PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
SCHOOL OF ENGINEERING

UNDERSAMPLED Q-SPACE RECONSTRUCTION METHODS FOR DIFFUSION SPECTRUM IMAGING

GABRIEL ENRIQUE VARELA MATTATALL

Thesis submitted to the Office of Research and Graduate Studies
in partial fulfillment of the requirements for the degree of
Doctor of Science in Engineering

Advisor:
PABLO IRARRÁZAVAL

Santiago de Chile, July 2019

PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
SCHOOL OF ENGINEERING

UNDERSAMPLED Q-SPACE RECONSTRUCTION METHODS FOR DIFFUSION SPECTRUM IMAGING

GABRIEL ENRIQUE VARELA MATTATALL

Members of the Committee:
PABLO IRARRÁZAVAL
MARCELO ANDÍA
STÉREN CHABERT
CRISTIÁN TEJOS
TONY STÖCKER
JORGE VÁSQUEZ

Thesis submitted to the Office of Research and Graduate Studies
in partial fulfillment of the requirements for the degree of
Doctor of Science in Engineering

Santiago de Chile, July 2019

Life is not about giving punches, it's about receiving them and still fight!

ACKNOWLEDGEMENTS

First of all, I want to say special thanks to my family and friends. Without their love and support, there is absolutely no doubt that I would have not accomplished my PhD. As I always say, they are my brightest stars during the darkest night. Kind regards to my dear friends from Piscola F.C, MGT and Touch Rugby Chile; and my friends since childhood: MG, SK and LSB.

I want to say thanks to my advisor and mentor, Pablo Irarrázaval Mena, A.K.A PIM, who offered me to do the PhD and since then, has trained me through a highly intensive methodology to do research. After surviving this process with him, I feel like willing to challenge anything.

I want to say thanks to Tony Stöcker, Alexandra Koch and Rüdiger Stirnberg, from the DZNE in Germany. We have been together since the beginning of my PhD program and thanks to them and to our research collaboration I was able to explore the world and open my mind. A fundamental quality to do research.

Finally, I want to say thanks to the Biomedical Imaging Center. We had good times and enjoyed from each other the best as possible. Special thanks to the PIM team (GdM, RC and CC), Carlos Milovic, Karis Letelier, Felipe Valenzuela and Manuel Chapa. For sure Im going to miss our coffee breaks, pranks and high quality memes.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

UNDERSAMPLED Q-SPACE RECONSTRUCTION METHODS
FOR DIFFUSION SPECTRUM IMAGING

Thesis submitted to the Office of Research and Graduate Studies in partial fulfillment of the requirements for the Degree of Doctor in Engineering Sciences by

GABRIEL ENRIQUE VARELA MATTATALL

Abstract

Diffusion magnetic resonance imaging is a technique that makes the signal sensitive to water diffusion at the microstructural level. It provides an indirect mean to describe the microstructure that restricts water diffusion. One important goal of diffusion magnetic resonance imaging is to obtain the ensemble average of water spins' displacements: the so-called diffusion propagator. The diffusion propagator contains the information from the diffusion signal and it provides biomarkers which could be used for clinical applications. Nonetheless, its full acquisition is infeasible in a routine clinical acquisition. Hence, its acceleration by means of advanced acquisition and reconstruction methods based on signal modeling is a key procedure to make the diffusion propagator attainable in time-restricted scans. The complete reconstruction of the diffusion propagator has been already accelerated by means of compressed sensing and by more sophisticated signal models. However, the main hypothesis of this thesis is that the acceleration can go even further. This thesis studies and proposes different reconstruction methods to improve the trade-off between signal reconstruction and the required number of samples. Reconstruction quality and propagator-based indices are used to evaluate the reconstruction performance applied to simulated and in vivo data. Additionally, we visually inspected reconstructions from anatomical regions that are known for their highly complex microstructure.

This thesis contains three main chapters describing three different works around the reconstruction problem. The first chapter is a proposal about combining both signal modeling and compressed sensing. In the proposed method, the multi-Gaussian model provides a "low-frequency" version of the diffusion propagator; and then compressed sensing reconstructs the differences, or the "high-frequency" components, between the model and the complete diffusion data. The proposed method improves the independent results of either the multi-Gaussian model or compressed sensing at low noise levels. The second chapter is a comparison of different reconstruction methods from the state-of-the-art. The methods are tested under different reconstruction settings in terms of under-sampling, noise level and microstructure orientation. Although the most efficient reconstruction method is based on adaptive dictionaries, the reconstruction method "mean apparent propagator" is less data-dependant and is more accurate for obtaining diffusion-based indices. Finally, the third chapter is the application of a compressed sensing to the mean apparent propagator. Results indicate that the combination of mean apparent propagator and compressed sensing provides a mean to obtain the complete diffusion propagator in clinical time-restricted scans.

Members of the Doctoral Thesis Committee:
PABLO IRARRÁZAVAL
CRISTIÁN TEJOS
MARCELO ANDÍA
STÉREN CHABERT
TONY STÖCKER
JORGE VÁSQUEZ

Santiago, July, 2019

Keywords: diffusion propagator, q-space, compressed sensing, mean apparent propagator

RESUMEN

UNDERSAMPLED Q-SPACE RECONSTRUCTION METHODS
FOR DIFFUSION SPECTRUM IMAGING

Tesis enviada a la Dirección de Investigación y Postgrado en cumplimiento parcial de los requisitos para el grado de Doctor en Ciencias de la Ingeniería

GABRIEL ENRIQUE VARELA MATTATALL

Abstract

La imagen de resonancia magnética de difusión es una técnica que hace la señal sensible a la difusión del agua a nivel microestructural. Esto proporciona un medio indirecto para describir la microestructura que restringe la difusión del agua. Un objetivo importante de las imágenes de resonancia magnética de difusión es obtener el conjunto promedio de los desplazamientos de los espins de agua: el llamado propagador de difusión. El propagador de difusión contiene la información de la señal de difusión y proporciona biomarcadores que podrían utilizarse en aplicaciones clínicas. No obstante, el tiempo de adquisición es inviable en una adquisición clínica rutinaria. Por lo tanto, su aceleración por medio de métodos avanzados de adquisición y reconstrucción en base a modelamiento de señales es un procedimiento clave para hacer que la adquisición del propagador de difusión sea alcanzable en tiempos adecuados. La reconstrucción completa del propagador de difusión ya ha sido acelerada por medio de sensado comprimido y por modelos más sofisticados. Sin embargo, la principal hipótesis de esta tesis es que se pueden alcanzar aceleraciones aún mayores. Esta tesis estudia y propone diferentes métodos de reconstrucción para mejorar el compromiso entre la reconstrucción de la señal y el número requerido de muestras. Los índices sobre calidad de reconstrucción y basados en el propagador son usados para evaluar el rendimiento de la reconstrucción aplicados a simulaciones y datos in vivo. Además, inspeccionamos visualmente las reconstruciones de regiones anatómicas que son conocidas por su microestructura altamente compleja.

Esta tesis contiene tres capítulos principales que describen tres trabajos sobre el problema de reconstrucción, tanto para la reconstrucción de la totalidad del propagador de difusión como de la totalidad del espacio-q. El primer capítulo es una propuesta con el propósito de combinar tanto el modelamiento de señales, como la teoría de sensado comprimido. En el método propuesto, el modelo multi-gaussiano proporciona una versión de "baja frecuencia" del propagador de difusión y luego, sensado comprimido reconstruye las diferencias , o los componentes de "alta frecuencia", entre el modelo y los datos completos de difusión. El método propuesto mejora los resultados independientes del modelo multi-Gaussiano o de sensado comprimido a niveles bajos de ruido. El segundo capítulo es una comparación de diferentes métodos de reconstrucción del estado del arte. Estos métodos son probados en diferentes configuraciones de reconstrucción en términos de nivel de submuestreo y ruido y orientación de la microestructura. Si bien el método de reconstrucción más eficiente se basa en diccionarios adaptativos, el método de reconstrucción "propagador medio aparente" depende menos de los datos y es más preciso para obtener índices basados en la difusión. Finalmente, el tercer capítulo es la aplicación de sensado comprimido al propagador medio aparente. Los resultados indican que la combinación de ambos provee un medio para obtener el propagador de difusión en adquisiciones clínicas acotadas en tiempo.
Miembros de la Comisión de Tesis Doctoral:
PABLO IRARRÁZAVAL
CRISTIÁN TEJOS
MARCELO ANDÍA
STÉREN CHABERT
TONY STÖCKER
JORGE VÁSQUEZ

Santiago, Julio, 2019

Palabras Claves: propagador de difusión, espacio-q, sensado comprimido, propagador medio aparente

1. INTRODUCTION

Magnetic resonance imaging (MRI) is a sophisticated technique that provides the opportunity to improve the diagnosis of human diseases and the follow-up of treatments in comparison to other imaging methods. The MR signal is sensitive to different physical properties such as density, susceptibility, flow, diffusion, etc. Therefore, it can generate several kinds of images; being a very unique imaging tool when compared to the single-contrast images from ultrasound and computational axial tomography. Its imaging richness may complement and even replace other imaging methods for medical decisions. Additionally, MRI is non-invasive and radiation-free. Hence, although MRI has challenges that need to be solved for the future, it offers much more opportunities for development that must be investigated.

One particular MRI technique is diffusion MRI (dMRI) in which the signal is sensitive to water diffusion at a microstructural level. Under this regime, unreachable for other imaging methods, the data from dMRI is used to infer about the state of the microstructure in healthy and in pathological stages. In particular, the state of neurons and axons that may be dying due to neurodegenerative diseases. Neurodegenerative diseases act at a microstructural level but they are only detectable at the macrostructural level. Therefore, dMRI may be the key to develop clinical strategies to detect, diagnose and monitor the treatment of these diseases.

A complete dMRI requires an infeasible acquisition time in order to sample the complete diffusion signal and it also requires very complex models to explain the microstructure behind the acquired signal. The clinically applicable dMRI is based on a Gaussian model for the diffusion in order to obtain a coarse characterisation of the signal in time-restricted scans. Although the value from Gaussian diffusion is undeniable, given its clinical utility, a more specific and accurate characterisation of the diffusion signal could provide new and valuable information to improve medical decisions. Finally, the challenge is to wisely reduce the acquisition time and preserve the quality of the complete diffusion signal.

This thesis focuses on the process of recovering the missing samples from the diffusion

signal due to under-sampling: the reconstruction process. By means of an improved reconstruction, the complete characterisation of the diffusion propagator may be attainable in reduced time scans. For that purpose, this thesis uses under-sampling strategies and completes the missing data through different reconstruction methods that exploit a-priori knowledge of the diffusion signal. Together with proposing new reconstruction techniques, this thesis also compares existing dMRI reconstruction methods and it analyses which one provides the highest reconstruction quality from minimum sampling. This thesis consists of three chapters. Each chapter corresponds to an investigation about reconstruction performance, either by proposing a new method or by comparing different reconstruction methods.

The first chapter corresponds to the proposal of reconstructing the complete diffusion propagator by combining multiple Gaussians and compressed sensing theory. Compressed sensing is used to recover the differences between the fit from multiple Gaussians and the complete diffusion data. The proposed method performs well under certain conditions, but it is, in general, outperformed by other reconstruction methods. Although the proposal could have been improved by including recent advances or by a better tuning of the optimisation parameters, its philosophy, which is the description of Gaussian and non-Gaussian diffusion, is represented by the reconstruction method known as mean apparent propagator, which is described in the second chapter and improved in the third chapter.

The second chapter corresponds to the comparison of three representative dMRI reconstruction methods based on compressed sensing theory, signal modeling and machine learning, respectively. The work in this chapter analysed the payoffs and drawbacks from the following representative methods: diffusion spectrum imaging accelerated by compressed sensing using either 1) the identity as sparsifying transform or 2) adaptative dictionaries for sparse representation; and 3) the mean apparent propagator, a physical model based on continuous basis-functions for an efficient representation of the diffusion signal. Although the most efficient reconstruction method turned out to be adaptive dictionaries, the mean apparent propagator is less data-dependant and is more accurate for obtaining diffusion-based indices.

The third and last chapter corresponds to an improved mean apparent propagator reconstruction using compressed sensing theory. Due to the fact that this physical model goes in agreement with compressed sensing fundamental requirements, this research shows that the proposed reconstruction method obtains higher reconstruction quality and obtains a better approximation of the complete diffusion signal than its previous version. Furthermore, with the application of compressed sensing it may operate in an acquisition time similar to the acquisition time used for more simplistic models such as the diffusion tensor model.
In what remains of this chapter we present a brief description of diffusion MRI and a review of the analysed reconstruction methods.

1.1. diffusion MRI: q-space imaging and the diffusion propagator

Diffusion MRI is a technique which has the ability to make the MR signal sensitive to water diffusion. By including a diffusion encoding after excitation, the MR signal is attenuated proportionally to the diffusion phenomenon occurring in each voxel. The image from the attenuated signal due to diffusion is a diffusion-weighted image. Since the attenuation depends on the time given to the diffusion process, and on the orientation and strength of the gradient, each diffusion-weighted image provides information about the microstructural setting that attenuates the signal. A significant number of diffusion-weighted images in all directions is required in order to obtain information about the microstructure. The complete set of diffusion-weighted images provides a mean to obtain the ensemble average of displacements from a population of water molecules and for different diffusion times. This is the so-called the diffusion propagator (Callaghan, 1991, 338 p.).
The diffusion encoding provides a way to label spins' positions and then observe the distance travelled during a period of time. Any gradient waveform after excitation and before the actual acquisition may be considered as diffusion enconding, as long as the net area for that encoding is zero (Bernstein, 2004). Overall, the diffusion encoding diminishes the MR signal from a T2-weighted image, E_0, to generate a diffusion-weighted image, E_i, as,

$$E_i = E_0 e^{i\phi(x,t)}, \tag{1.1}$$

where $\phi(x,t)$ is the phase from a spin located at position x over time t,

$$\phi(x,t) = \gamma \int_0^t G(\xi)x(\xi)d\xi; \tag{1.2}$$

and where γ is the gyromagnetic ratio and G is the strength of the gradient. In particular, the standard diffusion encoding from Stejskal and Tanner, known as the pulsed gradient spin echo (PGSE) (Stejskal & Tanner, 1965) sequence, allows to do such tracking of water spins' displacements between two instances of time in a very convenient manner. This PGSE sequence provides a clear distinction between the encoding time from the gradient and the diffusion time Δ (Johansen-Berg, 2009). The encoding time is used to label the position of spins whereas the diffusion time is used to track the displacement between labels. Furthermore, thanks to the principles used in the PGSE sequence, the Fourier relationship between the diffusion propagator and the complete set of diffusion-weighted images, known as q-space, results explicit. The PGSE diffusion encoding sequence is the only diffusion encoding used across this thesis because the objective in here is to reconstruct the diffusion propagator.

The PGSE sequence is modeled by two impulses whose amplitudes are G and $-G$; and separated by a period of time Δ. Hence, the phase from a single spin results in:

$$\phi(x,t) = \gamma G \int_0^t [\delta(\xi) - \delta(\xi - \Delta)]x(\xi)d\xi. \tag{1.3}$$

Given the sifting property of the impulse function, $\delta(\xi)$, the integral turns into a evaluation of the position of the spin at the starting time, $x(0)$, and its position at $x(\Delta)$ as,

$$\phi(x,t) = \gamma G[x(0) - x(\Delta)]; \tag{1.4}$$

and the complete expression so far is,

$$E_i = E_0 e^{i\gamma G[x(0)-x(\Delta)]}. \tag{1.5}$$

The use of impulse functions to characterise the displacement of the spin is known as the narrow pulse approximation (Callaghan, 1991, 338 p.); however, it is not given in a real experiment. The existence of a finite diffusion encoding time, historically known as δ, turns

the labels into the evaluation of centre-of-mass rather than the evaluation of position for a single spin (Mitra & Halperin, 1995; Lori, Conturo, & Le Bihan, 2003). If $\delta << \Delta$, the diffusion process during the encoding time may be neglected. The net phase $\Delta\phi$, results strictly dependant to $x(0)$ and $x(\Delta)$. If $x(0) = x(\Delta)$, the net phase vanishes. Nonetheless, in most cases $x(0) \neq x(\Delta)$ and it generates an incomplete cancellation of the net phase. Now, because a voxel contains a population of spins, the phase dispersion and/or spreading of phases among the randomly moving population of spins, that started at $x(0)$ and ended up in $x(\Delta)$ for a diffusion time Δ, produces an attenuated signal due to the incoherence in the orientations of individual magnetic moments (Johansen-Berg, 2009). Let us replace E_0 with the ensemble average of spins:

$$E_i = \int_{-\infty}^{\infty} p(x(\Delta)|x(0), \Delta) e^{i\gamma G[x(0)-x(\Delta)]} dx. \tag{1.6}$$

Now, if we replace in eq. (1.6) the travelled distance with $r = x(\Delta) - x(0)$, and if we add the definition of a wave-vector $q = \frac{\gamma}{2\pi}\delta G$, then,

$$E(q, \Delta) = \int_{-\infty}^{\infty} p(r|\Delta) e^{-i2\pi qr} dr, \tag{1.7}$$

the Fourier relationship between the diffusion-weighted signal, $E(q, \Delta)$, and the probability density function of displacements from the ensemble average of water molecules, $p(r|\Delta)$ (or diffusion propagator), becomes explicit. Hence, both domains are connected by the Fourier transform. Finally, because the diffusion propagator is expected to be positive and symmetric, it is possible to obtain it from q-space data as:

$$p(r|\Delta) = \int_{-\infty}^{\infty} |E(q)| e^{i2\pi qr} dq, \tag{1.8}$$

and this is known as diffusion spectrum imaging (DSI) (Wedeen et al., 2000, 2005).

Each sample of q-space corresponds to the magnitude of a diffusion-weighted image. Therefore, a complete q-space acquisition becomes infeasible for clinical protocols. Given the fact that the diffusion propagator could reveal relevant biomarkers for early detection, diagnosis and follow-up treatments for diseases that alter the tissue microstructure, there is

a need to accelerate the acquisition of the fully-sampled diffusion propagator. Furthermore, the acquisition needs to be inside time-restricted clinical scans.

1.2. Sampling and reconstruction of q-space

The first idea to accelerate the acquisition is to reduce the number of q-space samples below the Nyquist sampling rate. Obviously, this option generates a trade-off between signal quality and acquisition time that needs to be analysed from a signal processing perspective. In particular, random sampling is used in compressed sensing theory for signal acquisition and reconstruction.

The continuous signal is sampled such that the measured signal corresponds to the product between the original signal and a train of impulses as,

$$p[r] = p(r)\text{Ш}_T(r) \quad \text{with} \quad \text{Ш}_T(r) = \sum_{n=-\infty}^{\infty} \delta(r - nT), \tag{1.9}$$

where Ш_T is a train of impulses with a sampling period T. By Fourier properties, the product of functions in the image or Fourier domain corresponds to a convolution in the opposite domain (Irarrazaval, 2000). Furthermore, the Fourier transform of the Ш_T function is another function of the same style, Ш_f, but with a "sampling" period $f = T^{-1}$. Therefore, if we are sampling in q-space, the diffusion propagator is convolving its signal with a train of impulses which have a separation between impulses defined by $f = T^{-1}$. If the sampling period T increases as we try to under-sample the q-space, f decreases and the excessive under-sampling reaches a point where the convolutions between the diffusion propagator and the infinite impulses from Ш_f start to overlap. The overlap between convolutions is known as the aliasing artifact. This artifact is very troublesome because you can not distinguish between the overlapped convolutions. This is why a minimum sampling frequency is demanded to assure enough separation between convolutions and it is known as the Nyquist Sampling Criterion. It states that a signal needs to be acquired with at least twice the greatest frequency observed in the signal. Otherwise, aliasing appears and which is impossible to resolve, unless the acquisition system has enough acquisition devices to

disentangle the spatial origin of overlapped signals, as in parallel MRI (Lipton., 2008). In order to avoid the Nyquist Sampling Criterion, two main approaches are analysed in this thesis: 1) the application of a model to exploit directly the a-priori knowledge about the diffusion signal; and 2) the application of compressed sensing to exploit the sparse representation of the diffusion signal.

1.2.1. Models for q-space reconstruction

This section gives a brief explanation about the most relevant reconstruction methods that are analysed across the different chapters of this thesis. They are separated in two groups depending on whether the reconstruction method characterises diffusion as a Gaussian process (diffusion tensor imaging) or as a non-Gaussian process (mean apparent propagator and compressed sensing).

Diffusion Tensor Imaging (DTI)

The Gaussian model significantly speeds up the acquisition and the reconstruction process; and it is possible to obtain the diffusion tensor that provides tensor-based indices for clinical applications. Finally, the Fourier transform of a Gaussian function is another Gaussian function; hence, it is possible to analyse a "low-frequency" version of the diffusion propagator through Gaussian diffusion.
Diffusion Tensor Imaging (DTI) (Le Bihan et al., 2001) is probably the most common and standarised procedure to characterise the diffusion signal. In it,

$$E(\mathbf{q},\Delta) = E(\mathbf{q}|_{|\mathbf{q}|=0},\Delta)\exp\left(-4\pi^2\mathbf{q}\mathbf{D}\mathbf{q}^{\mathrm{T}}\right), \tag{1.10}$$

the diffusion signal is based on the diffusion tensor, $\mathbf{D}$. This tensor is a 3×3 covariance matrix; therefore, it is square and symmetric with six unknowns to estimate:

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix}. \tag{1.11}$$

The diffusion tensor is commonly decomposed with the eigendecomposition to obtain the eigenvalues and eigenvectors:

$$\mathbf{D} = \boldsymbol{R}\Lambda\boldsymbol{R}^{\mathrm{T}} = \begin{bmatrix} e_{1x} & e_{2x} & e_{3x} \\ e_{1y} & e_{2y} & e_{3y} \\ e_{1z} & e_{2z} & e_{3z} \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} \begin{bmatrix} e_{1x} & e_{2x} & e_{3x} \\ e_{1y} & e_{2y} & e_{3y} \\ e_{1z} & e_{2z} & e_{3z} \end{bmatrix}^{\mathrm{T}}; \tag{1.12}$$

that are used to obtain tensor-based indices for clinical applications. For example, the eigenvalues obtained from DTI may be used to compute the mean diffusivity (MD) and the fractional anisotropy (FA).

Mean Apparent Propagator (MAP)

The non-Gaussian diffusion component cannot be approximated by the Gaussian model. This non-Gaussian component appears: 1) at high q-space values; and 2) with complex microstructures such as crossing-fibres in the human brain. Consequently, the most difficult and challenging area in dMRI is the correct estimation of non-Gaussian diffusion.

Mean apparent propagator (MAP) (Ozarslan et al., 2013; Avram et al., 2016) corresponds to a physical model based on continuous basis-functions. The 1D basis-function, $\phi_n(\mu, q)$, is defined as,

$$\phi_n(\mu, q) = \frac{1}{i^n\sqrt{2^n n!}} \exp\left(-\frac{(2\pi\mu q)^2}{2}\right) H_n(2\pi\mu q), \tag{1.13}$$

where μ is an scaling parameter and H_n corresponds to the nth-order Hermite polynomial which is modulated by a Gaussian term. What is really interesting about this set of basis-functions, is that the first basis-function ($n = 0$) is the Gaussian diffusion whereas the rest of the basis-functions ($n > 0$) are the non-Gaussian diffusion. A second interesting property of $\phi_n(\mu, q)$ is that its Fourier transform also results in a Gauss-Hermite function, $\psi_n(\mu, r)$,

$$\psi_n(\mu, r) = \frac{1}{\mu\sqrt{2^{n+1}\pi n!}} \exp\left(-0.5\left(\frac{r}{\mu}\right)^2\right) H_n\left(\frac{r}{\mu}\right); \tag{1.14}$$

and this is the 1D basis-function that characterises the diffusion propagator. The reconstruction process for this method is the estimation of the coefficients, c, that weight the set of basis-functions. Furthermore, both sets of basis-functions, Φ and Ψ, share the same

coefficients, so it is possible to optimise data consistency from q-space samples, $\mathbf{E}$, while imposing properties of the diffusion propagator like non-negativity as:

$$\hat{\mathbf{c}} = \text{argmin}_{\mathbf{c}} ||\Phi^{\mathrm{T}}\mathbf{c} - \mathbf{E}||_2^2 \quad \text{s.t} \quad \Psi^{\mathrm{T}}\mathbf{c} \geq 0. \tag{1.15}$$

Reconstruction by means of compressed sensing theory

Compressed sensing (CS) (Candes & Romberg, 2007) is an efficient sampling and reconstruction strategy. Before CS, the Nyquist sampling criterion was a fundamental requirement for sampling technology and it could only be avoided under certain circumstances. Thanks to CS, now signals can be under-sampled below the Nyquist criterion and obtain an almost exact version of the fully-sampled signal.

If the three principles of randomness, sparsity and incoherence are fulfilled the under-sampled signal can be completely reconstructed with CS. Let us start with the principle of Randomness. As stated in the previous section, the conventional uniform under-sampling makes those infinite convolutions to overlap and make the aliasing artifact almost impossible to disentangle. Now, if random under-sampling is applied, it would not stop the overlapping between these convolutions; however, the overlapping will be completely disorganised and aliasing turns into a noise-like artifact. Now, the reconstruction problem from random under-sampling turns into a denoising problem, where this noise-like artifact is just energy that we need to re-distribute along the field-of-view of the original signal. From the infinite set of solutions to solve the ill-posed problem, we will pursue the solution that provides the sparsest representation of the signal in a certain domain. For example, if our time-dependant signal corresponds to a single cosine waveform with frequency f_1: we would have maximum sparsity of the signal in the Fourier domain because it results in only two coefficients different from zero; and correspond to f_1 and $-f_1$. Under this concept, we do not need to acquire all samples from the sampling space (time) if in the representation domain (Fourier) the information is already compressed in just a few coefficients (in this case two coefficients only). However, we need to wisely select the correct samples that compress the information and that allows to discard "irrelevant" samples. Finally, we need

a strategy to find that sparse solution for our problem and that is through the incoherence principle. Let us take for example the definition of the Fourier Transform from eq. (1.7). By definition, incoherence is the pattern-free effect of one sample in the other space; therefore, the image and Fourier domains show maximum incoherence between them because a single sample contributes to the complete signal in the other domain and vice versa. By means of incoherence, the optimisation pursues a sparse representation of the signal in the sparsifying domain whereas incoherence penalises incorrect results in the data consistency term. Therefore, the compressed sensing optimisation may be expressed as:

$$\text{argmin}_p \quad \|SFp - E\|_2^2 + \lambda\|\Omega p\|_1, \tag{1.16}$$

where Ω stands for the sparsifying domain and F stands for the Fourier operator (given the acquisition domain that is fixed to Fourier/q-space). S is the under-sampling operator, p is the diffusion propagator and E are the q-space samples. Finally, this two terms are balanced by the λ parameter. This optimisation may be written as a basis-pursuit optimisation, or as a lasso optimisation.

From the three principles, the most relevant principle is sparsity; and therefore, the most practical and common way to work with CS is by the selection of the most appropiate sparsifying basis for the application of CS. Diffusion Spectrum Imaging (DSI), eq. (1.8), is generally improved by means of compressed sensing as in (Menzel, Tan, Sperl, & King, 2011). The main idea in this line of research is to find the most sparsifiyng basis for the sparse representation of the diffusion propagator. This sparsifying basis may come from traditional sparsifying transforms like the identity-basis (S. L. Merlet & Deriche, 2013), Total Variation (Menzel et al., 2011), wavelets (Paquette, Merlet, Gilbert, Deriche, & Descoteaux, 2015), or from pre-defined dictionaries as sparsifying bases (Bilgic et al., 2012). This could be done by creating a sparse dictionary, from the information of thousands of diffusion propagators; or by creating an sparsifying basis that takes into account the a-priori knowledge from the diffusion signal. Thanks to CS, the diffusion propagator can be recovered with approximately one-fourth of the complete data set and it reduces the acquisition

time from tenths of minutes to a couple of them.

1.2.2. Summary of the research

Problem

The direct reconstruction of the diffusion propagator is, most of the time, infeasible for routinely clinical applications. Hence, the main problem tackled in this work is how to preserve high quality reconstruction of the complete diffusion propagator while the sampling protocol is modified to reduce the scan time.

Hypothesis

The hypothesis is that, by means of compressed sensing theory and the selection of the best sparsifying basis, it is possible to obtain a better reconstruction of the complete diffusion propagator from minimum sampling.

Objectives

The objectives are:

(i) To analyse different reconstruction methods from the state-of-the-art.

(ii) To propose strategies to improve the reconstruction methods from the state-of-the-art by means of compressed sensing theory.

(iii) To search for the best sparsifying basis for the application of compressed sensing.

Methods

Experiments are performed on simulations and in-vivo data. The simulations are done using Monte-Carlo simulations from the Camino software (Cook, Bai, Nedjati-Gilani, Seunarine, & Hall, 2006); and the in-vivo data comes from a collaboration with the German Center for Neurodegenerative Diseases (DZNE), in Bonn, Germany. In simulations, experiments are modified in terms of under-sampling, noise level and microstructure orientation in order to provide a thoroughly analysis from a known ground truth. In vivo data are used to show the direct applicability of results from simulations in real acquisitions.

The reconstructed diffusion propagators are evaluated in terms of reconstruction performance and in the extraction of propagator-based diffusion indices. The normalised mean squared error and Pearson's correlation coefficient indices are used for the quantitative evaluation of reconstruction performance; and the mean squared displacement and the return to zero probability indices are used for the quantitative evaluation of the extraction of propagator-based indices. Additionally, visual inspection from anatomical regions are used to observe the behaviour of the reconstructed diffusion propagators. The anatomical region of preference is the centrum semiovale because it contains many fibre-bundles with different orientations; therefore, it is a complex microstructure suitable for testing more sophisticated reconstruction methods.

Scientific Contribution

First, this thesis tries to obtain an original procedure to reconstruct the diffusion propagator from the minimum sampling possible. For that purpose, a thorough analysis from state-of-the-art reconstruction methods is performed in order to detect and improve these methods by means of compressed sensing theory. Finally, the improved method, at the end, resulted in the application of compressed sensing for the sparse fit of continuous basis-functions to q-space data.

This thesis generated two articles: one was accepted and another under revision. In this thesis nine scientific abstracts were accepted at annual meetings of the international society of magnetic resonance in medicine (ISMRM), at the international students congress from the IEEE-EMBS and at the biomedical engineering congress from the University of Concepcion.

2. COMBINED MODEL AND COMPRESSED SENSING RECONSTRUCTION FOR DIFFUSION SPECTRUM IMAGING

This Manuscript was rejected, but invited to re-submit a revised version of it by the journal Magnetic Resonance in Medicine (Magn Reson Med) in May, 2017.

2.1. Introduction

Diffusion, at a microscopic level, consists of the permanent and random displacements of water molecules. The set of possible displacements, $\mathbf{r}$, has a probability density function (pdf) which changes with time and with respect to the surrounding microstructure. The pdf of displacements for an average group of molecules given a diffusion time Δ, is the so-called diffusion propagator $p(\mathbf{r}|\Delta)$ (Callaghan, 1991, 338 p.). If the environment is free of constraints, then $p(\mathbf{r}|\Delta)$ is a Gaussian pdf. The pdf deviates from Gaussian when water molecules encounter physical constraints. Therefore, characterizing the deviation from a Gaussian pdf may provide valuable information about the underlying microstructure that restricts diffusion.

Diffusion MRI is a technique used to acquire the diffusion propagator for each voxel. It makes T2-weighted images sensitive to diffusion by including diffusion encoding gradients, with amplitude $|\mathbf{G}|$ and duration δ, separated by the previously mentioned diffusion time Δ. Thus, these images contain the effect of diffusion over Δ in the direction of $\mathbf{G}$ (Stejskal & Tanner, 1965). The data are acquired in the so-called q-space, $\mathbf{q} = \delta\mathbf{G}\gamma/2\pi$, which corresponds to the Fourier transform of the diffusion propagator. The direct reconstruction of the diffusion propagator from the magnitude of q-space data is called Diffusion Spectrum Imaging (DSI) (Wedeen et al., 2000, 2005). The diffusion propagator is defined as:

$$p(\mathbf{r}|\Delta) = \int_{-\infty}^{\infty} P(\mathbf{q}, \Delta) e^{i2\pi\mathbf{qr}} d\mathbf{q} \tag{2.1}$$

where $P(\mathbf{q}, \Delta)$ are the q-space samples. When q-space data are acquired under the Narrow Pulse Approximation ($\delta \rightarrow 0$) (Callaghan, 1991, 338 p.), the diffusion propagator is the direct representation of water molecule displacements. When the gradients duration are longer, ($\delta >> 0$), the diffusion propagator has information from the spins' centers of

mass (Mitra & Halperin, 1995; Lori et al., 2003). To fulfill the Nyquist sampling requirement, many q-space samples are needed and since each of them is an image, DSI becomes a time-expensive technique.

Undersampling methods like Compressed Sensing (CS) (Candes & Romberg, 2007) have been used to reduce the acquisition time in DSI below the Nyquist sampling rate. To do this, CS exploits the fact that data can be compressed in some domain. The first publication of CS in DSI (Menzel et al., 2011) proposed the use of Total Variation and Haar Wavelets as sparse domains. Soon after, other authors (Bilgic et al., 2012) proposed a sparse domain based on adaptive dictionaries. The pdf in its own space (the identity) has shown to be sparse enough for CS reconstruction (S. L. Merlet & Deriche, 2013). Recently, (Paquette et al., 2015) proposed a joint improvement by selecting the Cohen-Daubechies-Feauveau 9/7 Wavelet as sparse domain, and a sparse undersampling pattern with uniform angular distribution and randomly allocated samples in the radial profile. Research in this area is focused on the selection of both, the sampling pattern and the sparse domain.

When the full diffusion propagator is not needed, model-based estimations are enough to compute its most relevant properties. Generally, the propagator is approximated by a Gaussian model. Although the diffusion signal has a close connection with the Gaussian model because of the Brownian motion, a diffusion propagator based on it may not describe correctly, or completely, the underlying microstructure. For example, Diffusion Tensor Imaging (DTI) (Le Bihan et al., 2001) will fail when the voxel contains two or more non-parallel fibers. The multi-tensor model is an extension of DTI (Tuch et al., 2002), but it increases the number of parameters depending on the number of tensors used in the analysis. Another model technique is Diffusion Kurtosis Imaging (DKI) (Jensen, Helpern, & Ramani, 2005) which incorporates the non-Gaussian information of the diffusion signal as the kurtosis parameter. Finding the kurtosis increases the precision of the fit, but it is sensitive to noise; hence, DKI can only be used in high SNR conditions. Model fitting in diffusion MRI is normally used with a reduced number of parameters due to limitations in acquisition time and signal-to-noise ratio.

In this work, we propose to fit a model to DSI data with the use of CS to reconstruct the

complete diffusion propagator. The model represents the Gaussian part of the signal, whilst CS reconstructs the differences between the model and the DSI data. To do this, we firstly fit a model to randomly acquired samples in q-space, then we apply CS to complete the estimated diffusion propagator. Our hypothesis is that the combination of a model-based and CS will improve the reconstruction quality when compared to a strictly model-based method or compared to the conventional DSI using CS for a given undersampling rate.

Methods

Our proposed method is a two-step procedure. First, a model estimates the coarse part of the diffusion propagator from randomly distributed samples in q-space. Second, compressed sensing estimates the differences between the model and the data. In this work, the chosen model is the multi-tensor model in combination with compressed sensing to reconstruct the deviation from multiple Gaussians.

The combined model compressed sensing (MCS) method is compared to the multi-tensor model (M) and to compressed sensing (CS) alone. The three reconstruction methods (M, CS and MCS) are tested with retrospectively undersampled measurements which were obtained from simulated data (Monte Carlo simulations) and from actual data acquired from a diffusion phantom and an in-vivo brain. Results from the reconstructions were evaluated using the fully sampled data as ground truth. In section 3.1, the implementation of the three reconstruction methods is explained. In 3.2, the different data sets are described. And in 3.3, the quality indexes are defined.

Reconstruction

In what follows, we describe in more detail the multi-tensor model, compressed sensing and combined model compressed sensing:

1) Multi-Tensor Model (M):

We denote M as the method that obtains the diffusion propagator $\hat{m}(\mathbf{r}|\Delta)$ from the inverse

Fourier transform of the model defined in q-space:

$$\hat{m}(\mathbf{r}|\Delta) = F^{-1}\{M_{\mathbf{a}}(\mathbf{q},\Delta)\} \tag{2.2}$$

where $M_{\mathbf{a}}(\mathbf{q},\Delta) = \sum_{i=1}^{f} \alpha_i \text{exp}(-4\pi^2(\Delta - \delta/3)\mathbf{q}^T\mathbf{D}_i\mathbf{q})$ is the model fitted to the data with:

$$\hat{\mathbf{a}} = \text{argmin}_{\mathbf{a}} \|M_{\mathbf{a}}(\mathbf{q},\Delta) - P(\mathbf{q},\Delta)\|_2^2$$

where f is the number of tensors and $\mathbf{a}$ is a vector with parameters. For each tensor $\mathbf{D}_i$, $\mathbf{a}$ has entries $[\alpha_i, \lambda_i^1, \lambda_i^2, \lambda_i^3,\ \theta_i^x, \theta_i^y, \theta_i^z]$, where α_i is the fraction asociated to that tensor, $\lambda_i^{\{1,\,2,\,3\}}$ are the eigenvalues and $\theta_i^{\{x,\,y,\,z\}}$ are the directions of the eigenvectors such that $\mathbf{D}_i = Q(\theta_i^x, \theta_i^y, \theta_i^z)\ \Lambda(\lambda_i^1, \lambda_i^2, \lambda_i^3)\ Q(\theta_i^x, \theta_i^y, \theta_i^z)^{\mathrm{T}}$.

In order to estimate these parameters, we used a non-linear least-squares optimization algorithm with constrained values for $\lambda_i^{x,y,z} \in [0,\infty]$, $\theta_i^x \in [0,\pi]$ and $\theta_i^{y,z} \in [-\pi,\pi]$. The best fit of the multi-tensor model it is not necessarily a good estimation when the underlying distribution is not Gaussian-based; hovewer, we are only interested in an approximation such that the differences with the actual data will become sparser. A discussion on models and their validity in diffusion MRI can be found in (Hagmann, Jonasson, & Maeder, 2006; Ferizi et al., 2014).

2) Compressed Sensing (CS):

We denote CS as the method that estimates the diffusion propagator using compressed sensing:

$$\hat{p}(\mathbf{r}|\Delta) = \text{argmin}_{p(\mathbf{r}|\Delta)}\ \frac{1}{2}||SFp(\mathbf{r}|\Delta) - P(\mathbf{q},\Delta)||_2^2 + \gamma||\Psi p(\mathbf{r}|\Delta)||_1 \tag{2.3}$$

where the estimated diffusion propagator $\hat{p}(\mathbf{r}|\Delta)$ is obtained by minimizing the combination of a data consistency term, $0.5||SFp(\mathbf{r}|\Delta) - P(\mathbf{q},\Delta)||_2^2$, and a regularizer term, $||\Psi p(\mathbf{r}|\Delta)||_1$. Here SF is the undersampled Fourier operator. Data consistency assures the similarity between the estimation and the acquired data. The regularization constrains $p(\mathbf{r}|\Delta)$ to be represented as a sparse signal in the domain Ψ. The optimization balances by a weighting parameter γ.

For Ψ we used the Identity, Total Variation, the Cohen-Daubechies-Feauveau (CDF) 9/7 Wavelet and the Daubechies D8 Wavelet. The CDF 9/7 Wavelet toolbox of (Getreur, 2006) was modified to be used in 3D data. The Daubechies D8 Wavelet was implemented with the toolbox provided in (Baraniuk, Choi, Neelamani, & Ribeiro, 1993). The optimization was done using a homemade non-linear conjugate gradient algorithm. The weight γ was experimentally fixed and the undersampling pattern was randomly generated.

3) Combined Model and Compressed Sensing Reconstruction (MCS):

We denote MCS as the combination that uses the method in Eq.[2.2] to obtain the coarse structure of the pdf, $\hat{m}(\mathbf{r}|\Delta)$, and uses CS to estimate the differences, $\hat{d}(\mathbf{r}|\Delta)$, not captured by the model.

$$\hat{p}(\mathbf{r}|\Delta) = \hat{m}(\mathbf{r}|\Delta) + \hat{d}(\mathbf{r}|\Delta) \tag{2.4}$$

Since we can fit the model to arbitrarily located samples, we used randomly distributed positions in q-space so we could apply CS as the second step:

$$\hat{d}(\mathbf{r}|\Delta) = \text{argmin}_{d(\mathbf{r}|\Delta)} \, \|SFd(\mathbf{r}|\Delta) - D(\mathbf{q},\Delta)\|_2^2 + \gamma\|\Psi d(\mathbf{r}|\Delta)\|_1 \tag{2.5}$$

where $D(\mathbf{q},\Delta)$ are the differences between the data and the multi-tensor model. The sparse domains for Ψ are the same as in 2).

Data Sets

The simulated data were generated from Monte Carlo simulations using Camino software (Cook et al., 2006; Hall & Alexander, 2009). Simulations of one and two 60-degree crossing fibers were done with 10^3 time steps for 10^5 spins to diffuse in their respective substrates. We added Gaussian noise to each channel of the simulated measurements to obtain the magnitude of the diffusion signal as $P(\mathbf{q})_i = \sqrt{(P_{\text{simulation}}(\mathbf{q}) + \epsilon_i^1)^2 + (\epsilon_i^2)^2}$, where $\epsilon_i^1, \epsilon_i^2 \sim N(0,\sigma)$ and standard deviation $\sigma = 1/$ SNR (Paquette et al., 2015). For our simulations, we used $\sigma \in [0,1,2,3,4,5,6]\%$. Each noisy experiment was repeated one hundred times.

Phantom and in-vivo data were acquired in a 3T Siemens Prisma scanner with $G_{\text{max}} = 80$ mT/m and $SR_{\text{max}} = 200$ T/ms. For the 60-degree diffusion phantom (Moussavi-Biugui,

Stieltjes, & Fritzsche, 2011), manufactured by HQ Imaging, we acquired the images using a diffusion encoding with $\delta = 31$ ms, and $\Delta = 39$ ms, where either $q_{\text{max}} = 73.64$ mm^{-1} or $b_{\text{max}} = 8350$ s/mm^2, in a sequence with TE $= 139$ ms, TR $= 2$ s and resolution of 2.0x2.0x3.0 mm^3. From the diffusion phantom data, we selected two regions of interest in the phantom: One region with crossing fibers, and the other with a set of one fiber bundle (as shown on the left of Figure 2.1). For each region, we manually selected 20 voxels that shared structural similarity.

For the in-vivo acquisition, we used a diffusion encoding with $\delta = 31$ ms and $\Delta = 43.2$ ms, where either $q_{\text{max}} = 62.21$ mm^{-1} or $b_{\text{max}} = 6600$ s/mm^2, in a sequence with TE $= 88$ ms, TR $= 4.5$ s and isotropic resolution of 8.0 mm^3. We selected a region of interest from corpus callosum (as shown on the right of Figure 2.1 to test the three reconstruction methods. The acquisition was done after written informed consent approved by the Ethics Committee.

Simulations, phantom and in-vivo data acquisitions were done in an 11x11x11 q-space Cartesian grid. Each complete data set has 257 q-space measurements contained in a discrete half-sphere with a radius of 5 samples. The other half was obtained by symmetry since the q-space data come from the magnitude of the signal. To reduce the ringing artifacts due to insufficient acquisition, in all data sets, we multiplied the q-space data by a cosine window $W(\mathbf{q}) = \cos\left(0.5\pi\sqrt{(q_x/\|\mathbf{q}\|)^2 + (q_y/\|\mathbf{q}\|)^2 + (q_z/\|\mathbf{q}\|)^2}\right)$. We noticed that this window offers a better trade-off between artifacts and resolution than the Hanning window ($\cos^2$ window) used in the original work of DSI (Wedeen et al., 2005). It preserves more of the higher frequencies, reducing the blurring from the Hanning window without noticing ringing artifacts. This issue was also tackled in (Lacerda et al., 2016; Paquette, Gilbert, & Descoteaux, 2016), but we changed the window instead of removing the windowing. The ground truth for each data set corresponds to the inverse Fourier transform of the q-space data multiplied by the cosine window. In simulations, the q-space data are the noise-free measurements, $P_{\text{simulation}}(\mathbf{q})$, for one and two fibers.

From the fully sampled data, we used retrospective UnderSampling factors (US factors)

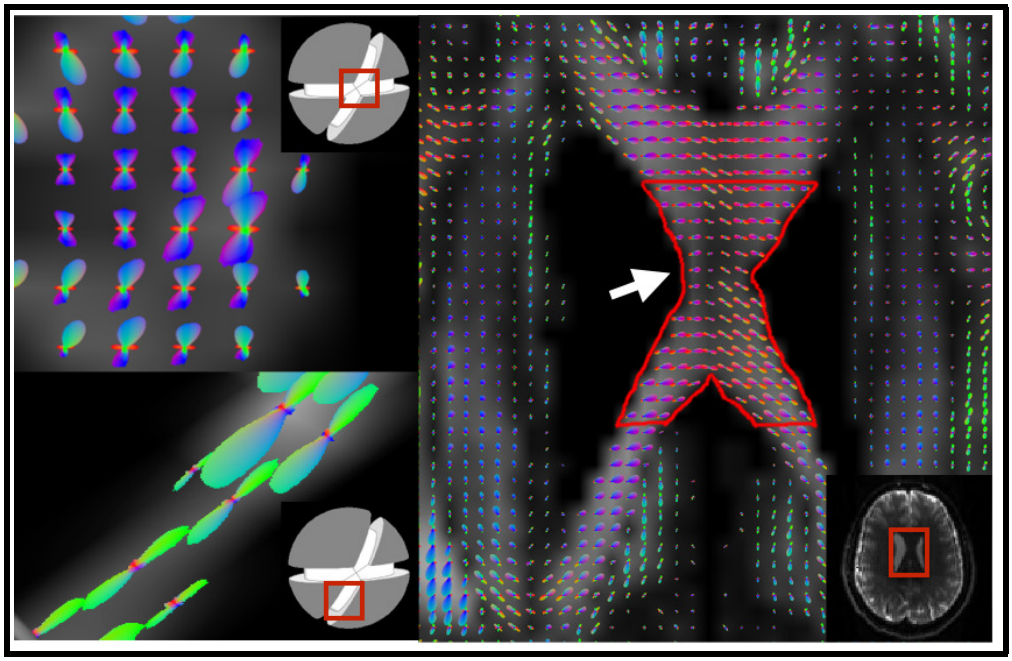

FIGURE 2.1. Visualization of the phantom and in-vivo data. Left side: Field of Oriented Distribution Functions (ODF) from two regions of interest in the diffusion phantom. Right side: Brain image depicting the position of the chosen region of interest from a slice of corpus callosum and their ODF.

from 2x to 8x for the reconstruction process of the different methods. The sampling pattern is a 3D Cartesian q-space grid where the central 3x3x3 measurements were always sampled and the rest of the measurements, depending on the US factor, were picked randomly with a linearly decreasing density distribution. This sampling pattern was used in the three reconstruction methods.

Additionally, we used a second sampling pattern which contains all the samples in a single-shell (located as in (24)) at $R_{\text{shell}} = 3/2\ q_{\text{max}}$, with the same number of measurements for each US factor as the Cartesian sampling pattern. We used a tri-linear interpolation from the Cartesian data to obtain the measurements for the spherical sampling pattern. This sampling pattern allowed us to compare the model fit from a random sampling pattern with a high angular resolution sampling pattern.

Quality Indexes

We treated the reconstruction and the ground truth as a probability density function (pdf). Hence, negative values were clipped (23) and later it was enforced that $\int p(r|\Delta) = 1$.

To quantitatively evaluate the reconstruction $\hat{p}(\mathbf{r}|\Delta)$ against the ground truth $p(\mathbf{r}|\Delta)$, we mainly used the Normalized-Root-Mean-Square-Error (NRMSE). Although, we also computed other indexes: Mean-Squared-Displacement (MSD) (Wu, Field, & Alexander, 2008), return-to-zero-probability (p0) (Assaf, Mayk, & Cohen, 2000) and Pearson's Correlation coefficient (PC). These indexes were chosen because they make sense as a way for comparing pdfs.

We defined the NRMSE as the RMSE between the reconstruction and the ground truth normalized by the maximum value in the pdf. So NRMSE $=$ RMSE$(\hat{p}(\mathbf{r}|\Delta), p(\mathbf{r}|\Delta))/\max(p(\mathbf{r}|\Delta))$. The MSD, second moment of the pdf, was defined as:

$$\text{MSD}(p(\mathbf{r}|\Delta)) = \int_{-\infty}^{\infty} p(\mathbf{r}|\Delta)\mathbf{r}^2 d\mathbf{r} \tag{2.6}$$

The relative MSD is ΔMSD $=$ (MSD$(\hat{p}(\mathbf{r}|\Delta) -$ MSD$(p(\mathbf{r}|\Delta))$)/MSD$(p(\mathbf{r}|\Delta))$. Positive ΔMSD indicates that the reconstruction overestimates the Mean-Squared-Displacement.

The return-to-zero-probability (p0) was defined as:

$$\text{p0} \equiv p(\mathbf{r} = 0 \mid \Delta) = \int_{-\infty}^{\infty} P(\mathbf{q}, \Delta) d\mathbf{q} \tag{2.7}$$

The relative return-to-zero-probability is Δp0 $= (\hat{\text{p0}} - \text{p0})/\text{p0}$, where $\hat{\text{p0}}$ is the return-to-zero-probability from the reconstruction. A positive difference indicates an overestimation of p0.

The Pearson's correlation coefficient was defined as

$$\text{PC}(p(\mathbf{r}|\Delta), \hat{p}(\mathbf{r}|\Delta)) = \frac{1}{N-1} \mathbf{Z}_{p(\mathbf{r}|\Delta)} \, \mathbf{Z}^{\mathbf{T}}_{\hat{p}(\mathbf{r}|\Delta)} \tag{2.8}$$

where $\mathbf{Z}_{p(\mathbf{r}|\Delta)} = \frac{p(\mathbf{r}|\Delta) - \text{mean}(p(\mathbf{r}|\Delta))}{\text{std}(p(\mathbf{r}|\Delta))}$ and N is the number of samples. We used the difference with 1, $(1 - \text{PC}(p(\mathbf{r}|\Delta), \hat{p}(\mathbf{r}|\Delta)))$, such that this index was zero for a perfect match.

Additionally, we visually inspected the results comparing the reconstructed pdfs and ODFs with those from the ground truth. The visualization of the ODFs and the ODF field (like in Figure 2.1) was done in the DSI Studio software (http://dsi-studio.labsolver.org).

Results

In this section we present the results from the M, CS and MCS reconstruction methods applied to simulations, diffusion phantom and in-vivo data. For M, we show the results using the random sampling pattern, and we show the fit when the sampling is spherical (denoted as Ms). For CS, we show the results of CS-I ($\Psi = \text{I}$ or Identity), and for MCS, we show the results of MCS-I. For both reconstruction methods, we used the same random sampling pattern.

We start by explaining why we chose $\Psi = \text{I}$ over other domains. Second, we show the results from simulations: for the correct model; for an incorrect model; and for different levels of noise. Third, we show the results from the diffusion phantom data. And fourth, we show the results from the in-vivo acquisition.

Sparsifying Transform

Figure 2.2 shows the quality indexes for CS and the MCS reconstructions with the different sparse domains used for the reconstruction of two-crossing fibers with noise standard deviation $\sigma = 6\%$ (SNR $= 16.66$) and undersampling (US) factor 8x. We chose $\Psi = \text{I}$ for CS because it produced the best results for the four quality indexes as shown in (A). For MCS (B), $\Psi = \text{I}$ was selected because it produced results very similar to the other domains, and it is significantly simpler. This behavior of CS and MCS with $\Psi = \text{I}$ were observed for all US factors and noise standard deviations.

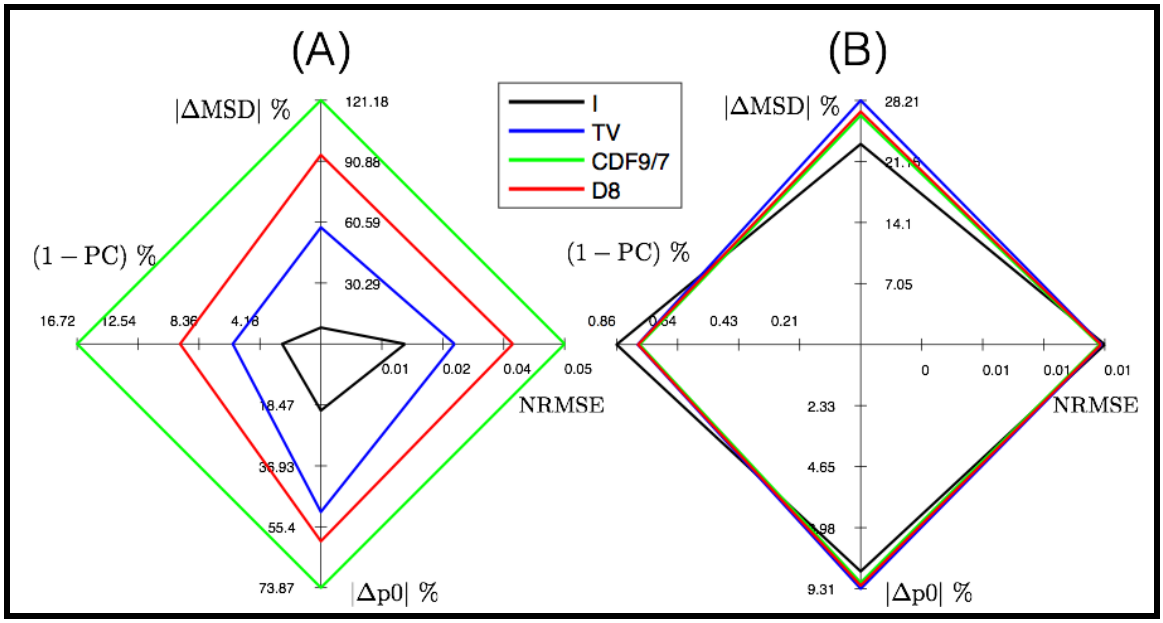

FIGURE 2.2. CS and MCS reconstructions using all sparse domains at 8x. These plots show in each axis one index, where a smaller index means a better approximation to the ground truth. A) CS-Ψ. B) MCS-Ψ. From this it was possible to conclude that $\Psi = I$ is a good sparse domain for both reconstruction methods.

Simulations

We used simulations to test the reconstruction methods. The ground truth is the noise-free fully sampled simulation. We simulated one and two fibers and for the reconstruction we used the model with one and two tensors.

Figure 2.3 shows the NRMSE for one fiber and two fibers (column (A) and (B)). When the model had the correct number of fibers, the proposed fitting with random samples (M) worked very well, outperforming the fitting to spherical samples (Ms). Here MCS performed equally or slightly better. Similar behaviors were witnessed using other indexes of comparison. For instance, at 8x undersampling for one fiber the ΔMSD measurement ranked first for both M and MCS with -7.47%; followed by, Ms with -17.53%; and finally, CS with -31.56%. PC indicated that all reconstructions had correlations above 99%. The Δp0 measurement ranked first for both M and MCS with -1.4%; followed by, Ms with 2%; and finally, CS with -3.5%. At 8x undersampling for two fibers, the ΔMSD ranked first MCS with -0.001%; followed by, M with -3.4%; CS with -14.46%; and finally, Ms

with -17.7%. PC indicated that all reconstruction methods had correlations above 98%. The Δp0 ranked first Ms with 3%; followed by, M with -4.48%; MCS with -4.84%; and finally, CS with -10.70%.

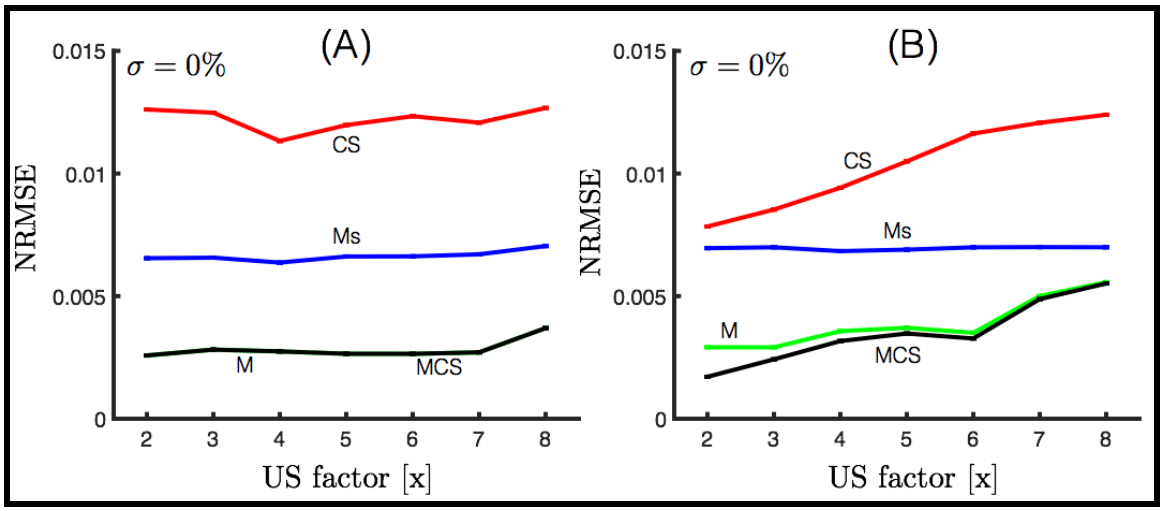

FIGURE 2.3. Results from simulations with the correct model. Column (A) is the NRMSE obtained from the noise-free reconstructions of one fiber at different undersampling factors. Column (B) is the result from the noise-free reconstructions of two-crossing fibers at different undersampling factors.

The following results show the behavior of the methods when the number of fibers in the model do not match the object. Figure 2.4 shows in column (A) the reconstruction of one fiber using $f = 2$ in the multi-tensor model, and in column (B) the reconstruction of two-crossing fibers using $f = 1$. From (B), it can be appreciated that the MCS method had the best NRMSE for all undersampling factors and, as expected, M did not perform well. Furthermore, the second stage in MCS was capable of correcting the reconstruction from the badly fitted model. As for the other indexes, at 8x undersampling, the ΔMSD ranked first MCS with 1.05%; followed by, CS with -14.46%; M with -23.58%; and finally, Ms with -40.18%. PC indicated that all reconstruction methods had correlations above 97%. The Δp0 ranked first Ms with 9.64%; closely followed by, CS with -10.70%; M with -11.32%; and finally, MCS with -13.79%.

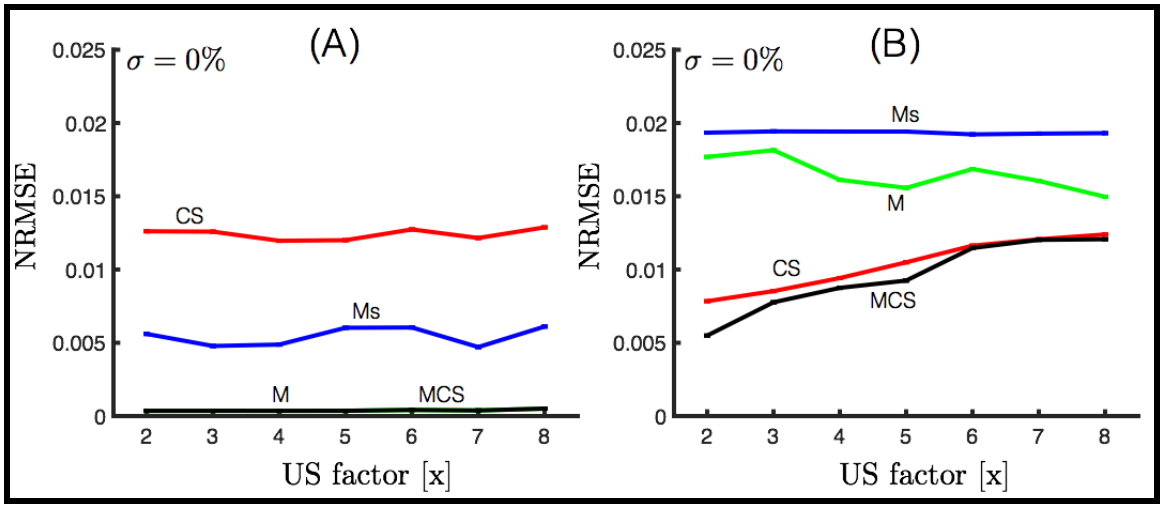

FIGURE 2.4. Results from simulations with an incorrect model. A) NRMSE obtained from noise-free reconstructions of one fiber, but using $f = 2$ in the multi-tensor model, at different undersampling factors. B) NRMSE obtained from noise-free reconstructions of two-crossing fibers, but using $f = 1$ in the multi-tensor model, at different undersampling factors.

In figures 2.3 and 2.4, it is possible to observe some cases with a non-monotonic behavior in the NRMSE measurement of some methods as the US factor was increased. This was explained by a random match between the sampling pattern and the data. To test this, we rotated the fiber orientation by different angles at 15-degree intervals. From figure 2.5, the reconstruction quality is dependent on the sampling pattern; however, the ranking of the methods was preserved.

Finally, we investigated how the quality of the reconstruction changes in the presence of noise. Figure 2.6 shows the NRMSE as a function of the standard deviation of the noise. Column (A) correspond to one fiber with $f = 1$, (B) is for two crossing-fibers with $f = 2$ and (C) to two-crossing fibers with $f = 1$. As expected, the NRMSE increased with the noise. In all cases, MCS performed very well, and similar to M when the model was correct. As also expected, for higher noise levels M performed better, since it does not suffer from over-fitting to noise. The other indexes showed a similar behavior. For example, in (A) at $\sigma = 6\%$, the ΔMSD measurement ranked first MCS with -6.01%; followed by,

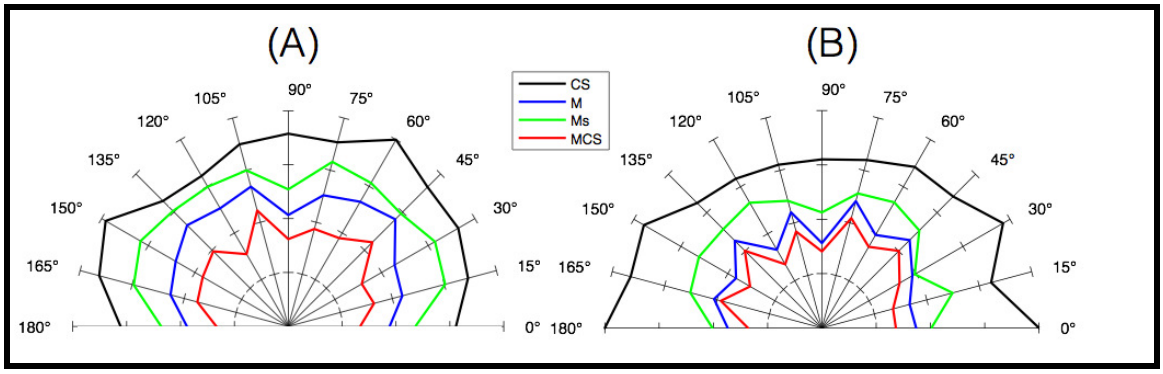

FIGURE 2.5. Evaluation of the reconstruction process according to the angle of the fiber bundle. A) Reconstruction with US factor 4x. B) Reconstruction with US factor 8x.

M with -8.39%; CS with -20.07%; and finally, Ms with -24.61%. PC indicated that all reconstruction methods had correlations above 98%.

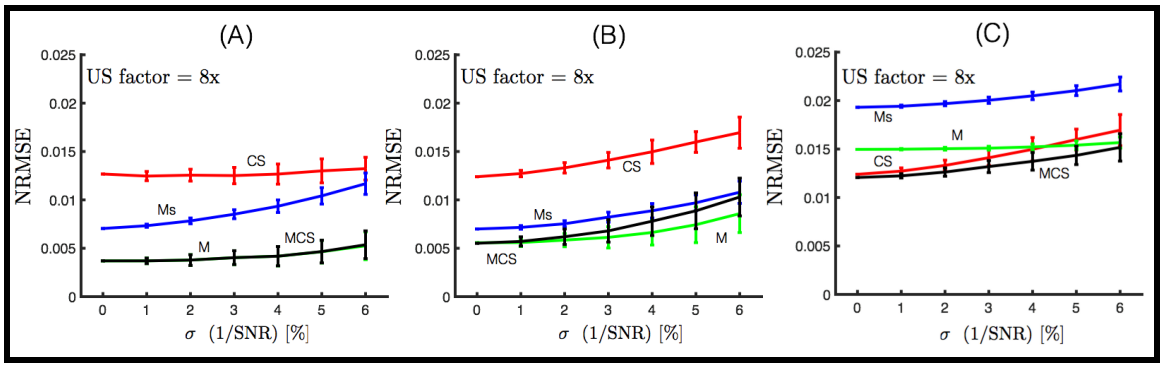

FIGURE 2.6. Quality indexes as a function of noise. Column (A) shows the NRMSE for reconstructions of one fiber for different noise level. Column (B) is for two-crossing fibers. And column (C) is for two-crossing fibers, but using $f = 1$ in M and Ms. Results were consistent with the noise-free experiments and the reconstruction quality was dependent of the correctness of the model.

Diffusion Phantom

Figure 2.7 shows the mean and standard deviation of the NRMSE of the reconstructions of the diffusion phantom for two regions, one containing one fiber, and the other

containing two fibers. For each group, we used 20 voxels to compute the mean and standard deviation. The reconstruction was done with $f = 1$ or $f = 2$ depending on the group. The ground truth is the inverse Fourier transform of the complete acquisition multiplied by the cosine window. As can be seen in figure 2.7, the MCS method had the best NRMSE for all the US factors regardless of the number of fibers. At 8x undersampling in column (A), the ΔMSD measurement ranked first MCS with -16.66%; followed by, CS with -25.93%; M with -34.66%; and finally, Ms with -50.08%. PC indicated that all reconstruction methods had correlations above 98%. For two fibers (column B), the indexes were very similar, preserving the ranking.

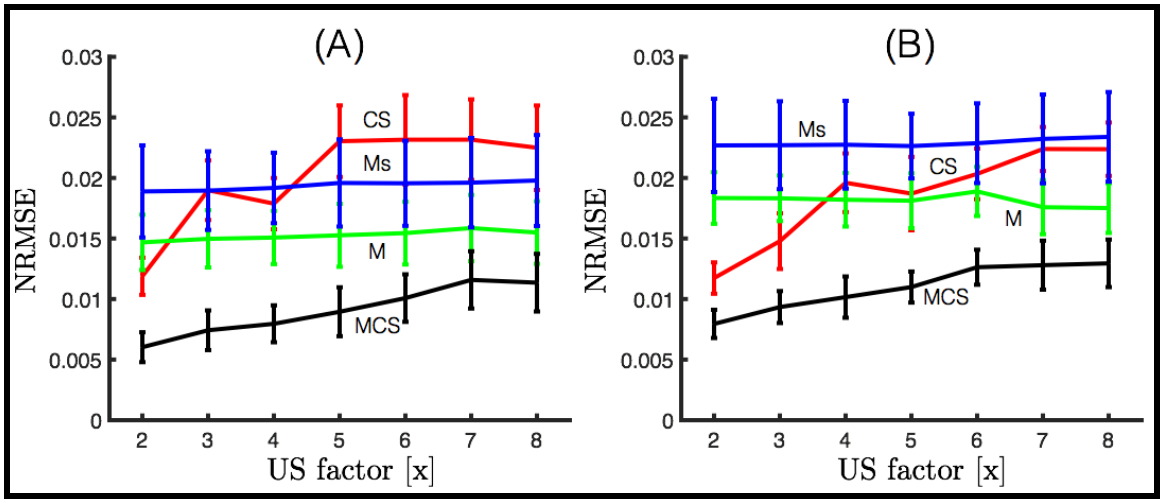

FIGURE 2.7. Results from the diffusion phantom. Column (A): mean and standard deviation of the NRMSE from 20 voxels with one fiber. Column (B): mean and standard deviation of the NRMSE from 20 voxels with two-crossing fibers.

Human Brain

Figure 2.8 shows the mean and standard deviation of the NRMSE for the region of interest in the corpus callosum (133 voxels with different number of fibers) of a human brain acquisition. For model fitting, we used two tensors for all voxels. The behavior was similar to the previous results. The proposed methods M and MCS had the lowest NRMSE

when compared to the other reconstruction methods. At 8x undersampling, the ΔMSD measurement ranked first MCS with -19.53%; followed by, CS with -29.36%; M with -36.34%; and finally, Ms with -48.04%. PC indicated that all reconstruction methods had correlations above 97%. The Δp0 ranked first MCS with -2.83%; followed by, M with -2.91%; Ms with 4.69%; and finally, CS with -10.41%.

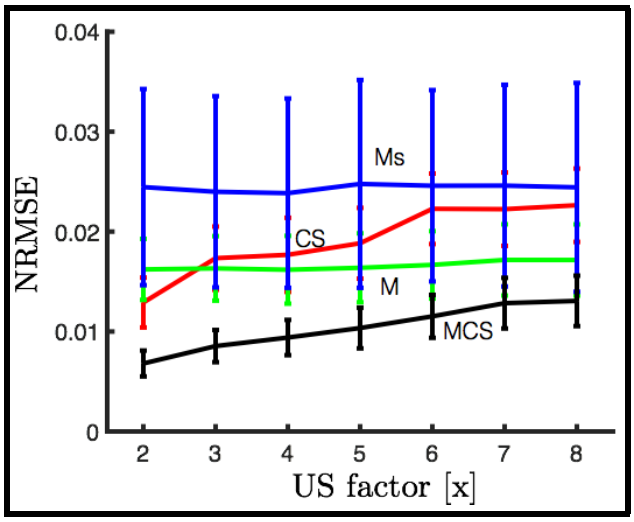

FIGURE 2.8. Results from the in-vivo data. Mean and standard deviation of the NRMSE measurement for 133 voxels inside the region of interest from corpus callosum for different US factors.

Finally, figure 2.9 shows the ODF field from the region of interest in corpus callosum at 4x. It is possible to observe that the ODFs from MCS are similar to the ground truth.

Discussion

The hypothesis of this work was that it could be possible to obtain higher acceleration factors and/or higher reconstruction quality by combining a model-based method and compressed sensing (CS) to reconstruct the diffusion propagator. While the model reconstructs the bulk of the data, CS completes the information not obtained by the model. In this way, both methods work together for a better solution and improve the estimation of

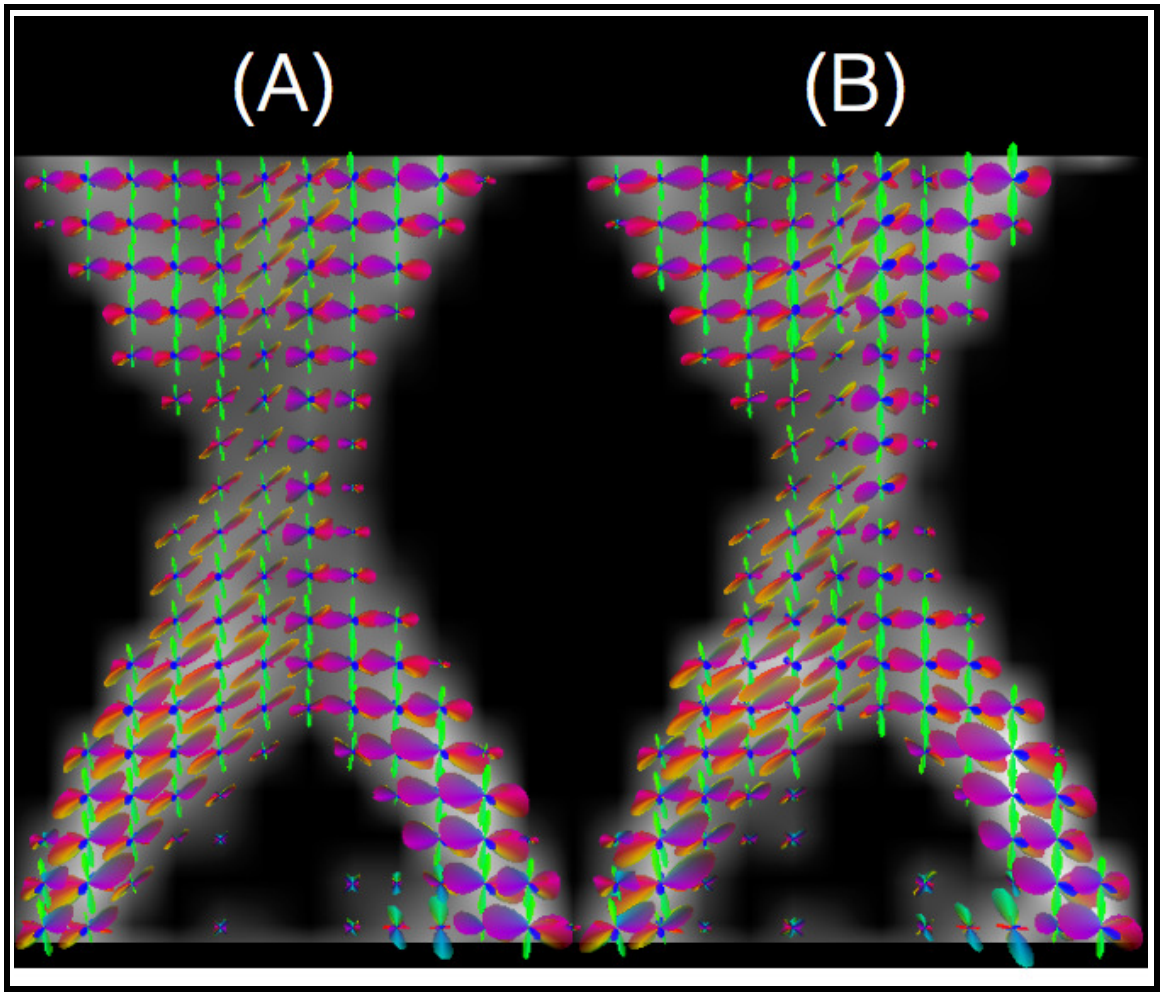

FIGURE 2.9. ODF field from a region of interest from corpus callosum. (A) Ground truth. (B) MCS at 4x undersampling.

the diffusion propagator.

To verify our hypothesis, we tested our proposed methods, M (multi-tensor model with random sampling) and MCS, against the multi-tensor model with spherical sampling and CS under different undersampling factors, sparse domains, noise standard deviations and number of fiber bundles. In order to apply CS after the model fit, we had to estimate the model using randomly distributed samples in q-space. This is an innovative approach since model-based methods like Diffusion Tensor Imaging are commonly used with uniformly distributed samples in a single-shell.

The results from simulations of one fiber bundle and two-crossing fibers showed that the model fit using randomly distributed samples gives a better NRMSE than using distributed samples in a single-shell. Acquiring random samples in q-space may be thought as a multi-shell acquisition. It is not the purpose of this work to compare with DTI since we are interested in the full pdf. Additionally, the acquisition parameters are quite different in terms of the diffusion encoding; which, may invalidate any conclusion. The MSD, p0 and PC also support an increased quality while reconstructing the data with sparse sampling patterns. We did not use other indexes, such as Fractional Anisotropy, Difference in the Number of fiber Compartments, and Angular Error because these measurements are useful in a model-based context, but not for model-free pdf comparison.

In simulations, where the ground truth was the noise-free data, model-based methods performed better than CS when the model was appropriately representing the microstructure of interest. MCS was able to only marginally improve M. Even though Camino simulations are not Gaussian-based, averaging center-of-masses and windowing makes the diffusion propagator resemble a Gaussian pdf. The multi-tensor fitting is a relatively good model under this scenario. However, when the model is incorrect and fails, MCS improved the final result. We observed that whenever the number of tensors f exceeded the real number of fibers in the object, it was possible to obtain a better estimation of the diffusion propagator. However, increasing excessively f may generate multiple answers. In this work we used the multi-tensor model with only one and two fibers. Although it is known that in cerebral regions it is possible to find three or even more fiber bundles per voxel, the purpose of this work was to show the feasibility of combining model-based with model-free methods. Our method, which is a combination of the multi-tensor model and CS, was capable to achieve similar results with the best reconstruction methods depending on the correctness of the model.

In phantom and in-vivo acquisitions, where the ground truth was fully sampled, despite noisy acquisition, our proposed methods showed better quantification of the pdf when measured with the NRMSE, the Pearson correlation coefficient and the Mean-Squared-Displacement. The key difference in the results between simulations and actual data come

from the presence of noise in the ground truth.

In general, model fitting performs relatively well in presence of noise, in which case the improvement of the second stage in MCS is only moderate. In any case it is appealing that the worst case scenario for MCS results in the model fit. The improvements obtained for a more correct estimation of the pdf, may lead to valuable information to describe the underlying microstructure. In this work, we chose the multi-tensor model, but in general other models can be used.

3. COMPARISON OF Q-SPACE RECONSTRUCTION METHODS FROM UNDERSAMPLED DIFFUSION SPECTRUM IMAGING DATA

This Manuscript was accepted by the journal Magnetic Resonance in Medical Sciences (Magn Reson Med Sci) in March, 2019.

3.1. Introduction

Diffusion, at a microscopic level, consists of molecules' random displacements given a diffusion time Δ. This set of possible displacements, $\mathbf{r}$, has associated a probability density function (pdf), $p(\mathbf{r}|\Delta)$, which is the so-called diffusion propagator (Callaghan, 1991, 338 p.). The reconstruction of the diffusion propagator is relevant to describe the complete diffusion process and to extract biomarkers without any assumption of the microstructure. Diffusion MRI is a technique that estimates the diffusion propagator for each imaging voxel. Usually, T2-weighted images are made sensitive to diffusion by including diffusion encoding gradients, with amplitude $|\mathbf{G}|$ and duration δ, that are separated by the diffusion time Δ. Thus, these images contain the effect of diffusion in the direction of $\mathbf{G}$ over Δ (Stejskal & Tanner, 1965). The data acquired in q-space, $\mathbf{q} = \delta\mathbf{G}\gamma/2\pi$, correspond to the Fourier transform of the diffusion propagator. The direct reconstruction of the center-of-mass (Mitra & Halperin, 1995) diffusion propagator from the magnitude of the complete q-space data is called Diffusion Spectrum Imaging (DSI) (Wedeen et al., 2000, 2005). Further details and recent improvements for DSI acquisition and reconstruction can be found in (Lacerda et al., 2016; Paquette et al., 2016; Tian et al., 2016).

The DSI reconstruction for each voxel is defined as

$$p(\mathbf{r}) = \int_{\mathcal{R}^3} |E(\mathbf{q})| e^{i2\pi\mathbf{q}\mathbf{r}} d\mathbf{q}, \tag{3.1}$$

where $p(\mathbf{r})$ is the diffusion propagator given the fixed diffusion time for a traditional q-space acquisition and $E(\mathbf{q})$ are the values of the q-space samples. The Fourier transform for DSI requires a minimum q-space sampling to fulfill the Nyquist criterion. Since each q-space sample is a diffusion-weighted scan, the DSI acquisition becomes time-expensive for clinical settings.

When the full diffusion propagator is not needed, model-based methods are suitable to

compute specific diffusion indices. Typically, the propagator is approximated by the diffusion tensor model. Even though diffusion tensor imaging (DTI) (Basser, Mattiello, & LeBihan, 1994) has shown relevance in clinical applications (Le Bihan et al., 2001; Dong et al., 2004), this approximation may not describe correctly, or completely, the tissue microstructure. For example, DTI provides a biased result, if the voxel contains two or more non-parallel fibers (Landman et al., 2012). Another model technique is diffusion kurtosis imaging (DKI) (Jensen et al., 2005) that incorporates the non-Gaussian diffusion using the kurtosis parameter. Unfortunately, DKI cannot be used in a complete q-space regime due to high signal to noise ratio requirements. A discussion on models in diffusion MRI and their assumptions can be found in (Hagmann et al., 2006; Assemlal, Tschumperl, Brun, & Siddiqi, 2011; Ferizi et al., 2014).

On the other hand, compressed sensing (CS) (Candes & Romberg, 2007) has been used to fully reconstruct the diffusion propagator from undersampled q-space data. CS uses as prior knowledge that the data should be represented sparsely in some domain. The first publication of CS in DSI (Menzel et al., 2011) proposed to use Total Variation and Haar wavelets as sparse domains. Soon after, (Bilgic et al., 2012) proposed to use a data-driven dictionary as sparse domain. This method had around 1.6-fold improvement in reconstruction quality, measured as the normalized mean squared error (NMSE) when compared to (Menzel et al., 2011); however, that work lacked simulations to describe in a more controlled environment the performance and limitations of both methods. Recently, (Paquette et al., 2015) proposed a joint improvement by selecting the Cohen-Daubechies-Feauveau 9/7 wavelet as sparse domain; and a sparse undersampling pattern using a uniform angular distribution with randomly allocated samples along radial profiles. That work analyzed most of the commonly used sparsifying transfoms like wavelets, Total Variation and the identity basis from (Menzel et al., 2011; Bilgic et al., 2012; S. Merlet, 2013); and different undersampling patterns as in (Menzel et al., 2011). However, a comprehensive comparison between CS-based reconstruction methods with reconstruction methods that fit the q-space signal to a highly efficient set of continuous basis functions is yet to be done. A representative method using continuous basis functions is the mean apparent propagator (MAP) (Ozarslan

et al., 2013; Avram et al., 2016), a method which approximates the diffusion propagator using Hermite functions. The first basis function corresponds to the diffusion tensor as in DTI; and the rest characterize any deviation from DTI. Up to some extent, MAP has a better characterization of q-space than other continuous basis functions because it takes advantage of the anisotropic nature of its scaling tensor, whereas other methods that use isotropic scaling would require more basis functions to obtain the same reconstruction quality (Fick, Wassermann, Caruyer, & Deriche, 2016).

In this work we compared the methods of MAP, compressed sensing using the identity as sparsifying transform (CSI) and compressed sensing using data-driven dictionaries (CSD). The novelty of this comparison is that we included MAP in addition to CSI and CSD methods, which were previously analyzed in (Bilgic et al., 2012; Paquette et al., 2015).

3.2. Methods

The three reconstruction methods were tested using retrospectively undersampled q-space data from Monte Carlo simulations and from an in vivo brain acquisition. The reconstructions were compared with the fully sampled data as ground truth to evaluate reconstruction quality and retrieval of propagator-based diffusion indices. Finally, the in vivo reconstructions were used for a visual analysis of the centrum semiovale.

3.2.1. Data sets

Simulations and in vivo data were done in an 11x11x11 q-space Cartesian grid. Each complete data set had 257 q-space samples contained in a discrete half-sphere with a radius of 5 samples. The other half was obtained by symmetry.

For the Monte Carlo simulations we used the crossing substrate from the Camino software (Cook et al., 2006; Hall & Alexander, 2009). The substrate corresponds to two fiber populations in interleaved planes; and one population was rotated with respect to the other population in order to resemble two crossing fiber bundles with a certain crossing angle. The fiber populations were done with a cylinder radius of 2 μm and a cylinder separation of 5.1 μm.

We used 10^3 time steps for 10^5 spins, initially uniform-distributed, to diffuse in the crossing substrates. We simulated two crossing fibers with crossing angles in $[0, 15, 30, ..., 90]^\circ$. The diffusion encoding parameters for the simulations were with $\delta = 31$ ms and $\Delta = 43.2$ ms, where $q_{\text{max}} = 62.21$ mm^{-1} (or $b_{\text{max}} = 6600$ s/mm^2); and TE $= 88$ ms. Gaussian noise was added to the real and imaginary part of the simulated measurements to obtain Rician noise in the magnitude of the diffusion signal as $E(\mathbf{q})_i = \sqrt{(E_{\text{simulation}}(\mathbf{q}) + \eta_{i,1})^2 + \eta_{i,2}^2}$, where $\eta_{i,1}, \eta_{i,2} \sim N(0, \sigma)$ and standard deviation $\sigma = 1/\text{SNR}$ (Paquette et al., 2015). For simulations, σ was chosen as $[0, 1, 2, ..., 10]$ % of the peak value. Each noisy experiment was repeated fifty times. The ground truth for simulations were the inverse Fourier transform of the noise-free and fully sampled measurements, $\mathscr{F}^{-1}\{E_{\text{simulation}}(\mathbf{q})\}$.

For the in vivo acquisition, we used fully sampled DSI acquisitions collected as in (Tobisch et al., 2018). Data were acquired in a 3T Prisma scanner at 1.5mm isotropic resolution (TE/TR = 105/6100 ms, bmax = 6800 s/mm ^{2}, Δ = 51.3 ms, δ = 20.1 ms) using a 64-channel head-neck coil. 257 diffusion weighted and 8 interleaved, non-weighted images were acquired with both anterior-to-posterior and posterior-to-anterior phase encoding. Finally, the acquired data were processed in FSL to estimate and correct for susceptibility geometric distortions, eddy currents and subject motion (Andersson, Skare, & Ashburner, 2003; Andersson & Sotiropoulos, 2015). The total scan time was 55 minutes. The ground truth for in vivo data were the inverse Fourier transform of the fully sampled q-space from each voxel.

The fully sampled data from simulations and acquisitions were retrospectively undersampled from 2x to 8x to test reconstruction performance. The sampling pattern was a 3D Cartesian q-space grid where the central 3x3x3 samples were always sampled and the rest, depending on the undersampling factor (USF), were picked randomly with a variable decreasing density distribution (Menzel et al., 2011; Bilgic et al., 2012). We repeated the reconstructions with ten different sampling patterns using the previously mentioned strategy in order to establish if the reconstruction quality was invariant to both, the microstructure orientation and the particular sampling pattern. To test dependency on the angle between undersampling pattern and fiber orientation, we also reconstructed simulations (as

described above) of a single fiber with USF = 4, and noise = 5% rotated in a plane using $[0, 15, ..., 165]^{\circ}$.

3.2.2. Reconstruction

The reconstruction methods are as follows.

Mean apparent propagator (MAP)

Mean apparent propagator (MAP) (Ozarslan et al., 2013; Avram et al., 2016) is a least-squares optimization that finds the best fit to q-space data as,

$$\hat{\mathbf{c}} = \text{argmin}_{\mathbf{c}} ||\Phi(\mathbf{A}, \mathbf{q})^{\mathrm{T}}\mathbf{c} - \mathbf{E}(\mathbf{q})||_2^2 \quad \text{s.t} \quad \Psi(\mathbf{A}, \mathbf{r})^{\mathrm{T}}\mathbf{c} \geq 0, \tag{3.2}$$

where $\mathbf{c}$ are the coefficients to estimate, $\Phi(\mathbf{A}, \mathbf{q})^{\mathrm{T}}$ is the basis in q-space and $\Psi(\mathbf{A}, \mathbf{r})^{\mathrm{T}}$ is the corresponding basis in pdf-space. Each basis function from Φ and Ψ are scaled by $\mathbf{A} = \text{diag}(\mu_1^2, \mu_2^2, \mu_3^2)$, where $\mu_{\{1,2,3\}} = \sqrt{2\lambda_{\{,2,3\}}\tau}$ is the square root of the mean displacement for each eigenvalue of the diffusion tensor, which was fitted to all q-space measurements. Because of the eigen decomposition of the diffusion tensor, each 3D basis function in Φ is the combination of three 1D orthogonal basis functions, $\phi_n(\mu, q)$, where each one of them corresponds to the th-order Hermite polynomial, , modulated by a Gaussian-like term. The 1D basis function is defined as

$$\phi_n(\mu, q) = \frac{1}{i^n\sqrt{2^n n!}} \exp\left(-\frac{(2\pi\mu q)^2}{2}\right) H_n(2\pi\mu q). \tag{3.3}$$

An interesting property of $\phi_n(\mu, q)$ is that its Fourier transform also results in a Gauss-Hermite function, $\psi_n(\mu, r)$, as

$$\psi_n(\mu, r) = \frac{1}{\mu\sqrt{2^{n+1}\pi n!}} \exp\left(-0.5\left(\frac{r}{\mu}\right)^2\right) H_n\left(\frac{r}{\mu}\right); \tag{3.4}$$

and this is the 1D basis function that generates each 3D basis function for Ψ. Furthermore, Φ and Ψ share the same coefficients, so it is possible to optimize data consistency while imposing properties of the diffusion propagator like non-negativity. Therefore, the first

basis function ($n = 0$) in Φ and Ψ correspond to Gaussian diffusion whereas the remaining basis functions ($n > 0$) correspond to non-Gaussian diffusion. The bases Φ and Ψ were constructed as in (Ozarslan et al., 2013) and according to recommendations from (Avram et al., 2016). Hence, we used a maximum basis order of six, $N_{max} = 6$, which corresponds to the first 50 basis functions from the truncated infinite series. The number of basis functions is defined by the following expression:

$$\Phi(\mathbf{A}, \mathbf{q}) = \sum_{N=0}^{N_{\max}} \sum_{\substack{i,jk>0 \\ i+j+k=N}} \phi_i(\mu_1, q_1)\, \phi_j(\mu_2, q_2)\, \phi_k(\mu_3, q_3), \tag{3.5}$$

in order to ensure that all the possible combinations for the respective maximum basis order are included to design both bases, Φ and Ψ. For symmetric signals, only even values of N are non zero; therefore, the number of basis functions is defined by $B(N_{\max}) = (N_{\max} + 2)(N_{\max} + 4)(2N_{\max} + 3)/24$. Finally, in this work the diffusion propagator from MAP was obtained as $\hat{p}(\mathbf{r}) = \mathscr{F}^{-1}\{\Phi(\mathbf{A}, \mathbf{q})^{\mathrm{T}}\hat{\mathbf{c}}\}$.

Compressed Sensing (CS)

Compressed sensing (CS) is a method that imposes sparsity in the representation of the diffusion propagator while maintaining data consistency:

$$\hat{p}(\mathbf{r}) = \text{argmin}_{p(\mathbf{r})} \ \frac{1}{2}||S\mathscr{F}p(\mathbf{r}) - E(\mathbf{q})||_2^2 + \lambda||\Omega p(\mathbf{r})||_1. \tag{3.6}$$

Data consistency ensures the similarity between the estimation and the acquired data $E(\mathbf{q})$, where $S\mathscr{F}$ is the undersampled Fourier operator. The regularization constrains $p(\mathbf{r})$ to be represented sparsely in the domain Ω. For the data sets in this work we used $\Omega = \mathrm{I}$ since it is suitable and computationally efficient (Bilgic et al., 2012; Paquette et al., 2015). Finally, the optimization was done using a non-linear conjugate gradient implementation. The parameter λ was empirically fixed according to the procedure explained in Appendix A.

CS using data-driven dictionaries (CSD)

CSD is a CS reconstruction that obtains the diffusion propagator by fitting coefficients to a trained dictionary (Bilgic et al., 2012). The method iteratively solves:

$$\begin{gathered} W_{j,j}^{t} = \operatorname{diag}\left(\left|x_j^t\right|^{1/2}\right) \\ s^t = \arg\min_s \|s\|_2^2 \text{ such that } S\mathscr{F}\mathbf{D}\mathbf{W}^t\mathbf{s} = E(\mathbf{q}) \\ \mathbf{x}^{t+1} = \mathbf{W}^t\mathbf{s}^t \end{gathered} \tag{3.7}$$

where $\mathbf{D}$ is the dictionary obtained from a training set of diffusion propagators using the ksvd algorithm (Aharon, Elad, & Bruckstein, 2006). The spirit of this dictionary is to sparsely concentrate the variance between propagators in the first k-th atoms from the dictionary. $\mathbf{W}^t$ is a diagonal weighting matrix whose jth diagonal entry is denoted $W_{j,j}^t$ and this matrix multiplied with $\mathbf{s}^t$ equals $\mathbf{x}^{t+1}$, which are the corresponding coefficients after the last iteration t. The focal underdetermined system solver (FOCUSS) promotes ℓ_1-sparsity on the coefficients $\mathbf{x}^{t+1}$ through reweighted ℓ_2-optimizations on the auxiliary variable $\mathbf{s}^t$. For more information see (Ye, Tak, Han, & Park, 2007). Implementation was downloaded from `http://martinos.org/~berkin/software.html`.

For the simulation case, the training was done with simulated noise-free single fibers rotated across the three axes in the pdf-space. For this purpose, we simulated in Camino single fibers using a wide range of diameters and separation between cylinders. The experiments were done using this dictionary, but with simulations of two crossing fibers (as described above). For the in vivo case, the training was done with propagators from fully-sampled q-space data of one axial slice. The in vivo training was done as in (Bilgic et al., 2012). The reconstruction was done along a coronal slice to avoid any bias in favor of the CSD method. The particular axial slice used for training the dictionary was removed from the comparison. The diffusion propagator was obtained as $\hat{p}(\mathbf{r}) = \mathbf{D}\mathbf{x}^{t+1}$.

Finally, in all reconstruction methods the negative values found in each diffusion propagator were clipped to zero (Paquette et al., 2016).

3.2.3. Quality indices and visualization

To compare the reconstruction $\hat{p}$ against the ground truth p, we used the normalized mean squared error, NMSE$\{\hat{p}, p\} = ||\hat{p} - p||_2^2/||p||_2^2$, and the Pearson's correlation coefficient (PC) (Paquette et al., 2015). We also compared two propagator-based diffusion indices: the mean squared displacement (MSD) (Wu et al., 2008) and the return to zero probability (p0) (Assaf et al., 2000; Ozarslan et al., 2013). The MSD is the second moment of the pdf and the relative MSD error was defined as $\Delta\text{MSD}\{\hat{p}, p\} = (\text{MSD}(\hat{p}) - \text{MSD}(p))^2/\text{MSD}(p)^2$. The relative return to zero probability error was defined as $\Delta\text{p0}\{\hat{p}, p\} = (\text{p0}(\hat{p}) - \text{p0}(p))^2/\text{p0}(p)^2$.

The in vivo reconstructions were loaded into DSI Studio (`http://dsi-studio.labsolver.org`) to visually analyze their performance around the centrum semiovale. This region contains the intersection of multiple white matter bundles; therefore, it is a complex area in the brain suitable to evaluate multiple crossing angles (Tian et al., 2016).

3.3. Results

3.3.1. Simulations

Figure 3.1 shows the mean and standard deviation of the NMSE index for the reconstructions of a single fiber rotated at $[0, 15, ..., 165]^\circ$ in a plane with USF $= 4$. Figure 3.1A shows that the reconstructions were slightly biased for angles aligned with the Cartesian grid (0° and 90°), i.e. the reconstructions are not entirely independent of the microstructure orientations. This result agrees with what was obtained in (Lacerda et al., 2016), where it is stated that the orientation distribution function reconstruction depends on q-space acquisition and resolution. Figure 3.1B shows a slice of the sampling pattern. To avoid this orientational bias between the Cartesian grid and the orientation from simulations, we rotated the noise-free simulations by $\pi/4$ in the three axes.

Figure 3.2 shows the mean and standard deviation of the NMSE and PC indices (50 reconstructions) of two crossing fibers while varying noise level, undersampling factor and crossing angle. Both reconstruction quality indices are consistent, ranking the methods in

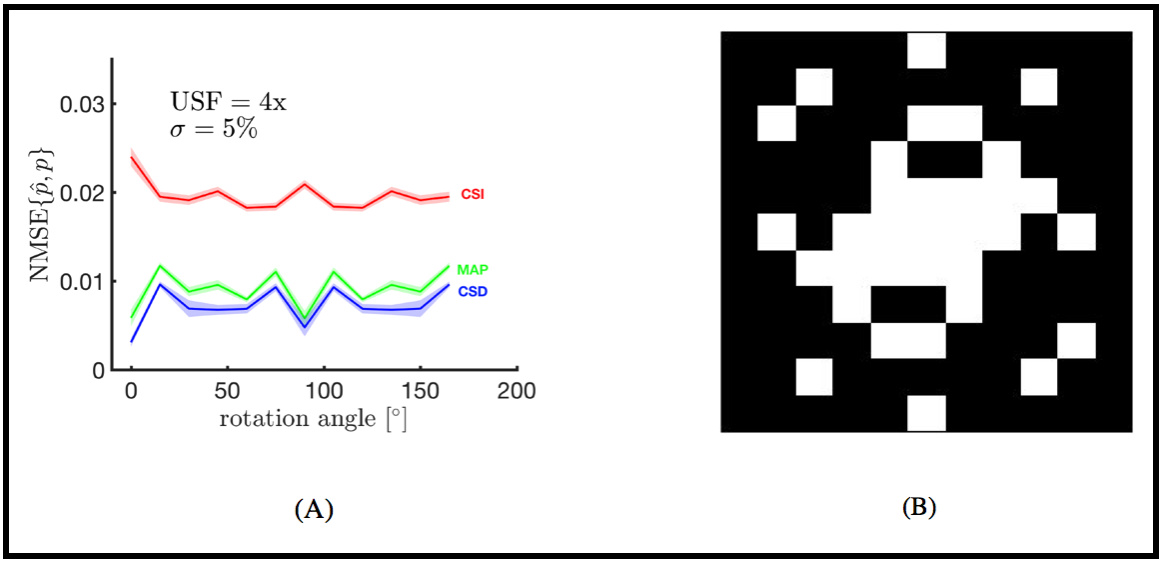

FIGURE 3.1. Effect of the fiber orientation in the reconstruction of a single fiber rotated over in pdf-space. The NMSE as a function of the rotation angle in (A) shows that the reconstructions were slightly biased for rotation angles aligned with the Cartesian grid (0° and 90°). The corresponding central plane of the 3D q-space sampling pattern at 4x is shown in (B) as reference.

the order: CSD, MAP and CSI. Figure 3.2A depicts NMSE as a function of noise. At $\sigma = 5\%$, NMSE obtained with MAP and CSI were 1.1-fold and 1.5-fold higher than the NMSE from CSD; which agrees with the results in (Bilgic et al., 2012). The same behavior can be seen in figure 3.2D for the PC. Figure 3.2B depicts NMSE as a function of USF and it shows that the CSD method performed better than the other methods for undersampling factors above 4x; which again agrees with (Bilgic et al., 2012). Between 2x and 8x, the mean NMSE from CSD increased around 2% while MAP and CSI mean NMSE increased around 8% and 5% respectively. This is also shown for PC in figure 3.2E. Figure 3.2C depicts NMSE as a function of the crossing angle. The ranking of reconstructions is preserved and it is worth mentioning the increased error from CSI at lower crossing angles. More prominent Gibbs ringing were observed at small crossing angles, which may be an explanation for the increased errors. Furthermore, resolving small crossing angles is generally known to be challenging, even with advanced diffusion MRI, which may as well explain increased NMSE values.

The previous results were reproduced for 9 additional undersampling patterns in order to

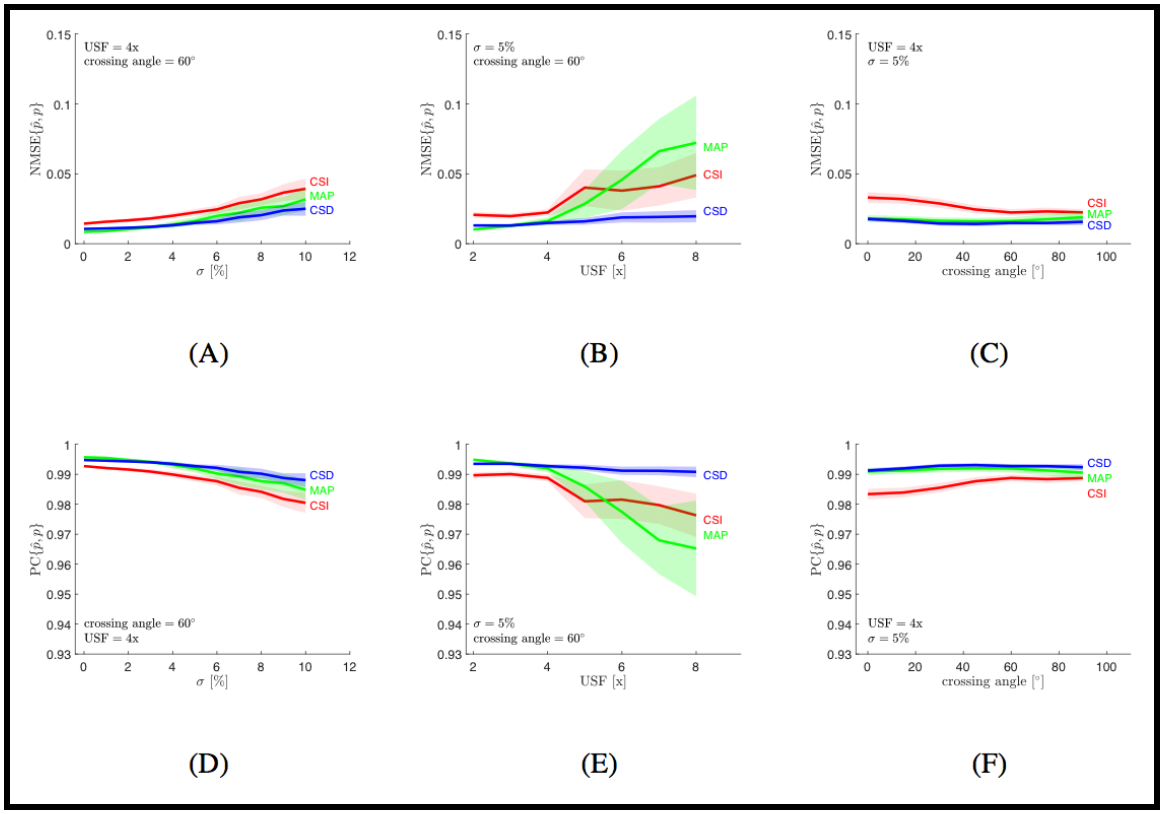

FIGURE 3.2. Reconstruction quality indices for different settings from simulations of two crossing fibers. The first row corresponds to NMSE as a function of noise level σ (panel A); as a function of undersampling factor USF (panel B); and as a function of crossing angle (Panel C). The second row corresponds to PC as a function of the same variables (panels D,E and F).

verify whether these results were influenced by the selection of the specific sampling pattern. Although there is a natural variance from randomized patterns, the ranking of the reconstruction methods was preserved and most of them agree quantitatively too, which can be seen in Appendix B.

Finally, figure 3.3 shows the mean and standard deviation of the propagator-based relative index errors as functions of noise, undersampling and crossing angle from the reconstructions of crossing fibers. Figures 3.3A,D show almost exact recovery (error below 2%) of the index as a function of noise. Furthermore, figure 3.3A demostrates that CSI may provide the best relative MSD error even at a relatively high undersampling factor of 4, if the

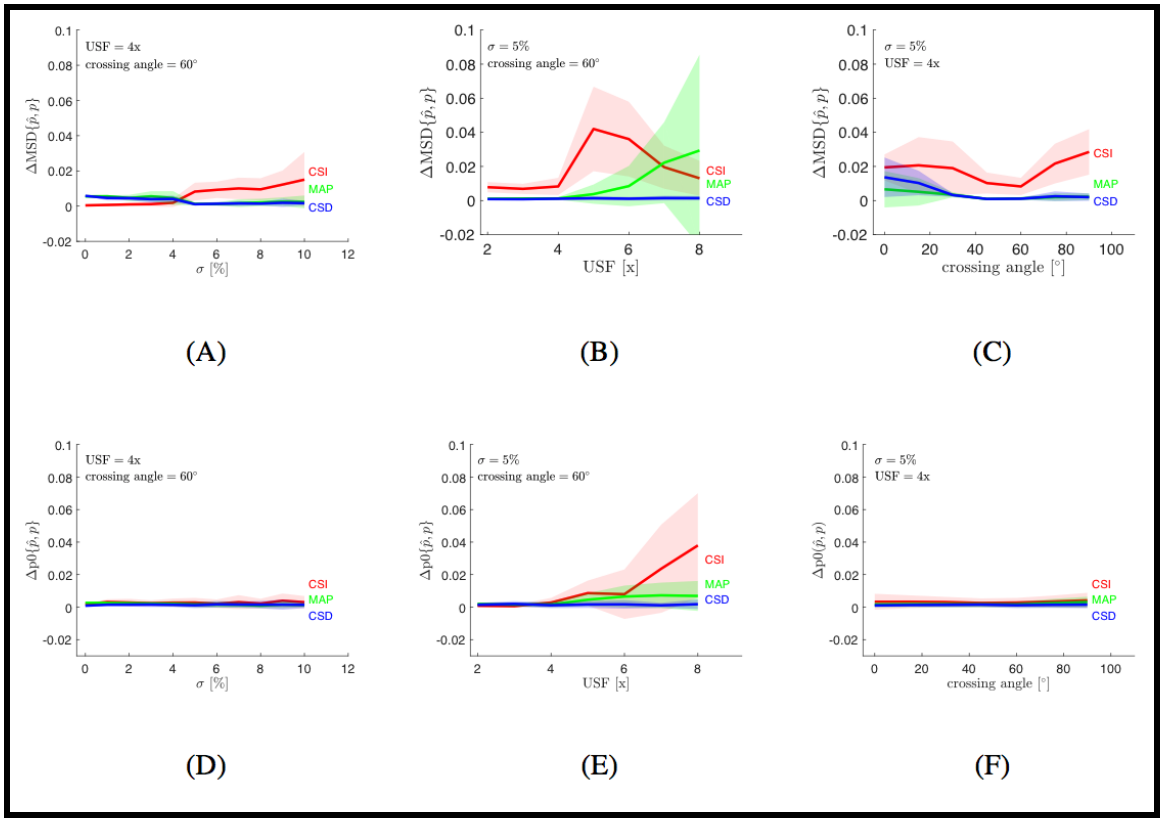

FIGURE 3.3. propagator-based diffusion indices for different settings from simulations of two crossing fibers. The first row corresponds to relative MSD error as a function of noise level σ (panel A); as a function of undersampling factor USF (panel B); and as a function of crossing angle (Panel C). The second row corresponds to relative p0 error as a function of the same variables (panels D,E and F).

noise level is relatively low ($\sigma \leq 4\%$). Figures 3.3B,E show that both indices suffer considerably for undersampling factors above 4x in the cases of CSI and MAP. Figure 3.3C depicts the relative MSD error as a function of the crossing angle and it shows that the relative MSD obtained from CSI has around a 2% error with respect to the ground truth. This may indicate that at USF$= 4$, CSI reconstruction is obtaining a close approximation of the shape, but not an exact recovery of the ground truth which slightly influences the retrieval of the MSD index. Finally, figure 3.3F shows an excellent estimation of p0 for all crossing angles.

3.3.2. In vivo data

Figure 3.4 shows spatial maps of NMSE, (1-PC), relative MSD error and relative p0 error for the reconstruction methods using USF $= 4$ along a coronal slice. For visual clarity the Pearson's correlation coefficient has been subtracted from one, i.e. lower values indicate greater correlation to the ground truth. The axial slice used to train the dictionary of the CSD method was removed from the comparison (red line). In general terms, the three methods show reconstruction quality below five percent error; except in areas containing cerebrospinal fluid or deep brain areas. On the other hand, the methods show high reconstruction quality for the corticospinal tract, the corpus callosum and the corona radiata in comparison with other regions. Additionally, reconstruction quality indices indicate a correlation with the SNR which, for small parallel imaging factors and hence negligible geometry factors, can be assumed to be linked to the receive coil sensititvity map (decreasing towards the center of the brain). With regard to the estimation of the relative MSD and relative p0 errors, the MAP reconstruction shows higher quality when compared with CSI and CSD methods. We included the spatial maps of NMSE for the three reconstruction methods using USF $= 4, 5, 6$ and 8 in Appendix C.

Finally, figure 3.5 shows the generalized fractional anisotropy maps from the ground truth and the reconstruction methods under investigation. From each map, we zoomed the region around the centrum semiovale for visual inspection of the directional information from the propagators obtained for those voxels. All methods show high similarity with the ground truth, but there are areas where directional information has errors (see for example the upper left corner of the zoomed areas).

3.4. Discussion

In this work, three methods to reconstruct the diffusion propagator from undersampled q-space data were compared. These methods were MAP, CSI (using the identity as sparsifying transform) and CSD (using data-driven dictionaries for sparse representation). The

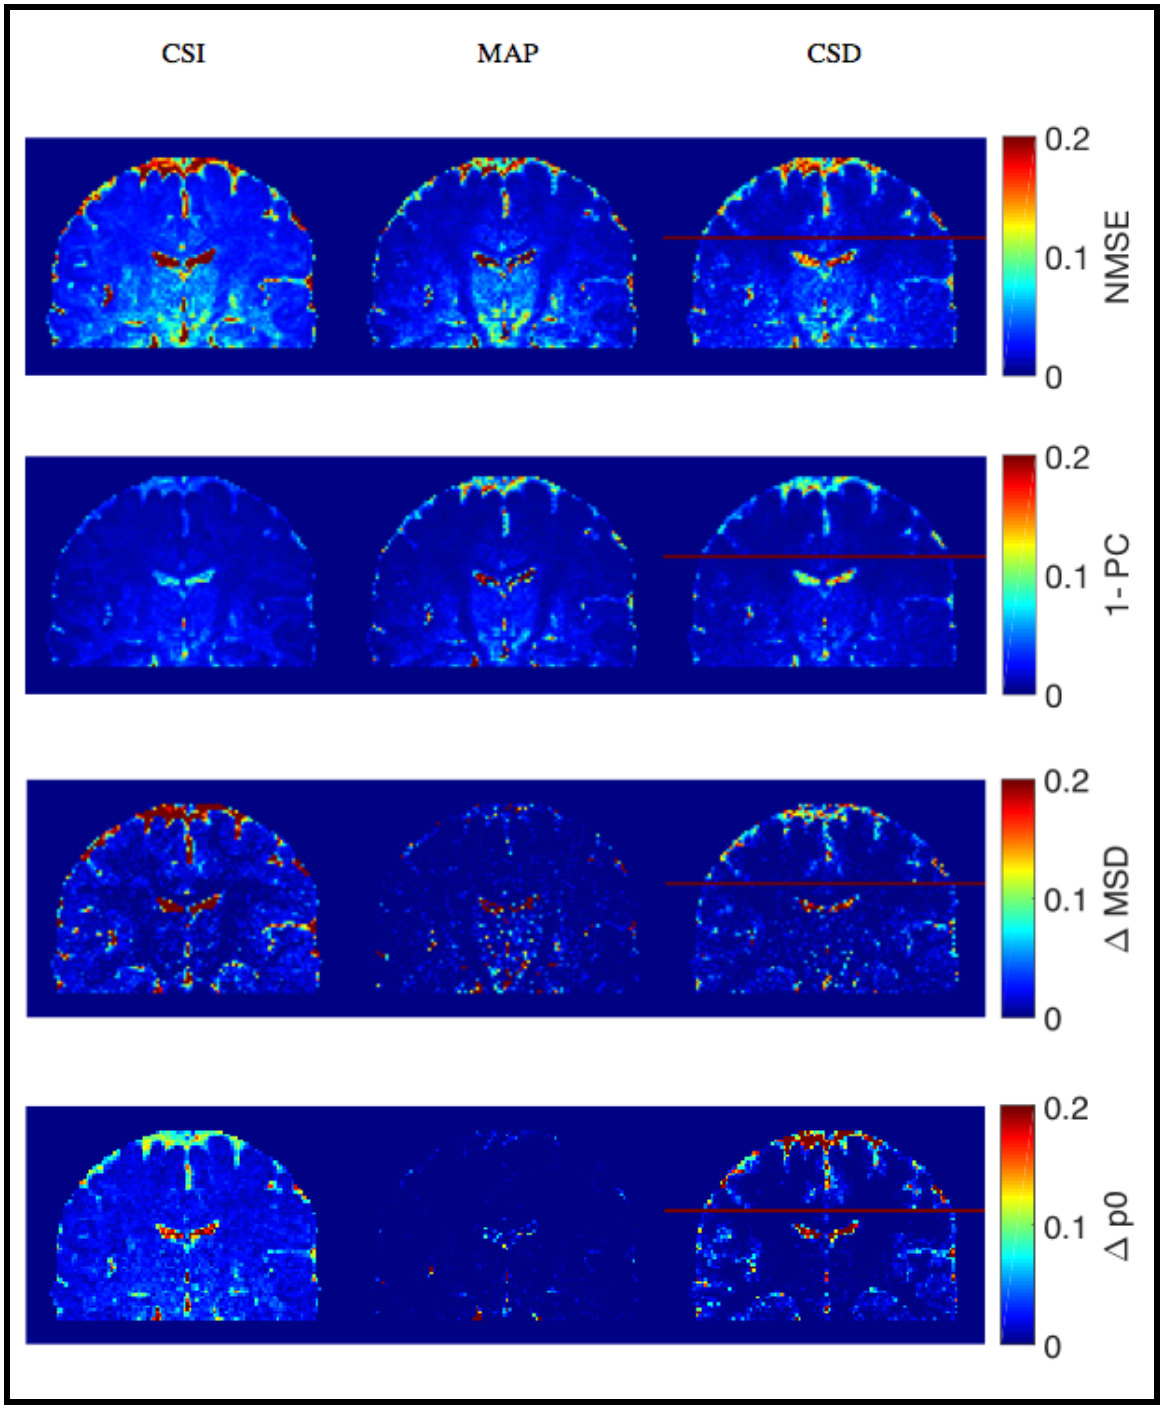

FIGURE 3.4. Index-based maps from in vivo reconstructions. The first row corresponds to the NMSE maps; the second row to (1-PC) maps; the third row to relative MSD error maps; the fourth row to relative p0 error maps. The columns are the reconstruction methods: CSI, MAP and CSD. From the NMSE and (1-PC) maps, all methods are very similar and show reconstruction errors below five percent, except in areas containing cerebrospinal fluid or areas deeper in the brain. The real difference comes on the estimation of diffusion indices, where MAP shows more accuracy in recovering the MSD and p0 indices. CSI reconstruction provides the worst recovery, but its error is around five percent only.

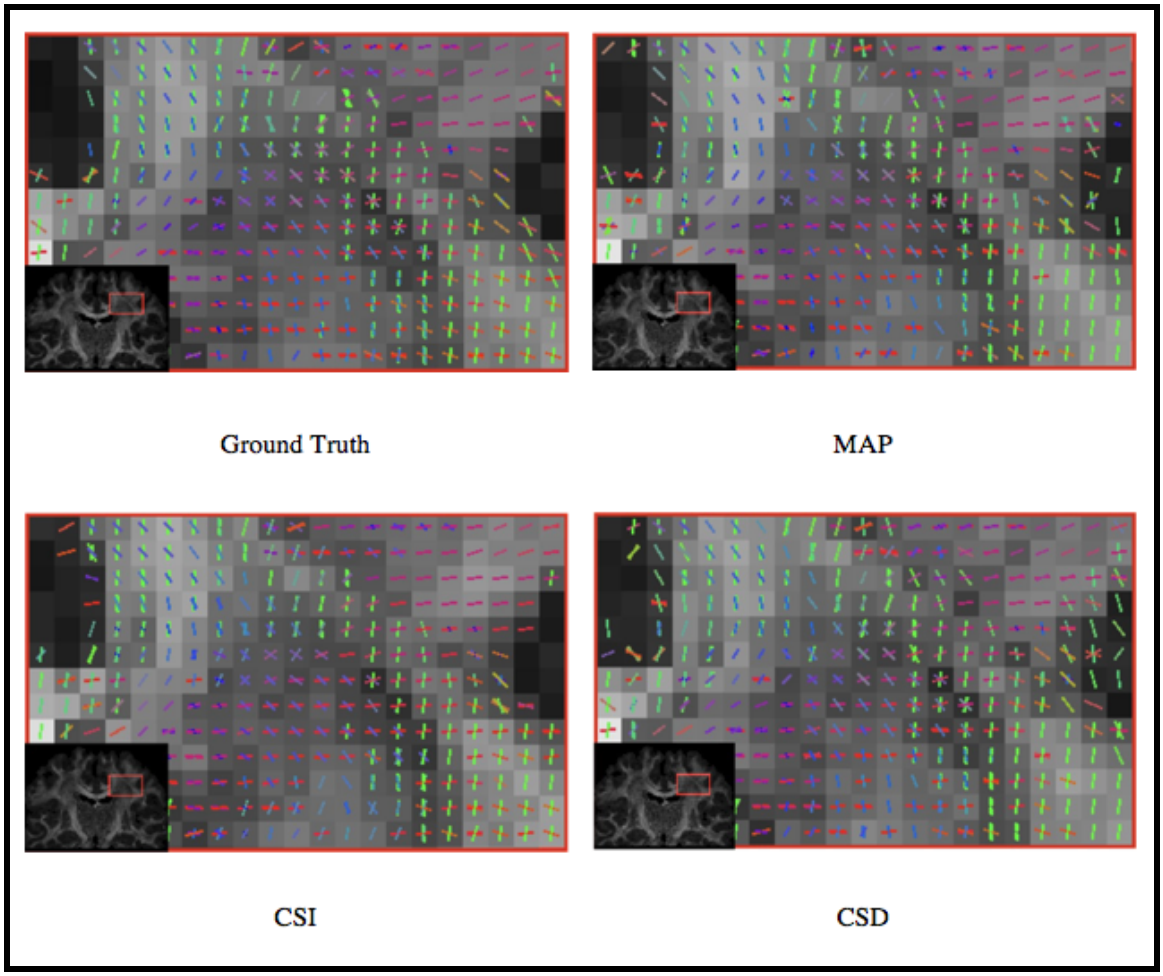

FIGURE 3.5. generalized fractional anisotropy maps from the ground truth and the reconstruction methods under investigation. From each map, we zoomed the region around the centrum semiovale for visual inspection of the directional information from the propagators obtained for those voxels. All methods show high similarity with the ground truth, but there are areas where directional information has errors (see for example the upper left corner of the zoomed areas).

comparison was done in terms of the reconstruction quality (mean squared error and Pearson's correlation coefficient) and in terms of propagator-based indices (MSD and p0). The novelty of this comparison is that we included MAP in addition to CSD and CSI methods, which were previously analyzed (Bilgic et al., 2012; Paquette et al., 2015). We included in the simulations different reconstruction settings of noise, undersampling factor and crossing angles. In this way, we provided a thorough analysis to determine the advantages and limitations for the different methods; and complement the previous work. Our study has

the same spirit as (Hutchinson et al., 2017), but applied to reconstruction methods using Cartesian data with high q-space samples. Across the experiments, we observed that the Cartesian nature of the acquisition could influence the reconstruction of the directional information from diffusion MRI, so sampling patterns should be designed as isotropic as possible, taking this fact into consideration as in (Paquette et al., 2015; Tobisch et al., 2014).

The three reconstruction methods are representative procedures for recovering q-space from undersampled data. One method is based on fitting q-space signal to a highly efficient set of continuous basis functions; and two CS-based methods that sparsely encode q-space by means of a sparsifying transform or by a dictionary constructed from the same measurements. MAP (Ozarslan et al., 2013) is a set of basis functions where the first basis function corresponds to Gaussian diffusion (DTI) and the extra basis functions are used to characterize non-Gaussian diffusion; however, there is a trade-off between additional basis functions and the required number of q-space samples for a robust optimization. This work reports that MAP using 50 basis functions worked well in terms of reconstruction quality and for recovering the propagator-based indices at low undersampling factors. CS reconstruction using the identity as sparse domain worked very similar when compared to MAP and CSD, but there was a consistently higher error at realistic noise levels. Although results from CS reconstruction could change by the selection another sparsifying transform or the value of the Lagrangian multiplier, in our experiments the identity was enough. This is confirmed since our results were similar to the results obtained in (Bilgic et al., 2012) where they used wavelets and Total Variation as sparsifying transforms. Furthermore, in (Paquette et al., 2015) it is shown that it produces satisfactory reconstruction performance in comparison to other sparsifying transforms. CSD, on the other hand, assures the sparsity requirement in CS theory by using an ideal sparse dictionary constructed from the same measurements by means of the ksvd algorithm. Additionally, it provides an iterative parameter-free reconstruction method.

The effectiveness of CSD reconstruction is related to the quality of the data used for training. In simulations, the training was done with single fibers of different geometrical properties and rotated across the entire pdf-space. This is a complete training set for characterizing multiple fibers as a lineal combination of single fibers. The CS reconstruction using data-driven sparse dictionary provided better results for the different settings and data when compared to MAP or CS using the identity. Quantitatively speaking, MAP and CSI obtained 10 and 50% more NMSE as compared to CSD. The observed reconstruction quality from the three reconstruction methods did not affect in a significant manner the estimation of propagator-based diffusion indices. The extraction of MSD and p0 were quite stable for different noise levels and crossing angles; nonetheless, MAP and CSI deviated from the ground truth, in both reconstruction quality and extraction of the propagator-based diffusion indices, at undersampling factors above 4x. (Hutchinson et al., 2017) justifies similar behavior on its results due to sampling dependency mainly, but in these experiments it may be more likely that the requirements for a robust optimization were not met for those undersampling factors and the reconstruction error propagated to the diffusion indices. In the in vivo data, it is difficult to assure how rich or representative the slice used for training the dictionary was. The difference in results about the extraction of propagator-based diffusion indices between simulations and in vivo data highlights the disadvantage of CSD in terms of determining how rich was the data used for training. Because q-space exploration is far from being defined, it is difficult to establish which dictionary should be used if the specific parameters of the acquisition are not known. Furthermore, the dictionary could change across healthy and patient populations, or across age groups (Bilgic et al., 2012). One alternative is concatenating dictionaries from different acquisition schemes or objective-designed dictionaries.

MAP reconstruction gave higher quality in the extraction of propagator-based diffusion indices. The errors from the three methods were below five to eight percent, although MAP was the best; followed by CSD reconstruction and then CSI. The superior extraction of diffusion indices from MAP could be important if the objective of the study is related to the characterization of microstructure. Future work could include further propagator-based

quality measures, such as the angular error in crossing angle to complement NMSE, PC, MSD and p0; in order to increase the characterization of the propagator reconstruction quantitatively.

Finally, this comparison could be improved by using recent advances of the corresponding methods, like the Laplacian regularized MAP (Fick et al., 2016), joint k-q reconstruction (Ourselin, Alexander, Westin, & Cardoso, 2015), or the implementation of deep learning for CSD like in (Rasmussen et al., 2018).

4. MAPCS: Q-SPACE RECONSTRUCTION USING MEAN APPARENT PROPAGATOR AND COMPRESSED SENSING

This Manuscript was submitted to the journal Magnetic Resonance in Medicine (Magn Reson Med) in July, 2019.

4.1. Introduction

The diffusion propagator (Callaghan, 1991, 338 p.) is the complete characterisation of the diffusion phenomenon for a given diffusion time. From the complete information, it is possible to extract more sensitive diffusion-based indices that could be used as biomarkers in order to improve early detection, diagnosis and follow-up treatments for neurodegenerative diseases. However, the acquisition and reconstruction of the diffusion propagator is, most of the time, not easy. In Diffusion Spectrum Imaging (DSI) (Wedeen et al., 2000, 2005), the diffusion propagator is obtained by directly taking the Fourier transform to q-space. However, this would require a long acquisition time to acquire all the q-space samples, which is unfeasible in a clinical setting. A first option to reduce the acquisition time for a DSI protocol is by means of compressed sensing (CS) (Candes & Romberg, 2007). CS is an efficient sampling and reconstruction protocol based on randomness and sparse representation of information to recover a complete signal from a number of samples below the Nyquist criterion. Since the appearance of DSI-CS (Menzel et al., 2011), the main idea has been to identify the optimal sparsifying basis to compress the diffusion propagator and reconstruct it from a reduced number of q-space samples for feasible clinical acquisitions (Bilgic et al., 2012; Paquette et al., 2015).

A second option is by means of physical, signal-driven or biophysical models (Hutchinson et al., 2017). Models approximate the diffusion propagator and provide diffusion-based indices for specific clinical applications. By far the most common and useful approximation is the Gaussian model used in diffusion tensor imaging (DTI) (Basser et al., 1994; Le Bihan et al., 2001). The Gaussian model greatly reduces the acquisition time and the required number of q-space samples. However, in order to capture non-Gaussian diffusion, the model has to be sophisticated at the cost of more q-space samples. We seek a model that describes the complete q-space, using an efficient set of continuous basis-functions such that the diffusion signal is well represented with only a few coefficients and provides an

appropriate sparsifying bases to be used in CS. Depending on the model, CS can improve the trade-off between the fit to q-space data and the required number of the q-space samples for the proper reconstruction of the diffusion propagator. One such model is the simple harmonic oscillator-based reconstruction and estimation (SHORE) model (Ozarslan et al., 2013). Recently, we compared SHORE-CS (S. L. Merlet & Deriche, 2013) with DSI-CS; and found that the DSI-CS reconstruction using the identity as sparsifying basis outperformed SHORE-CS (Tobisch et al., n.d.). However, in this work we explore another model that can benefit from the CS formalism: Mean Apparent Propagator (MAP)(Ozarslan et al., 2013; Avram et al., 2016; Fick et al., 2016; Benjamini, Komlosh, Williamson, & Basser, 2018).

MAP reconstructs the diffusion propagator from the q-space data using a set of continuous basis-functions based on Hermite functions. Up to some extent, MAP has a better characterisation of q-space than SHORE because of the anisotropic nature of the scaling tensor that uses to generate the basis-functions, whereas the isotropic scaling of SHORE would require more basis-functions to obtain the same reconstruction quality (Ozarslan et al., 2013; Fick et al., 2015). The idea of using CS with MAP comes from the realisation that its coefficients are sparse (Avram et al., 2016). Therefore, we propose to use an ℓ_1 regulariser in a CS setting for MAP. This allows to find a similar number of coefficients, but in a much wider range of basis-functions which would not be present without CS.

Theory

Mean Apparent Propagator (MAP) (Ozarslan et al., 2013; Avram et al., 2016; Fick et al., 2016; Benjamini et al., 2018) is a physical model whose basis-functions are Gauss-Hermite functions. The first basis-function corresponds to Gaussian diffusion as in DTI, and the following basis-functions correspond to non-Gaussian diffusion; hence, it can represent well the complete diffusion propagator.

An important property of the MAP basis-functions is that their Fourier transforms are also Gauss-Hermite functions (Ozarslan et al., 2013), so there is an analytical representation for either q-space, $E()$, or the diffusion propagator, $p(\mathbf{r})$. Furthermore, the nth-basis-function

in both spaces share the same coefficient c_n. This is convenient, because it allows to refer to both spaces in the same formulation. Therefore, we can write the Fourier relationship as,

$$E() \;\leftrightarrow\; p(\mathbf{r})$$
$$\textstyle\sum_{n=0}^{B(N_{\text{max}})} c_n \phi_n(\mathbf{A},) \;\leftrightarrow \textstyle\sum_{n=0}^{B(N_{\text{max}})} c_n \psi_n(\mathbf{A}, \mathbf{r}),$$

where $\mathbf{A}$ corresponds to the scaling tensor to generate both sets of basis-functions; and $\mathbf{r}$ correspond to each vector space; N_{max} corresponds to the maximum order of the basis-functions and $B(\alpha)$ is the total number of basis-functions from the maximum order. Since in general q-space is considered to be an even function, only even numbers of N_{max} are needed and the total number of basis-functions is $B(N_{\text{max}}) = (N_{\text{max}}+2)(N_{\text{max}}+4)(2N_{\text{max}}+3)/24$. Nevertheless, to consider flow and motion it would be necessary to include both odd and even orders (Benjamini et al., 2018). Using matrix notation, the bases are $\mathbf{\Phi}_{B,Q}$ and $\mathbf{\Psi}_{B,R}$; where B is the number of basis-functions, Q is the number of acquired q-space samples and R is an arbitrary number of r-space positions.

The MAP optimisation (Ozarslan et al., 2013; Avram et al., 2016) is:

$$\min_{\mathbf{c}} \|S\mathbf{\Phi}^{\text{T}}\mathbf{c} - \|_2^2 \quad \text{s.t} \quad \mathbf{\Psi}^{\text{T}}\mathbf{c} \geq \mathbf{0}, \tag{4.1}$$

where S is the under-sampling operator so that the size of $\mathbf{\Phi}^{\text{T}}$ corresponds to the number of acquired q-space samples . The constraint is to impose non-negativity in the diffusion propagator. For non-Gaussian diffusion, such as in crossing-fibres, the maximum order N_{max} is commonly set to six, meaning that MAP uses $B(6) = 50$ basis-functions for the optimisation (Ozarslan et al., 2013; Avram et al., 2016). The number of basis-functions is a trade-off between the expected fit and the number of q-space samples for a robust optimisation.

In (Fick et al., 2016) it is proposed to remove the non-negativity constraint and to add a Laplacian regulariser. This formulation, MAPL, promotes smoothness:

$$\min_{\mathbf{c}} \|S\mathbf{\Phi}^{\text{T}}\mathbf{c} - \|_2^2 + \lambda_{\text{L}}\mathbf{c}^{\text{T}}\mathbf{U}\mathbf{c}, \tag{4.2}$$

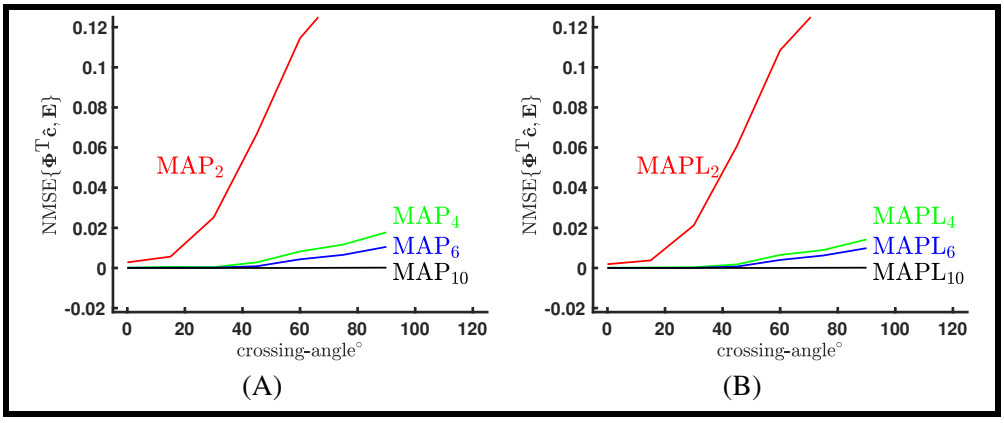

FIGURE 4.1. Reconstruction error from MAP (panels (A)) and MAPL (panels (B)) in noise-free simulations of two crossing-fibres. They show the normalised mean squared error (NMSE) as a function of the crossing-angle for different number of maximum orders. Differences in reconstruction quality are reduced as the number of the maximum order increases.

where $\mathbf{U}$ is the Laplacian matrix based on the scaling tensor $\mathbf{A}$ and the order from each basis-function (Fick et al., 2016). This formulation provides a closed-form optimum such that $\hat{\boldsymbol{c}} = (\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathrm{T}} + \lambda_{\mathrm{L}}\mathbf{U})^{-1}\boldsymbol{\Phi}\mathbf{y}$, where λ_{L} is a tuning parameter. This parameter is obtained by the generalised cross validation algorithm (Fick et al., 2016). MAPL has shown superior performance than MAP and the Laplacian-regularised modified Spherical Polar Fourier basis (Fick et al., 2016). Additionally, the estimated coefficients from MAPL have been successfully used to extrapolate the diffusion signal and provide some type of denoising procedure as pre-processing step for microstructural modeling (Fick et al., 2016).

In Figure 4.1 we show how the quality of the reconstruction improves as the maximum order of the basis increases, particularly for crossing angles close to 90. The NMSE was computed with MAP and MAPL in noise-free simulations of two crossing-fibres using different number of orders. However, in order to use more basis-functions, more q-space samples are required and the acquisition becomes longer. This is the motivation for using CS as a way to increase the order without lengthening the scan.

Given that MAP coefficients have a sparse representation for complex diffusion signals (Avram et al., 2016), we propose to include an ℓ_1 regulariser in order to improve the reconstruction quality by adding more basis-functions without increasing the number of q-space samples. The ℓ_1-regularised MAP, called MAPCS, is:

$$\min_{\mathbf{c}} \|S\boldsymbol{\Phi}^{\mathrm{T}}\mathbf{c} - \|_2^2 + \lambda_{\mathrm{cs}}\|\mathbf{c}\|_1. \tag{4.3}$$

The ℓ_1 regulariser allows a sparse selection of coefficients from a larger set of basis-functions. The balance between the regulariser and the data consistency term is defined by the Lagrange weighting λ_{cs}. In (Fick et al., 2018), compressed sensing was used to make a time-sensitive MAPL with the first fifty basis-functions that jointly reconstructs q-space for multiple diffusion times. However, in this work we use CS to improve the trade-off between the fit to q-space data and the required number of q-space samples for the proper reconstruction of the diffusion propagator.

4.2. Methods

Implementation

The implementation requires first to fit a diffusion tensor for scaling, then to compute the basis-functions and finally to perform the specific optimisation for the three techniques. First, from the available q-space samples we fit the diffusion tensor model:

$$E_{\mathrm{model}}(\mathbf{q}, \tau) = C + \exp(-4\pi^2\tau\mathbf{D}^{\mathrm{T}}), \tag{4.4}$$

where τ is the effective diffusion time from the diffusion encoding ($\tau = \Delta - \delta/3$ ms) and $\mathbf{D}$ is the diffusion tensor. C is a positive constant to characterise the noise-floor from the Rician-distributed magnitude of the diffusion signal at low SNR (usually the case for in vivo acquisitions for diffusion spectrum imaging). The optimisation is done using a non-linear least-squares optimisation in the eigen decomposition of the diffusion tensor, $\mathbf{D} = \boldsymbol{R}(\boldsymbol{\theta})\Lambda(\boldsymbol{\lambda})\boldsymbol{R}(\boldsymbol{\theta})^{\mathrm{T}}$, to enforce non-negative eigenvalues. We set $\lambda_{\{1,2,3\}} \in [0, \infty]$; and

we constrain the angles as $\theta_1 \in [0, \pi]$ and $\theta_{\{2,3\}} \in [-\pi, \pi]$. Second, we compute the scaling tensor $\mathbf{A}$,

$$\mathbf{A} = 2\Lambda\tau = \begin{bmatrix} \mu_1^2 & 0 & 0 \\ 0 & \mu_2^2 & 0 \\ 0 & 0 & \mu_3^2 \end{bmatrix} \text{ with } \mu_{\{1,2,3\}} = \sqrt{2\lambda_{\{1,2,3\}}\tau}, \tag{4.5}$$

in order to generate the basis-functions for MAP, MAPL and MAPCS. The techniques share the same basis-functions but they differ in which and how many are used.

We implement MAP, MAPL and MAPCS techniques in Matlab 2017 using CVX (Grant & Boyd, 2014, 2008). In particular, MAPCS (eq. (4.3)) is implemented as

$$\min_{\mathbf{c}} \|\mathbf{c}\|_1 \quad \text{s.t } \|S\boldsymbol{\Phi}^{\mathrm{T}}\mathbf{c} - \|_2^2 < \epsilon, \tag{4.6}$$

or,

$$\min_{\mathbf{c}} \|\mathbf{c}\|_1 \quad \text{s.t} \begin{cases} \|S\boldsymbol{\Phi}^{\mathrm{T}}\mathbf{c} - \|_2^2 < \epsilon \\ \boldsymbol{\Psi}^{\mathrm{T}}\mathbf{c} \geq \mathbf{0} \end{cases}, \tag{4.7}$$

where the optimisations pursuit a sparse representation of $\mathbf{c}$. The selection between eq. (4.6) and eq. (4.7) depends on the noise level. In other words, the selection of the optimisation comes from whether the constant C from eq. (4.4) is enabled to estimate the diffusion tensor or not. The data consistency term is relaxed by the parameter ϵ according to noise level, under-sampling, nature of the scaling tensor and number of basis-functions for the optimisation.

Datasets and analysis

We use 600 simulations from high-contrast q-space data from two 90° crossing-fibres as in (Tobisch et al., n.d.). These simulations have the same microstructure but different orientations in 3D space. We choose the 90° angle since it is the more demanding condition for the correct estimation of the scaling tensor. The simulations are made with 10^3 time steps for 10^5 spins to diffuse in the crossing substrate. We use the following parameters for the simulations: $\Delta = 65.9$ ms, $\delta = 57.4$ ms and $G_{\text{max}} = 26.79$ mT/m, where

$q_{\text{max}} = 65.47$ mm^{-1} (or $b_{\text{max}} = 7912$ s/mm^2); and TE $= 135.5$ ms and SNR of 20. We use these simulations to obtain the sample mean and the sample standard deviation from each quality index for the MAP, MAPL and MAPCS techniques.

For the in vivo acquisition we use a fully-sampled diffusion spectrum imaging scan from (Tobisch et al., 2018). Data are from a 3T Prisma scanner at 1.5 mm isotropic resolution (TE/TR = 105/6100 ms, bmax = 6800 s/mm ^{2}, Δ = 51.3 ms, δ = 20.1 ms) using a 64-channel head-neck coil. 257 diffusion weighted and 8 interleaved, non-weighted images are acquired with both anterior-to-posterior and posterior-to-anterior phase encoding. The acquired data are processed in FSL to estimate and correct for susceptibility geometric distortions, eddy currents and subject motion (Andersson et al., 2003; Andersson & Sotiropoulos, 2015). Finally, we also use a denoising procedure based on PCA decomposition to further improve the data at high q-space samples (Veraart, Fieremans, & Novikov, 2016). The redundancy in multi-directional and multi-radial q-space samples can identify and reduce the contribution of noise in the diffusion signal. We use the implementation from MRtrix (`http://www.mrtrix.org`).

Simulations and in vivo datasets are acquired in an $11 \times 11 \times 11$ q-space Cartesian grid. These complete datasets have 258 q-space samples contained in a discrete half-sphere with a radius of 5 samples and are the ground truth in our experiments. We use retrospective under-sampling, isotropic sparse distributions as in (Tobisch et al., n.d.), in both datasets to generate the corresponding basis-functions to evaluate the reconstruction performance given a number of q-space samples and to evaluate the effect of the estimated coefficients in the coefficient-based diffusion indices. In simulations we use $[32, 48, 64, ..., 128]$ q-space samples in an isotropic sparse pattern to measure the under-sampling effect. MAP and MAPL optimisations are done with $N_{\text{max}} = 6$ (50 basis-functions) as in (Ozarslan et al., 2013; Fick et al., 2016); whereas our MAPCS optimisation is done with $N_{\text{max}} = 10$ (161 basis-functions). For MAP and MAPL the number of basis-functions cannot be greater than the number of q-space samples. The advantage of MAPCS is that the limit of required samples is much lower. The constant C from eq. (4.4) is disabled because the high contrast q-space simulation reaches $E(\mathbf{q}, \tau) = 0$ in most directions. Finally, in eq. (4.6) we

use directly the noise from the magnitude and subtract it from the ground truth to estimate ϵ.

In the in vivo acquisition, we retrospectively choose 32 q-space samples with the isotropic sparse distribution. This is similar to the number of samples frequently used in diffusion tensor imaging. In the in vivo acquisition, the constant C from eq. (4.4) enables the characterisation of the noise-floor (Farooq et al., 2016) which is equivalent to add a one-function, $1(\mathbf{q}) = 1\ \forall \mathbf{q}$, as an extra basis-function. In this realistic clinical acquisition, the number of basis-functions needs to be reduced. For MAP and MAPL we use 8 basis-functions, corresponding to $N_{\text{max}} = 2$ plus the one-function; whereas MAPCS uses 51 basis-functions, corresponding to $N_{\text{max}} = 6$ plus the one-function. The maximum orders are defined after exhaustIve inspections for robust optimisations in most voxels along a coronal slice. In eq. (4.7) we use the residual from MAP optimisation using 8 basis-functions to estimate ϵ. Although we do not know the noise-free and artefact-free ground truth from the in vivo acquisition, it is possible to analyse the similarity between the fully-sampled data and the reconstruction from under-sampled data. Finally, from the known centrum semiovale it is possible to discuss if the different optimisations are able to resolve crossing-fibres from the intersection of the corpus callosum, the corticospinal tract and the corona radiata neural networks (Tian et al., 2016). For that purpose, we use the coefficients from each optimisation to analytically generate each oriented distribution function (Ozarslan et al., 2013; Avram et al., 2016).

Quality indices

To quantitatively evaluate each q-space reconstruction we use the normalised mean squared error (NMSE) and the Pearson's correlation coefficient (PC) (Paquette et al., 2015). The NMSE is defined as,

$$\text{NMSE}\{\boldsymbol{\Phi}^{\text{T}}\hat{c},\} = \frac{\|\boldsymbol{\Phi}^{\text{T}}\hat{c} - \|_2^2}{\|\|_2^2},$$

where $\mathbf{\Phi}^{\mathrm{T}}\hat{c}$ is the reconstructed q-space from the estimated coefficients and is the corresponding dataset that we use as ground truth. The PC is defined as,

$$\mathrm{PC}\{\mathbf{\Phi}^{\mathrm{T}}\hat{c},\} = \frac{1}{n-1}\mathbf{Z}_{\mathbf{\Phi}^{\mathrm{T}}\hat{c}} \cdot \mathbf{Z}^{T}, \quad \text{with } \mathbf{Z}_{\mathbf{x}} = \frac{\mathbf{x} - \mathrm{mean}(\mathbf{x})}{\mathrm{std}(\mathbf{x})},$$

where n is the number of discrete samples in q-space, $\mathbf{Z}$ is the vector of standard score of each sample and std stands for standard deviation.

After each reconstruction, we compute coefficient-based diffusion indices: non-Gaussianity (NG), return to zero probability (p0) and the mean squared displacement (MSD) (Ozarslan et al., 2013; Avram et al., 2016; Fick et al., 2016) and calculate the NMSE of each index with respect to the same indices computed from the fully-sampled and noise-free 90° crossing-fibres. The properties from the microstructure are the same regardless of the orientation.

4.3. Results

Simulations

Figure 4.2 shows the sample mean and the sample standard deviation of the reconstruction quality indices. Panel (A) shows the normalised mean squared error (NMSE); and panel (B) shows the Pearson's correlation coefficient (PC), as a function of the number of q-space samples from the reconstructions for the three techniques. MAP and MAPL use a maximum order $N_{\text{max}} = 6$ whilst MAPCS uses $N_{\text{max}} = 10$. Means and standard deviations above 10% are removed for a clearer visualisation. The results indicate that MAPCS improves both the NMSE and the PC in up to 3% for the complete range of q-space samples that we use in the simulations. Furthermore, the standard deviation of NMSE for MAPCS reconstructions using less than 64 q-space samples is smaller than for MAP and MAPL.

Figure 4.3 shows the NMSE for the coefficient-based diffusion indices of the different

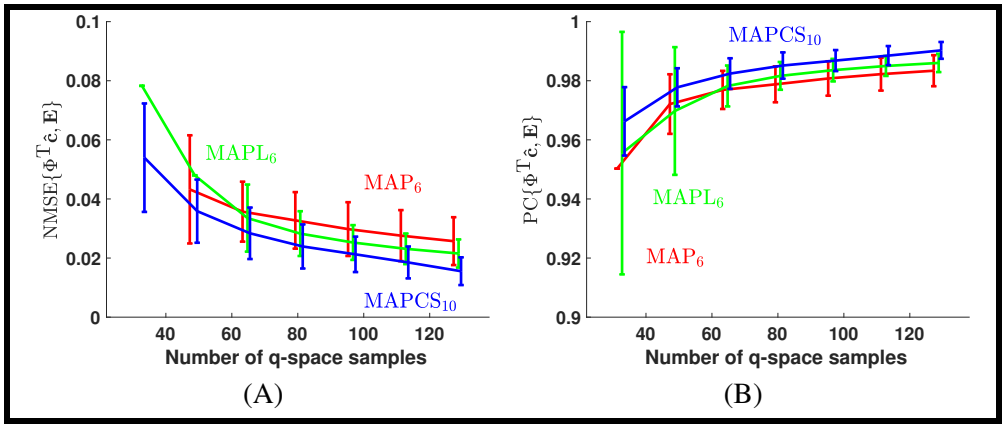

FIGURE 4.2. Reconstruction quality indices of the different techniques as a function of the number of q-space samples at SNR of 20. MAP and MAPL consider a maximum order $N_{\text{max}} = 6$; and MAPCS considers a maximum order $N_{\text{max}} = 10$. The normalised mean squared error (NMSE) in panel (A) and the Pearson's correlation coefficient (PC) in panel (B) indicate that MAPCS produces better reconstructions in comparison to MAP and MAPL.

optimisations. Means and standard deviations above 10% are removed for a clearer visualisation. The results indicate that MAPCS performs well for these indices, keeping an error no greater than 2% in comparison with MAP which is the one that performs better. These results are expected, since using more basis-functions (MAPCS) improves the fit to the data, but an unwanted effect is that it also reduces the quality of indices that depend on the smoothness of the propagator.

In vivo acquisition

To show the in vivo results we compare them to the fully-sampled reconstruction in a coronal slice that contains part of the centrum semiovale. Figure 4.4 shows the spatial maps for reconstruction quality indices as for coefficient-based diffusion indices. Panels (A) and (B) show the spatial maps of the NMSE and (1-PC) for MAP, MAPCS and MAPL from 32 q-space samples. For visual clarity the PC index has been subtracted from one, such that lower values indicate greater correlation to the fully-sampled data. In general terms, the three techniques show reconstruction quality below ten percent error; however, MAPCS

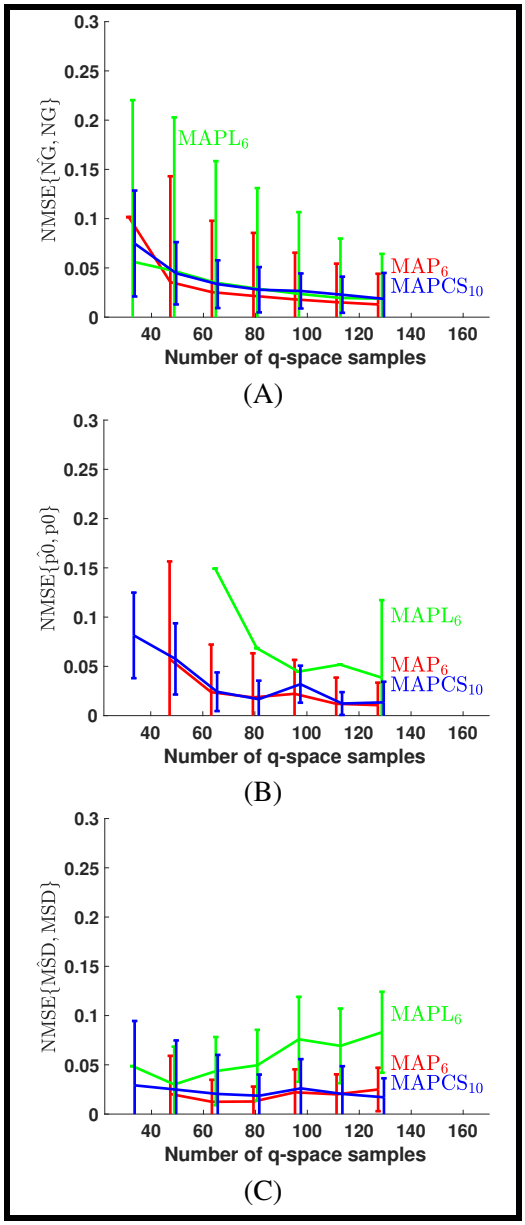

FIGURE 4.3. Normalised mean squared error (NMSE) from coefficient-based diffusion indices of the different optimisations as a function of the number of q-space samples at SNR of 20. MAP and MAPL consider a maximum order $N_{\text{max}} = 6$; and MAPCS considers a maximum order $N_{\text{max}} = 10$. Panel (A) is the NMSE of the non-Gaussianity (NG) index. Panel (B) is the NMSE of the return to zero probability (p0) index; and panel (C) is the NMSE of mean squared displacement (MSD) index.

reduces the reconstruction error and has a more homogenous reconstruction quality along the coronal slice (Panels (A) and (B)). MAP and MAPL show higher reconstruction errors around the centrum semiovale. This may be explained because most of the energy from the coefficients is concentrated in the basis-function which represents the Gaussian diffusion and it is where the tensor model is known to be inaccurate due to the multiple crossing-fibres. Panels (C), (D) and (E) show the spatial maps for the non-Gaussianity (NG), the return to zero probability (p0) and the mean squared displacement (MSD) indices. In panel (C), the NG spatial maps indirectly show how the optimisations distribute the energy from the coefficients along their respective bases. In MAP and MAPL ($N_{\text{max}} = 2$), the energy is mostly concentrated in the Gaussian coefficient and it reduces its correspondence with the anatomy in the coronal slice. MAPCS ($N_{\text{max}} = 6$) shows a better visual correlation with the expected anatomy. Additionally, the higher maximum order raises the NG index. In panel (D), the p0 spatial maps show higher visual correlation between each method. Finally, panel (E) shows the MSD spatial maps. MAP and MAPCS MSD maps differ from the MSD map obtained from MAPL estimated coefficients.

Finally, figure 4.5 shows a magnified area around the centrum semiovale to analyse the field of oriented distribution functions (ODFs). Panel (A) shows the ODFs from MAP using $N_{\text{max}} = 2$, panel (B) shows the ODFs from MAPCS using $N_{\text{max}} = 6$, and panel (C) shows the ODFs from MAPL using $N_{\text{max}} = 2$. These results show how an increased number of basis-functions is relevant to resolve crossing-fibres (indicated with the red arrow), which could not be done without CS for 32 q-space samples.

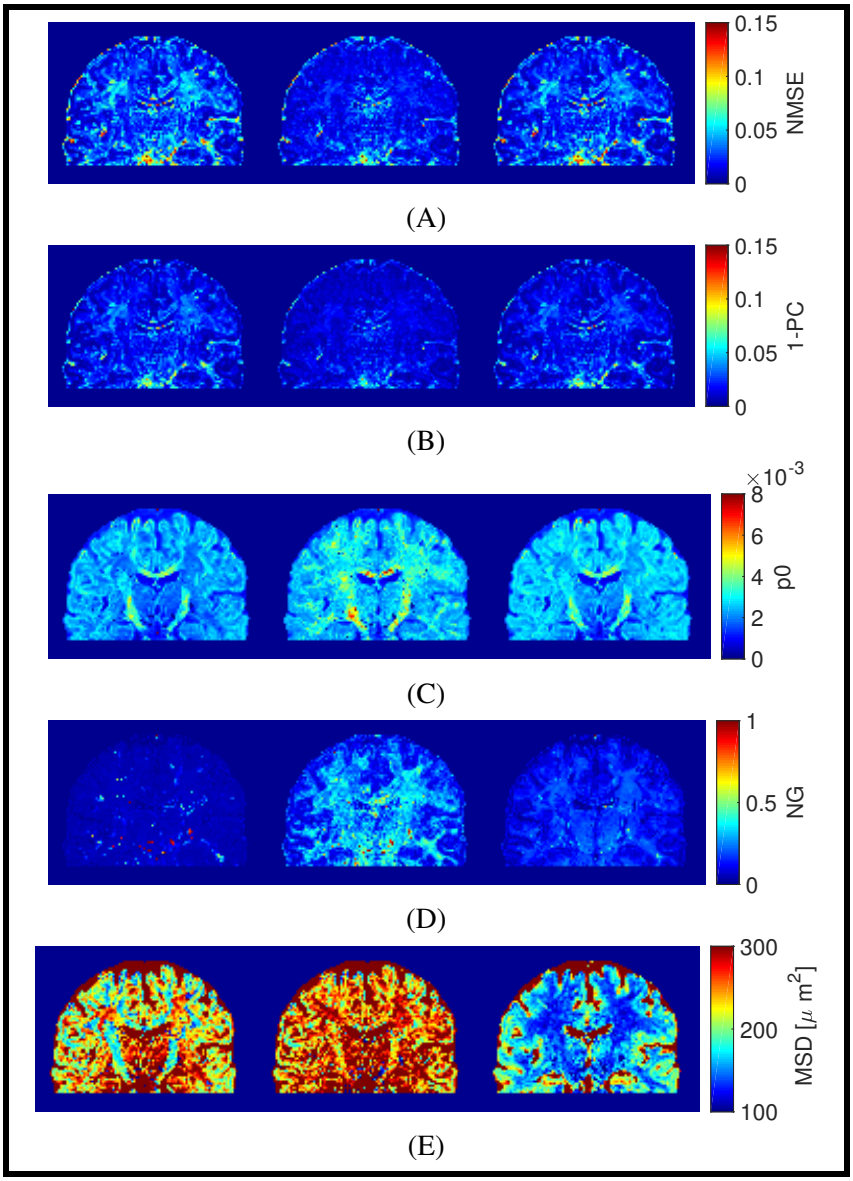

FIGURE 4.4. Spatial maps of reconstruction quality indices and coefficient-based diffusion indices from in vivo reconstructions. Panel A corresponds to the NMSE of the data; Panel B is (1-PC); Panel C is the Non-Gaussianity (NG); Panel D is p0; and Panel E is the MSD. The columns are the three techniques: MAP $_2$, MAPCS $_6$ and MAPL $_2$, in that order.

4.4. Discussion

In this work, we propose the use of the ℓ_1-regulariser for MAP optimisation, named MAPCS, to improve the trade-off between the number of basis-functions and the required

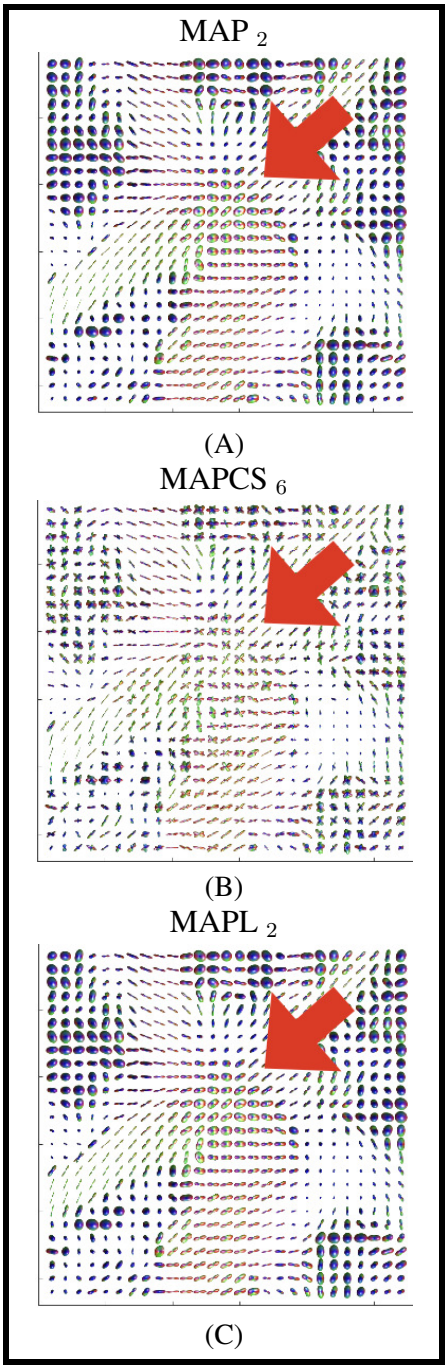

FIGURE 4.5. Visualisation of oriented distribution functions (ODFs) around the centrum semiovale from MAP, MAPCS and MAPL optimisations using just 32 q-space samples. Panel (A) shows the ODFs from MAP using a maximum order 2 (which results in 8 basis-functions) . Panel (C) shows the ODFs from MAPL with the same number of basis-functions. Panel (B) shows the ODFs from MAPCS using a maximum order 6 which results in 51 basis-functions. The red arrows highlight a region where multiple fibres intersect.

number of q-space samples for a robust optimisation. To evaluate our proposal, we compare MAPCS to the existing MAP and MAPL optimisations. We test the three methods in terms of reconstruction quality in simulated and in vivo data. In this work we use the Cartesian isotropic distribution from (Tobisch et al., n.d.) which has previously demonstrated to improve CS reconstruction for diffusion spectrum imaging data. But, the technique is not restricted to Cartesian sampling patterns, and it may be used with other patters such as multi-shell.

The simulations correspond to high-contrast q-space data of two fibres with a 90° crossing-angle, where the scaling tensor is isotropic and more basis-functions are required to preserve reconstruction quality. Although 90° crossing-angles seldomly appear in in vivo acquisitions, the main objective is to highlight how compressed sensing improves MAP by improving this trade-off and makes the method less sensitive to the data-driven scaling tensor found in (Fick et al., 2016; Hutchinson et al., 2017). In an in vivo acquisition we use minimum sampling to demonstrate complete q-space reconstruction with a similar number of samples as diffusion tensor imaging, but with the ability to correctly resolve crossing-fibres. Hence, Gaussian and non-Gaussian information from the diffusion signal are attainable in realistic clinical acquisition times.

MAP is a reconstruction that uses a least-squares optimisation subject to a non-negativity constraint. We notice that the constraint prevents over-fitting to noise, which is very convenient for high q-space samples or under-sampled acquisitions where SNR is low. This constraint implies a direct evaluation at the corresponding $\mathbf{r}$ positions; and for assuring it in the complete r-space a high computational load is required. On the other hand, MAPL removes the non-negativity constraint and includes a Laplacian regulariser which imposes smoothness and it avoids over-fitting to noise as the non-negativity contraint. MAPL demands more q-space samples to preserve reconstruction quality. MAPL also requires the estimation of λ_{L} by means of the generalised cross validation method which may be time-consuming voxelwise, but a fixed λ_{L} provides an optimisation as fast as a least-squares optimisation (Fick et al., 2016). Finally, reconstruction quality is reduced when the scaling

tensor is isotropic and more basis-functions are needed for both MAP and MAPL techniques.

Given the sparse representation of the coefficients (Avram et al., 2016), we propose an ℓ_1-regularised MAP optimisation based on compressed sensing (CS). By means of CS, it is feasible to select a sparse set of coefficients from a larger basis without sampling penalisation; or in other words, we improve the trade-off between the number of basis-functions and q-space samples. Our proposal requires the correct estimation of the Lagrange weighting to relax the data consistency term and avoid over-fitting to noise. This parameter is possible to estimate since it is related to SNR, under-sampling and the number of basis-functions. Finally, even though is not shown in here, the ℓ_1 regulariser and its estimated coefficients allow the correct extrapolation of q-space data. It can replace the Laplacian regulariser for that purpose without imposing excessive smoothness in the q-space data.

In simulations, the experiments show that MAPCS achieves better reconstructions in terms of the similarity of the propagators (data error and correlation) while maintaining a good estimation of the coefficient-based diffusion indices, very similar to those obtained with MAP. It is also clear that the excessive smoothing introduced by MAPL has a negative effect on the mean and standard deviation of these indices. The variance in the estimation of the coefficient-based diffusion indices is related to the variance of the estimated coefficients for each basis-functions, as reported in (Hutchinson et al., 2017).

In the in vivo reconstructions with just 32 q-space samples, we observe that MAPCS provides higher similarity to the fully-sampled data than MAP and MAPL. From the maps for the coefficient-based diffusion indices, we observe that the NG reports higher values and more contrast for MAPCS than MAP and MAPL. This es expected since the non-gaussianity is associated with more complex fiber configurations. In terms of p0 the spatial map do not differ significantly, which could be explained because p0 is associated to the integral of the data in the Fourier domain, and with only a few basis-function a good estimation of the integral can be obtained. The MSD spatial maps differ significantly between all three methods, since this parameter is very sensitive to the regularization or smoothing of the data. Finally, as can be seen in the ODFs, the region around the centrum semiovale

shows that a high number of basis-functions is relevant to properly describe crossing-fibres. A maximum order of two (8 basis-functions) is not enough to capture them, and with 32 q-samples it is not possible to use more in MAP or MAPL. CS provides a mean to maintain the recommended 50 basis-functions to describe that kind of microstructures with 32 q-samples, an acquisition length similar to that used for diffusion tensor. In summary, MAPCS may be used to describe Gaussian and non-Gaussian diffusion in realistic acquisition times.

5. CONCLUSIONS AND FUTURE WORK

This thesis investigated different strategies to reconstruct the complete diffusion propagator from minimum q-space sampling. The objective is to reduce the scan time while still being able to obtain the complete diffusion propagator. The reconstruction of this propagator would allow new biomarkers that could detect and distinguish neurodegenerative diseases and their different pathological stages. With that purpose in mind, a thorough analysis about reconstruction performance from the different methods used along this thesis was done on simulations under different levels of under-sampling and noise and orientation of the microstructure. Additionally, the reconstruction methods were tested on in vivo data. We qualitatively evaluated the centrum semiovale, a region of the human brain known for its complex anatomy. From the different evaluations, the application of compressed sensing for the sparse fit of the set of continuous basis-functions, presented in here as MAPCS in chapter 3, appears to be the most promising method given the results obtained across the different works. Future work should analyse the reconstruction process whilst applying forward q-space under-sampling, because it decreases the signal-to-noise-ratio from the acquisition and it may influence the lower bound for q-space sampling. Additionally, tuning parameters and optimisation bounds have to be validated over different acquisition schemes and different diffusion encoding parameters.

References

Aharon, M., Elad, M., & Bruckstein, A. (2006, November). $rm K$-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation. *IEEE Transactions on Signal Processing*, *54*(11), 4311–4322. doi: 10.1109/TSP.2006.881199

Andersson, J. L. R., Skare, S., & Ashburner, J. (2003, October). How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage*, *20*(2), 870–888.

Andersson, J. L. R., & Sotiropoulos, S. N. (2015, November). Non-parametric representation and prediction of single- and multi-shell diffusion-weighted MRI data using Gaussian processes. *NeuroImage*, *122*, 166–176.

Assaf, Y., Mayk, A., & Cohen, Y. (2000). Displacement imaging of spinal cord using q-space diffusion-weighted mri. *Magnetic Resonance in Medicine*, *44*, 713 – 722.

Assemlal, H.-E., Tschumperl, D., Brun, L., & Siddiqi, K. (2011). Recent advances in diffusion mri modeling: Angular and radial reconstruction. *Medical image analysis*, *15*, 369–396.

Avram, A. V., Sarlls, J. E., Barnett, A. S., Ozarslan, E., Thomas, C., Irfanoglu, M. O., ... Basser, P. J. (2016, February). Clinical feasibility of using mean apparent propagator (MAP) MRI to characterize brain tissue microstructure. *NeuroImage*, *127*, 422–434. doi: 10.1016/j.neuroimage.2015.11.027

Baraniuk, R., Choi, H., Neelamani, R., & Ribeiro, v. e. a. (1993). *Rice wavelet toolbox 3.0.* (http://dsp.rice.edu/software)

Basser, P., Mattiello, J., & LeBihan, D. (1994, January). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, *66*(1), 259–267. Retrieved 2018-10-16, from http://linkinghub.elsevier.com/retrieve/pii/S0006349594807751

Benjamini, D., Komlosh, M. E., Williamson, N. H., & Basser, P. J. (2018). Generalized mean apparent propagator (GMAP) MRI to measure and image advective and dispersive flows in medicine and biology. *IEEE Transactions on Medical Imaging*, 1–1.

Bernstein, K. K. . Z. X., M.A. (2004). *Handbook of mri pulse sequences*. Academic Press.

Bilgic, B., Setsompop, K., Cohen-Adad, J., Wedeen, V., Wald, L. L., & Adalsteinsson, E. (2012). Accelerated diffusion spectrum imaging with compressed sensing using adaptive dictionaries. *Med Image Comput Assist Interv*, *15*, 1–9.

Callaghan, P. T. (1991, 338 p.) *Principles of nuclear magnetic resonance microscopy*. New York: Oxford University Press Inc.

Candes, E., & Romberg, J. (2007). Sparsity and incoherence in compressive sampling. *Inverse Problems*, *23*, 969–985.

Cook, P., Bai, Y., Nedjati-Gilani, S., Seunarine, K., & Hall, M. G. e. a. (2006). Camino: Open-source diffusion-mri reconstruction and processing. In (p. 2759). Seattle, USA.

Dong, Q., Welsh, R. C., Chenevert, T. L., Carlos, R. C., Maly-Sundgren, P., Gomez-Hassan, D. M., & Mukherji, S. K. (2004). Clinical applications of diffusion tensor imaging. *Journal of magnetic resonance imaging*, *19*, 6–18.

Farooq, H., Xu, J., Nam, J. W., Keefe, D. F., Yacoub, E., Georgiou, T., & Lenglet, C. (2016, December). Microstructure Imaging of Crossing (MIX) White Matter Fibers from diffusion MRI. *Scientific Reports*, *6*, 38927.

Ferizi, U., Schneider, T., Panagiotaki, E., Nedjati-Gilani, G., Zhang, H., Wheeler-Kingshott, C., & D., A. (2014). A ranking of diffusion mri compartment models with in vivo human brain data. *Magnetic Resonance Imaging*, *72*, 1785–1792.

Fick, R. H. J., Petiet, A., Santin, M., Philippe, A.-C., Lehericy, S., Deriche, R., & Wassermann, D. (2018, January). Non-parametric graphnet-regularized representation of dMRI in space and time. *Medical Image Analysis*, *43*, 37–53. Retrieved 2018-07-09, from `http://www.sciencedirect.com/science/article/pii/S1361841517301330` doi: 10.1016/j.media.2017.09.002

Fick, R. H. J., Wassermann, D., Caruyer, E., & Deriche, R. (2016, July). MAPL: Tissue microstructure estimation using Laplacian-regularized MAP-MRI and its application to HCP data. *NeuroImage*, *134*, 365–385. doi: 10.1016/j.neuroimage.2016.03.046

Fick, R. H. J., Zucchelli, M., Girard, G., Descoteaux, M., Menegaz, G., & Deriche, R. (2015, April). Using 3d-SHORE and MAP-MRI to obtain both tractography and microstructural constrast from a clinical DMRI acquisition. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)* (pp. 436–439).

Getreur, P. (2006). *Wavelet cdf9/7 implementation.* (http://www.getreuer.info/home/waveletcdf97)

Grant, M., & Boyd, S. (2008). Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, & H. Kimura (Eds.), *Recent advances in learning and control* (pp. 95–110). Springer-Verlag Limited. (`http://stanford.edu/~boyd/graph_dcp.html`)

Grant, M., & Boyd, S. (2014, March). *CVX: Matlab software for disciplined convex programming, version 2.1*. http://cvxr.com/cvx.

Hagmann, P., Jonasson, L., & Maeder, P. e. a. (2006). Understanding diffusion mr imaging techniques: From scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *RadioGraphics*, *26*, 205–224.

Hall, M., & Alexander, D. (2009). Convergence and parameter choice for monte-carlo simulations of diffusion mri. *IEEE Transactions on Medical Imaging*, *28*, 1354–1364.

Hutchinson, E. B., Avram, A. V., Irfanoglu, M. O., Koay, C. G., Barnett, A. S., Komlosh, M. E., ... Pierpaoli, C. (2017, November). Analysis of the effects of noise, DWI sampling, and value of assumed parameters in diffusion MRI models: Cross-Model Analysis of Noise and DWI Sampling. *Magnetic Resonance in Medicine*, *78*(5), 1767–1780. Retrieved 2018-10-16, from http://doi.wiley.com/10.1002/mrm.26575

Irarrazaval, P. (2000). Mcgraw - Hill Interamericana.

Jensen, J., Helpern, J., & Ramani, A. e. a. (2005). Diffusional kurtosis imaging: The quatinfication of non-gaussian water diffusion by means of magnetic resonance imaging. *Magnetice Resonance in Medicine*, *53*, 1432–1440.

Johansen-Berg, B. T., H. (2009). *Diffusion mri: From quantitative measurement to in-vivo neuroanatomy*. Academic Press.

Lacerda, L., Sperl, J., Menzel, M., Sprenger, T., Barker, G., & Dell'Acqua, F. (2016). Diffusion in realistic biophysical systems can lead to aliasing effects in diffusion spectrum imaging. *Magnetic Resonance in Medicine*, *76*, doi 10.1002/mrm.26080.

Landman, B., Bogivic, J., Wan, H., ElShahaby, F., Bazin, P., & Prince, J. (2012). Resolution of crossing fibers with constrained compressed sensing using diffusion tensor mri. *Neuroimage*, *59*, 2175–2186.

Le Bihan, D., Mangin, J. F., Poupon, C., Clark, C. A., Pappata, S., Molko, N., & Chabriat, H. (2001). Diffusion tensor imaging: concepts and applications. *Magnetic Resonance Imaging*, *13*, 534–546.

Lipton., M. (2008). Totally accessible mri. In (chap. Parallel Imaging: Acceleration with SENSE and SMASH). Springer.

Lori, N., Conturo, T., & Le Bihan, D. (2003). Definition of displacement probability and diffusion time in q-space magnetic resonance measurements that use finite-duration diffusion-encoding gradients. *Journal of Magnetic Resonance*, *165*, 185–195.

Menzel, M., Tan, E., Sperl, J., & King, K. e. a. (2011). Accelerated diffusion spectrum imaging in the human brain using compressed sensing. *Magnetic Resonance in Medicine*, *66*, 1226–1233.

Merlet, S. (2013). *Compressive sensing in diffusion mri* (Unpublished doctoral dissertation). Universite Nice Sophia Antipolis.

Merlet, S. L., & Deriche, R. (2013, July). Continuous diffusion signal, EAP and ODF estimation via Compressive Sensing in diffusion MRI. *Medical Image Analysis*, *17*(5), 556–572. doi: 10.1016/j.media.2013.02.010

Mitra, P., & Halperin, B. (1995). Effects of finite gradient-pulse widths in pulsed-field-gradient diffusion measurements. *Journal of Magnetic Resonance*, *113*, 94–101.

Moussavi-Biugui, A., Stieltjes, B., & Fritzsche, K. e. a. (2011). Novel spherical phantoms for q-ball imaging under in vivo conditions. *Magnetic Resonance in Medicine*, *65*, 190–194.

Ourselin, S., Alexander, D. C., Westin, C.-F., & Cardoso, M. J. (Eds.). (2015). *Information Processing in Medical Imaging: 24th International Conference, IPMI 2015, Sabhal Mor Ostaig, Isle of Skye, UK, June 28 - July 3, 2015, Proceedings*. Springer International Publishing.

Ozarslan, E., Koay, C., Shepherd, T., Komlosh, M., Irfanoglu, M., Pierpaoli, C., & Basser, P. (2013). Mean apparent propagator (map) mri: A novel diffusion imaging method for mapping tissue microstructure. *NeuroImage*, *78*, 16–32.

Paquette, M., Gilbert, G., & Descoteaux, M. (2016). Optimal dsi reconstruction paramenter recomendations: Better odfs and better connectivity. *NeuroImage*, *142*, 1–13.

Paquette, M., Merlet, S., Gilbert, G., Deriche, R., & Descoteaux, M. (2015). Comparison of sampling strategies and sparsifying transforms to improve compressed sensing diffusion spectrum imaging. *Magnetic Resonance in Medicine*, *73*, 401–416.

Rasmussen, K. G. B., Kristensen, M. J., Blendal, R. G., Ostergaard, L. R., Plocharski, M., O'Brien, K., ... Bollmann, S. (2018, March). DeepQSM - Using Deep Learning to Solve the Dipole Inversion for MRI Susceptibility Mapping. *bioRxiv*.

Stejskal, E. O., & Tanner, J. E. (1965). Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, *42*, 288–292.

Tian, Q., Rokem, A., Folkerth, R. D., Nummenmaa, A., Fan, Q., Edlow, B. L., & McNab, J. A. (2016, December). Q-space truncation and sampling in diffusion spectrum imaging. *Magnetic Resonance in Medicine*, *76*(6), 1750–1763. doi: 10.1002/mrm.26071

Tobisch, A., Schultz, T., Stirnberg, R., VarelaMattatall, G., Knutsson, H., Irarrzaval, P., & Stcker, T. (n.d.). Comparison of basis functions and q-space sampling schemes for robust compressed sensing reconstruction accelerating diffusion spectrum imaging. *NMR in Biomedicine*, *0*(0), e4055. Retrieved 2019-01-17, from `https://onlinelibrary.wiley.com/doi/abs/10.1002/nbm.4055` doi: 10.1002/nbm.4055

Tobisch, A., Stirnberg, R., Harms, R. L., Schultz, T., Roebroeck, A., Breteler, M. M. B., & Stöcker, T. (2018). Compressed Sensing Diffusion Spectrum Imaging for Accelerated Diffusion Microstructure MRI in Long-Term Population Imaging. *Frontiers in Neuroscience*, *12*.

Tobisch, A., Varela-Mattatall, G., Stirnberg, R., Knutsson, H., Schultz, T., Irarrazaval, P., & Stöcker, T. (2014). Isotropic q-space sampling distribution for compressed sensing in dsi. In *Proceedings of the 22th scientific meeting of the international society for magnetic resonance in medicine* (p. 4451).

Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., & Wedeen, V. (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, *48*, 577–582.

Veraart, J., Fieremans, E., & Novikov, D. S. (2016, November). Diffusion MRI noise mapping using random matrix theory: Diffusion MRI Noise Mapping. *Magnetic Resonance in Medicine*, *76*(5), 1582–1593.

Wedeen, V. J., Hagmann, P., Tseng, W.-Y. I., Reese, T. G., , & Weisskoff, R. M. (2005). Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance in Medicine*, *54*, 1377–1386.

Wedeen, V. J., Reese, T. G., Tuch, D. S., Weigel, M. R., Dou, J., Weiskoff, R. M., & Chessler, D. (2000). Mapping fiber orientation spectra in cerebral white matter with fourier-transform diffusion mri. In (p. 82). Milan, Italy.

Wu, Y., Field, A., & Alexander, A. (2008). Computation of diffusion function measures in q-space using magnetic resonance hybrid diffusion imaging. *IEEE Transactions on Medical Imaging*, *27*, 858 – 865.

Ye, J. C., Tak, S., Han, Y., & Park, H. W. (2007, April). Projection reconstruction MR imaging using FOCUSS. *Magnetic Resonance in Medicine*, *57*(4), 764–775.

APPENDIX A. SELECTION OF TUNING PARAMETER FOR COMPRESSED SENSING

Compressed sensing is the only method from the comparison that needs to tune a parameter (λ); and which may vary depending on the acquired data. In order to use equation A.1,

$$\hat{p}(\mathbf{r}) = \text{argmin}_{p(\mathbf{r})} \ \frac{1}{2}||S\mathscr{F}p(\mathbf{r}) - E(\mathbf{q})||_2^2 + \lambda||\Omega p(\mathbf{r})||_1, \tag{A.1}$$

we applied two different heuristics:

Simulations. First, we analyzed the ℓ-curve to obtain the λ that had a good balance between data consistency and the sparsity regularization. Second, we defined a threshold based on the NMSE in which the reconstruction error should be less than five percent error with regard to the ground truth. λ was changed until this was satisfied or the maximum number of iterations was exceeded. The maximum number of iterations was 6.

In vivo data. We repeated the heuristic from simulations on certain voxels along the coronal slice. Because the λ parameter was similar across the slice, we took the mean value from λs and applied it to the complete slice.

APPENDIX B. EVALUATION OF THE SAMPLING PATTERN

We evaluated 10 different sampling patterns with variable density distribution to analyze whether the behavior of the reconstruction quality was influenced by them. Figure B.1 shows the results of the NMSE with respect to noise for the different patterns. The 8th pattern is not included here since it is the one used in the main document. Figure B.2 shows the differences in reconstruction errors between sampling patterns specified in the rows and columns, $(\text{NMSE}_{\text{column}} - \text{NMSE}_{\text{row}})^2 \ / \ \text{NMSE}^2_{\text{row}}$. From the three methods, the DictCS method shows more similar results across the different patterns (panel C).

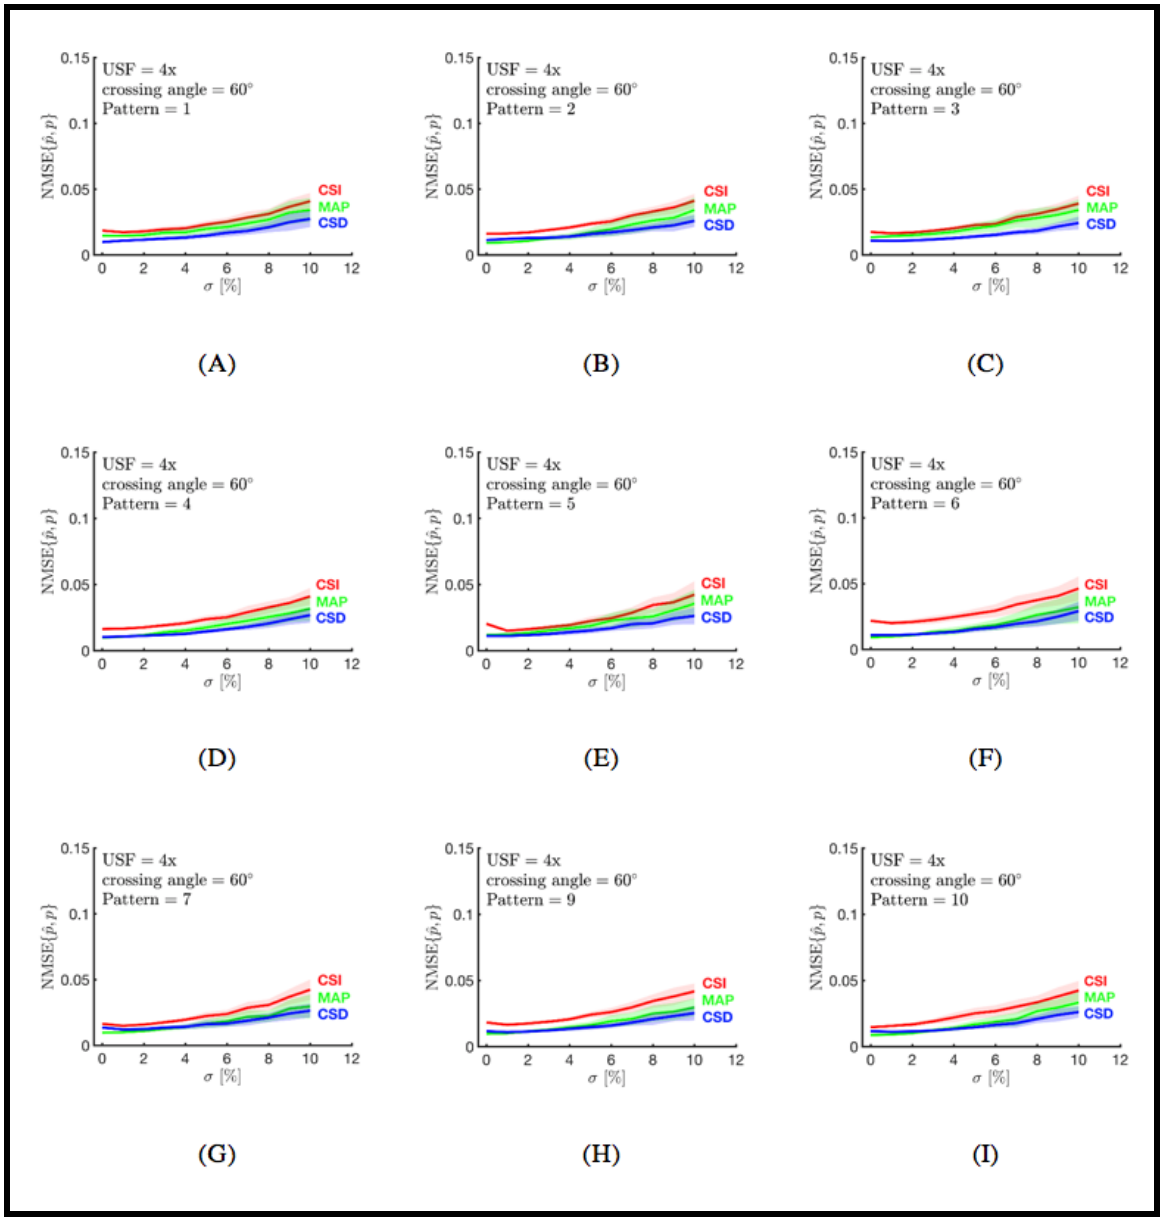

FIGURE B.1. NMSE as a function of noise level σ for 9 different sampling patterns.

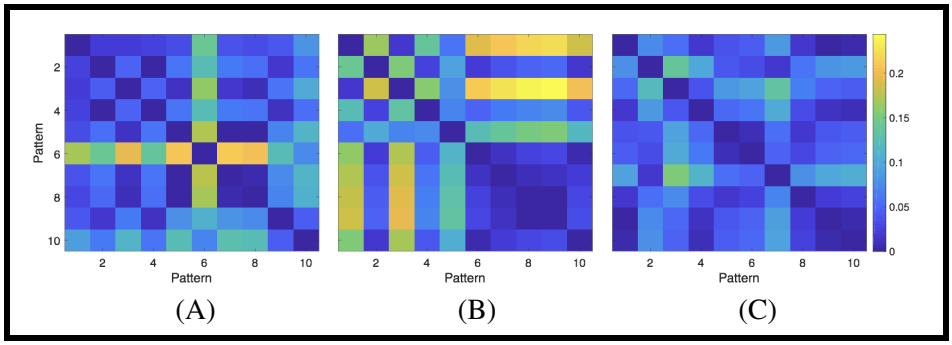

FIGURE B.2. NMSE difference matrix of the reconstruction methods. The colors indicate how much was the difference in reconstruction quality between the sampling patterns in the rows and columns of the matrix. Panel (A) corresponds to the NMSE difference matrix from CS reconstruction; panel (B) from MAP reconstruction; and panel (C) from DictCS reconstruction.

APPENDIX C. IN VIVO SPATIAL MAPS OF NMSE FOR CS, MAP AND DICTCS RECONSTRUCTIONS AT DIFFERENT UNDERSAMPLING FACTORS

Figure 8 shows the spatial maps of NMSE for the three reconstruction methods using USF = 4, 5, 6 and 8; as a way to validate how is the behavior of the reconstruction process in the in vivo data as the USF was increased.

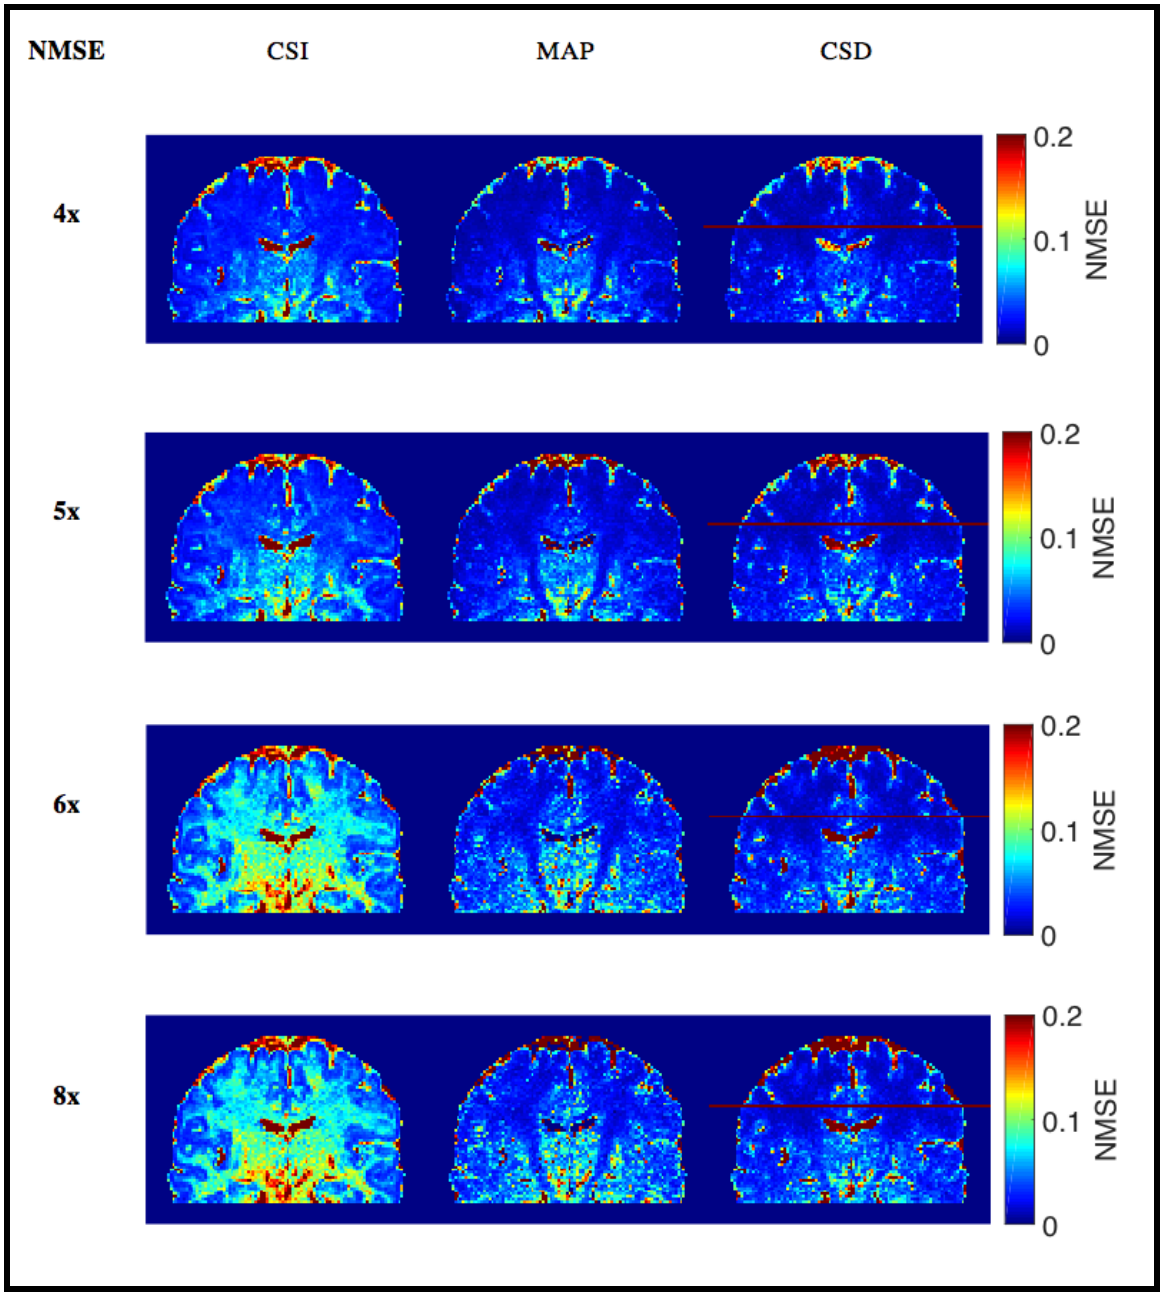

FIGURE C.1. NMSE spatial maps for CS, MAP and DictCS reconstruction methods at different undersampling factors factors