusion weighted imaging," *Magnetic Resonance in Medicine*, vol. 32, no. 3, pp. 379–387, 1994.

[38] P. Mansfield, "Multi-planar image formation using NMR spin echoes," *Journal of Physics C: Solid State Physics*, vol. 10, no. 3, pp. L55–L58, 1977.

[39] R. Turner and D. Le Bihan, "Single-shot diffusion imaging at 2.0 tesla," *Journal of Magnetic Resonance (1969)*, vol. 86, no. 3, pp. 445–452, 1990.

[40] K. M. Lüdeke, P. Röschmann, and R. Tischler, "Susceptibility artefacts in NMR imaging," *Magnetic Resonance Imaging*, vol. 3, no. 4, pp. 329–343, 1985.

[41] J. J. Van Vaals and A. H. Bergman, "Optimization of eddy-current compensation," *Journal of Magnetic Resonance (1969)*, vol. 90, no. 1, pp. 52–70, 1990.

[42] M. S. Graham, I. Drobnjak, M. Jenkinson, and H. Zhang, "Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI," *PLOS ONE*, vol. 12, no. 10, pp. 1–25, 2017.

[43] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. J. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, R. K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J. M. Brady, and P. M. Matthews, "Advances in functional and structural MR image analysis and implementation as FSL," *NeuroImage*, vol. 23, Supplement 1, pp. S208–S219, 2004.

[44] J. L. R. Andersson, S. Skare, and J. Ashburner, "How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging," *NeuroImage*, vol. 20, no. 2, pp. 870–888, 2003.

[45] L. Ruthotto, S. Mohammadi, C. Heck, J. Modersitzki, and N. Weiskopf, "Hyperelastic Susceptibility Artifact Correction of DTI in SPM," in *Bildverarbeitung für die Medizin 2013*, pp. 344–349, Springer Berlin Heidelberg, 2013.

[46] J. Modersitzki, *FAIR : flexible algorithms for image registration*. Society for Industrial & Applied Mathematics, 2009.

[47] J. L. Andersson and S. N. Sotiropoulos, "An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging," *NeuroImage*, vol. 125, p. 1063–1078, Jan 2016.

[48] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "Sense: sensitivity encoding for fast mri," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.

[49] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (grappa)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.

[50] J. Xu, S. Moeller, E. J. Auerbach, J. Strupp, S. M. Smith, D. A. Feinberg, E. Yacoub, and K. Uğurbil, "Evaluation of slice accelerations using multiband echo planar imaging at 3 t," *NeuroImage*, vol. 83, pp. 991–1001, 2013.

[51] K. Setsompop, B. A. Gagoski, J. R. Polimeni, T. Witzel, V. J. Wedeen, and L. L. Wald, "Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty," *Magnetic Resonance in Medicine*, vol. 67, no. 5, pp. 1210–1224, 2012.

[52] A. L. Alexander, J. E. Lee, M. Lazar, and A. S. Field, "Diffusion tensor imaging of the brain," *Neurotherapeutics*, vol. 4, no. 3, pp. 316–329, 2007.

[53] J.-D. Tournier, S. Mori, and A. Leemans, "Diffusion tensor imaging and beyond," *Magnetic Resonance in Medicine*, vol. 65, no. 6, pp. 1532–1556, 2011.

[54] V. Wedeen, T. Reese, D. Tuch, M. Weigel, J. Dou, R. Weiskoff, and D. Chessler, "Mapping fiber orientation spectra in cerebral white matter with fourier-transform diffusion mri," in *Proceedings of the 8th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, (Denver, Colorado, USA), p. 82, ISMRM, 2000.

[55] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly, "Direct estimation of the fiber orientation density function from diffusion-weighted mri data using spherical deconvolution," *NeuroImage*, vol. 23, no. 3, pp. 1176–1185, 2004.

[56] J.-D. Tournier, F. Calamante, and A. Connelly, "Robust determination of the fibre orientation distribution in diffusion mri: Non-negativity constrained super-resolved spherical deconvolution," *NeuroImage*, vol. 35, no. 4, pp. 1459–1472, 2007.

[57] M. Chamberland, E. P. Raven, S. Genc, K. Duffy, M. Descoteaux, G. D. Parker, C. M. Tax, and D. K. Jones, "Dimensionality reduction of diffusion mri measures for improved tractometry of the human brain," *NeuroImage*, vol. 200, pp. 89–100, 2019.

[58] T. Hendriks, A. Vilanova, and M. Chamberland, "Neural spherical harmonics for structurally coherent continuous representation of diffusion mri signal," in *Computational Diffusion Magnetic Resonance Imaging*, pp. 1–12, 2023.

[59] H.-E. Assemlal, D. Tschumperlé, and L. Brun, "Efficient and robust computation of pdf features from diffusion mr signal," *Medical Image Analysis*, vol. 13, no. 5, pp. 715–729, 2009.

[60] G. E. Andrews, R. Askey, and R. Roy *Special functions*, 1999.

[61] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, "Regularized, fast, and robust analytical q-ball imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 3, p. 497–510, 2007.

[62] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng, "Fourier features let networks learn high frequency functions in low dimensional domains," *NeurIPS*, 2020.

[63] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics*, vol. 41, pp. 102:1–102:15, July 2022.

[64] Y. Sun, J. Liu, M. Xie, B. Wohlberg, and U. S. Kamilov, "Coil: Coordinate-based internal learning for tomographic imaging," *IEEE Transactions on Computational Imaging*, vol. 7, pp. 1400–1412, 2021.

[65] R. Liu, Y. Sun, J. Zhu, L. Tian, and U. S. Kamilov, "Recovery of continuous 3D refractive index maps from discrete intensity-only measurements using neural fields," *Nature Machine Intelligence*, vol. 4, no. 9, pp. 781–791, 2022.

[66] "Nist digital library of mathematical functions, §14.30: Spherical and spheroidal harmonics," 2025.

[67] H. C. Project, "Wu-minn consortium of the nih human connectome project. 1200 subjects reference manual – appendix i," tech. rep., Human Connectome Project, 2017.

[68] J. L. Andersson, S. Skare, and J. Ashburner, "How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging," *NeuroImage*, vol. 20, p. 870–888, Oct 2003.

[69] B. Fischl, "Freesurfer," *NeuroImage*, vol. 62, p. 774–781, Aug 2012.

[70] M. Deserno, "How to generate equidistributed points on the surface of a sphere," *MPIP*, Sep 28 2004.

[71] E. Caruyer, C. Lenglet, G. Sapiro, and R. Deriche, "Design of multishell sampling schemes with uniform coverage in diffusion mri," *Magnetic Resonance in Medicine*, vol. 69, p. 1534–1540, Apr 2013.

[72] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 12 2014.

[73] C. Liao, U. Yarach, X. Cao, S. S. Iyer, N. Wang, T. H. Kim, Q. Tian, B. Bilgic, A. B. Kerr, and K. Setsompop, "High-fidelity mesoscale in-vivo diffusion mri through gslider-buda and circular epi with s-loraks reconstruction," *NeuroImage*, vol. 275, p. 120168, 2023.

[74] L. I. Sacolick, F. Wiesinger, I. Hancu, and M. W. Vogel, "B1 mapping by bloch-siegert shift," *Magnetic Resonance in Medicine*, vol. 63, p. 1315–1322, Apr 2010.

[75] M. Deprez, A. Price, D. Christiaens, G. Lockwood Estrin, L. Cordero-Grande, J. Hutter, A. Daducci, J.-D. Tournier, M. Rutherford, S. J. Counsell, M. B. Cuadra, and J. V. Hajnal, "Higher order spherical harmonics reconstruction of fetal diffusion mri with intensity correction," *IEEE Transactions on Medical Imaging*, vol. 39, no. 4, pp. 1104–1113, 2020.

[76] E. Thompson, A. Mohammadi-Nejad, E. Robinson, J. Andersson, S. Jbabdi, M. Glasser, M. Bastiani, and S. Sotiropoulos, "Non-negative data-driven

mapping of structural connections with application to the neonatal brain," *NeuroImage*, vol. 222, p. 117273, 2020.

[77] S. Wilson, M. Pietsch, L. Cordero-Grande, A. N. Price, J. Hutter, J. Xiao, L. McCabe, M. A. Rutherford, E. J. Hughes, S. J. Counsell, and et al., "Development of human white matter pathways in utero over the second and third trimester," *PNAS*, vol. 118, no. 20, 2021.