UNIVERSIDADE DE LISBOA
FACULDADE DE CIÊNCIAS
DEPARTAMENTO DE FÍSICA

Comparing single-shell and multi-shell free-water fraction estimation algorithms

Maria Oliveira Sobral Lopes

Mestrado em Engenharia Biomédica e Biofísica

Dissertação orientada por:
Rita Nunes
Raquel Conceição

2022

Abstract

Diffusion Weighted (DW) MRI is a medical imaging modality which can be used to model the displacement of water molecules as they diffuse through the brain, allowing the microstructural architecture of brain tissues to be explored *in vivo*. This technique has been widely applied to the study of many brain pathologies. However, the presence of extracellular free water affects the diffusion measurements, potentially leading to wrong interpretation about the underlying microstructural changes. Free-water elimination (FWE) is an alternative to more traditional approaches to model DWI data, which divides the signal into an extracellular compartment (which depending on tissue type can be isotropic) representing free water and another compartment representing tissue (usually anisotropic). A recent method by Neto Henriques, et al, to estimate free water fraction using multiple diffusion-weighting shells has been shown to reduce the bias in parameter estimates. However, as clinical protocols often use a single diffusion-weighting (single-shell data) to reduce exam times, it becomes relevant to investigate if introducing prior knowledge in the estimation, as proposed in the work of Pasternak, et al, could enable reliable free water elimination when applied to single-shell data.

The goal of this project is to compare the performance of these two free water elimination algorithms when applied to the same data. A large dataset of multi-shell DWI data was acquired as part of a longitudinal study led by the Norwegian University of Science and Technology (NTNU) in Trondheim. This dataset includes 78 healthy controls. The data was pre-processed and the multi-shell algorithm applied to eliminate free water contamination. In this project, the same data was processed after removal of the high diffusion-weighting shell, by applying an open-source implementation of the single-shell algorithm presented on the work of Golub, et al.

The methods used are fully detailed in this work, including the participants, image acquisition and image preprocessing phases. Tract-based spatial statistics (TBSS) were used for registration and alignment of images for all studied parameters and voxelwise statistics was performed in order to learn which voxels were significantly different between images processed with the two algorithms. White matter, cerebral cortex and subcortical masks were used to understand how the algorithms behave at a regional level.

The results include a comparison of the original data used, where differences can be observed and discussed. The statistical results - corrected p-value images for each parameter - are presented and discussed: considering the multishell algorithm as gold standard, for both fractional anysotropy (FA) and free water (FW) the singleshell algorithm seems to underestimate the white matter values and overestimate the gray matter values. For mean diffusivity (MD), it is the opposite: the single-shell algorithm seems to overestimate the white matter values and underestimate the gray matter values. These results are supported by further analysis, where each subject images for all parameters (FA, MD and FW) and for both algorithms (single and multi-shell) was used to get an average value of the voxels (excluding null values from this average), in order to understand how these values differ according to the algorithm used. These results, presented as boxplots for the regions being studied, also indicate that the values for both

algorithms are significantly different in almost all regions and parameters.

This study allows to understand the FWE-DTI application in single-shell data. The comparison with the multi-shell algorithm for FA and FW showed an underestimaion of WM values and overestimation of GM values. For MD, the values are conditioned by the prior and overestimated for all tissue type. Besides, the values obtained with the single-shell algorithm are considered significantly different from the ones obtained with multi-shell agorithm for both WM and GM in most parameters and regions.

Keywords

Diffusion Weighted MRI - Free-water - Single/Multi-shell - Mean diffusivity - Fractional anysotropy

Resumo

A ressonância magnética de difusão (*Diffusion-Weighted Magnetic Resonance Imaging*, dwMRI) é uma categoria de imagem médica que pode ser usada para modelar a deslocação de moléculas de água que se difundem pelo cérebro, permitindo a exploração da arquitetura microestrutural dos tecidos cerebrais *in vivo* - o contraste das imagem de difusão reflete a diferença da taxa de difusão entre os tecidos. Num meio livre, o movimento das moléculas não tem preferência direcional e é classificado como isotrópico. No entanto, no tecido biológico, as moléculas de água são constantemente restringidas por obstáculos como membranas celulares e macromoléculas - neste caso, a difusão é designada anisotrópica. A imagem por tensor de difusão (*diffusion tensor imaging*, DTI) é uma modificação da dwMRI que permite representar vetores correspondentes à intensidade e direção do movimento da água nos tecidos, para caracterizar os efeitos anisotrópicos da difusão no tecido biológico, o que não é possível com uma única imagem de difusão. Esta técnica de imagiologia tem sido amplamente aplicada ao estudo de muitas patologias cerebrais, nomeadamente a doença de Parkinson, doença de Alzheimer, alterações na substância branca após acidente vascular cerebral grave, esquizofrenia, tecido peritumoral, bem como diferenças na substância branca devidas à idade, entre outros.

O esquema de amostragem é um fator essencial que dita a qualidade de reconstrução da imagem de difusão. O termo "single-shell" refere-se a um esquema de amostragem uniforme (em formato esférico) de uma coleção de pontos com um único valor de ponderação da difusão ("b-value"); já o termo "multi-shell" implica 2 ou mais esferas. Embora as vantagens dos esquemas "multi-shell" com um b-value máximo mais alto ($\geq$ 2000s/mm^2) tenham já sido exploradas por autores como Christiaens et al, 2018, e estes esquemas sejam, em alguns casos, necessários, os protocolos clínicos impõem, por norma, um esquema de amostragem DTI típico "single-shell" com um único "b-value" ($\leq$ 1200s/mm^2) e cerca de 30 amostras. Neste sentido, é relevante investigar se a introdução de informação prévia no modelo proposto no trabalho de Pasternack, et al, 2009 poderia permitir a eliminação eficaz de água livre quando aplicada a dados "single-shell".

A presença de água livre (*free water*, FW) extracelular também pode afetar a qualidade dos parâmetros de difusão. A FW pode causar alterações nas medidas de difusão, e consequentemente levar a interpretações erradas sobre as mudanças microestruturais subjacentes. A eliminação de água livre é uma alternativa às abordagens mais tradicionais para modelação de dados de imagens de difusão, que divide o sinal num compartimento extracelular isotrópico - isto é, a água livre - e um compartimento anisotrópico, que representa o tecido.

Os algoritmos de estimativa/eliminação de FW para corrigir as medidas de difusão afetadas são hoje um tema em exploração. Uma maneira de abordar o problema é modificar o modelo DTI[1] (*diffusion tensor imaging*) padrão, largamente utilizado, e adicionar um segundo compartimento que representa a componente isotrópica - de forma a que a fração FW seja um parâmetro adicional a ser estimado. Hoy et

[1] Algoritmo padrão utilizado em imagens de difusão que permite a representação de vetores correspondentes à força e direção do movimento da água nos tecidos.

al., 2014, fizeram alteraçoes num modelo de regressão linear (*Weighted Linear Least Squares*) no sentido de torná-lo compatível com o modelo de dois compartimentos. Neto-Henriques et al., 2017 desenvolveu mais tarde uma implementação deste algoritmo. No entanto, este algoritmo aplica-se a dados "multi-shell", que são menos comuns em ambiente clínico, uma vez que implicam tempos de aquisição mais longos - se o número de direções por *shell* permanecer o mesmo, uma aquisição de dois levaria mais tempo. Portanto, a nível clínico, as aquisições "single-shell" (normalmente adquiridas com $b = 1000 s/mm^2$) são preferidas - o que levou Pasternak et al., 2009 a propor e implementar as estratégias para dados "single-shell": este algoritmo consiste em adicionar restrições espaciais aos dados e realizar um método Gradient Descent (GD) de forma a que seja possível obter uma solução única. Golub et. al, 2018 fez modificações nas várias etapas do algoritmo com o intuito de melhorar a robustez da técnica.

O objetivo deste projeto é comparar o desempenho desses dois algoritmos de eliminação de água livre (DTI para multi-shell, e o algoritmo para dados single-shell criado por Golub et. al) quando aplicados aos mesmos dados. Um grande conjunto de dados de difusão "multi-shell" foi adquirido como parte de um estudo pela Universidade Norueguesa de Ciência e Tecnologia (NUST) em Trondheim, Noruega. Este conjunto de dados inclui 78 controlos saudáveis. Os dados já foram pré-processados e o algoritmo "multi-shell" foi aplicado para eliminar a água livre. Neste projeto, os mesmos dados serão processados após a remoção da camada de alta difusão, através da implementação do algoritmo "single-shell" do trabalho de Golub et. al, 2018 . Os resultados do algoritmo "single-shell" serão comparados com os resultados do algoritmo "multi-shell" (atualmente o estado da arte na área) para os mesmos dados, a fim de avaliar a eficácia dos indicadores de difusão estimados pela abordagem "single-shell".

Os métodos utilizados estão detalhados neste trabalho, incluindo os participantes, a aquisição de imagens e o pré-processamento das imagens utilizadas. Uma ferramenta de estatística espacial (TBSS, *tract-based spatial statistic*) foi usada para registo e alinhamento de imagens para todos os parâmetros estudados e estatística aplicada a voxéis (*voxelwise statistics*) foi realizada com o intuito de identificar quais os voxéis considerados significativamente diferentes entre as imagens em que foi implementado o algoritmo "multi-shell" e as que foram corridas com o algoritmo "single-shell". Os pârametros estudados neste trabalho incluem: fração anisotrópica (FA, *fractional anysotropy*), água livre (FW) e difusividade média (MD, *mean diffusivity*). Máscaras da substância branca, córtex cerebral e região subcortical foram usadas para entender o comportamento dos algoritmos a nível regional.

Os resultados incluem uma comparação das imagens originais utilizadas, onde algumas diferenças são identificadas e discutidas. Os resultados estatísticos - imagens de valor-p (também chamado de nível descritivo ou probabilidade de significância) corrigidas para cada parâmetro - são apresentados e discutidos: considerando o algoritmo "multi-shell" como *gold-standard*, tanto para FA quanto para FW o algoritmo "single-shell" parece subestimar os valores da substância branca e sobrestimar os valores da substância cinzenta. Para MD, é o oposto: o algoritmo "single-shell" parece sobrestimar os valores da substância branca e subestimar os valores da substância cinzenta. Estes resultados são suportados por uma análise mais aprofundada, onde as imagens de cada sujeito para todos os parâmetros (FA, MD, FW) e para ambos os algoritmos ("single" e "multi-shell") foram usadas para obter uma média do valor dos voxéis, com o intuito de identificar quais os valores significativamente diferentes entre os dois algoritmos ("single" e "multi-shell"). Esses resultados, apresentados como boxplots (ou caixa de bigodes) para cada região em estudo, também parecem indicar que os valores para ambos os algoritmos são significativamente diferentes em quase todas as regiões e parâmetros.

Este estudo debruça-se sobre a aplicação de FWE-DTI em dados "single-shell". A comparação com o algoritmo "multi-shell" para FA e FW mostrou uma subestimação dos valores de matéria branca e sobrestimação dos valores de matéria cinzenta. Para MD, os valores são condicionados pelo "prior" e

sobrestimados para todos os tipos de tecido. Além disso, os valores obtidos com o algoritmo ”single-shell” são considerados significativamente diferentes dos obtidos com o algoritmo multi-shell tanto para matéria branca quanto para matéria cinzenta em quase todos os parâmetros e regiões.

Palavras-chave

Ressonância Magnética ponderada por difusão - Água livre - Single/Multi-shell - Difusividade média - Anisotropia fracionária

Acknowledgments

First, I would like thank my supervisors, Prof. Rita Gouveia Nunes and Prof. Marta Correia, for their availability, knowledge and patience that were essential for the development of this thesis and to Prof. Raquel Conceição, whose help and encouraging words eased the difficult days.

I also want to acknowledge the Department of Medicine at the University of Cambridge and thank Dr. Stefan Winzeck for helping me with the preprocessing of the data used in this study. Furthermore, I want to thank Marc Golub, one of the authors of the paper on which this work is based, for his availability and insights into some details of the algorithm he developed.

Finally, I want to thank the friends and family that supported me and cheered for the day I would succeed.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation, structure and goal of the current work

1.1.1 Motivation

The term free water (FW) refers to water molecules that do not flow and are not constrained by their surroundings. Free water is found in the human brain as cerebrospinal fluid (CSF), which is confined to the ventricles and the brain parenchyma. Due to processes such as tumors, brain trauma, or hemorrhage that cause ruptures in the blood–brain barrier, free water may also accumulate in the form of vasogenic edema within the brain parenchyma in the extracellular space [1][2].

CSF contamination is a type of partial volume effect that occurs in voxels shared by CSF and brain tissue along the contour lines of the ventricles and around the perimeters of the brain parenchyma [3][4]. As a result, CSF-contaminated voxels have decreased fractional anisotropy (FA) values. A white matter voxel contaminated by free water will most likely be represented by a relatively isotropic diffusion tensor, preventing it from being classified as white m atter. CSF contamination has been shown to disturb the delineation of fibers that run near the ventricles, such as the fornix, cingulum, and parts of the corpus callosum [5][6], and to impact the comparisons of DTI-related quantities using voxel-based and histogram analysis [7].

Both pathological neuro degeneration or healthy aging of the brain can cause ventricle augmentation and degeneration processes and increase partial volume effects (PVE). Thus, correcting for FW contamination is also critical in studies involving elderly subjects.

1.1.2 Structure

This work starts by introducing the teorical concepts necessary to understand the diffusion processes that happen in the brain and how they are used in magnetic ressonance imaging (MRI) - including the diffusion phenomenon, diffusion tensor imaging and diffusion tensor. The mathematics behind the scalar invariants that are used to characterize the anisotropic effects of diffusion on biological tissue is also explored. This is followed by the introducion of free water elimination DTI (FWE-DTI) and the alterations necessary to the mathematical equations behind the single-shell algorithm in order to separate the water and tissue components[1]. At this point, the state-of-the-art is presented and the context of evolution of FWE-DTI is further explained. Next, the difference between a single-shell and multi-shell schemes is introduced, followed by the introduction of regularized gradient descent fitting procedure

[1]In his work, Golub, et al, 2018, further explores the mathematical context which lead him to the implementation of the algorithm.

that allows to apply the FWE-DTI to single-shell data. This first introdutory chapter is followed by the methods used in this work, from the characerization of the participants of the dataset images, the details of image acquisition and processing and the process of statistical analysis used for the comparison of the algorithms.

The results include a comparison of the original data used, where differences can be observed and discussed. The statistical results - corrected p-value images for each parameter (FA, FW and MD) - are presented and discussed and are supported by further analysis, presented as boxplots for the regions being studied.

1.1.3 Goal

The goal of this project is to compare the performance of two free water elimination algorithms when applied to the same data. A large dataset of multi-shell DWI data has been acquired as part of a longitudinal traumatic brain injury study led by the NTNU in Trondheim [8]. This dataset includes 78 healthy controls whose data has already been pre-processed (further explained), and the multi-shell algorithm applied to eliminate free water contamination. In this project, the same data will be processed after removal of the high diffusion-weighting shell, by applying an open-source implementation of the single-shell algorithm [9]. The results will be compared to the output of the multi-shell algorithm (currently the state-of-the art in the field) for the same data, in order to assess the reliability of the diffusion metrics estimated by the single-shell approach.

1.2 Diffusion-Weighted Magnetic Resonance Imaging

1.2.1 Diffusion Phenomenom

The thermal motion of all particles (liquid or gas) at temperatures above absolute zero is designated molecular diffusion. Metabolism and respiration of living beings rely in part upon diffusion (in addition to bulk or active processes) given that in cell biology, diffusion is a main form of transport of necessary materials such as aminoacids within cells [10]. In a free medium, this movement of molecules has no directional preference and it is classified as isotropic, represented in Figure 1.1 A. However, in biological tissue, water molecules are constantly being hindered by obstacles such as cell membranes and macromolecules - in this case, the diffusion is anisotropic[2] and no longer Gaussian (depending on the observation time), represented in Figure 1.1 B. As a consequence, the diffusion coefficient (D) depends on the direction of measurement and the geometry of the underlying structures.

[2] Anisotropy is the property of a material which allows it to change or assume different properties in different directions

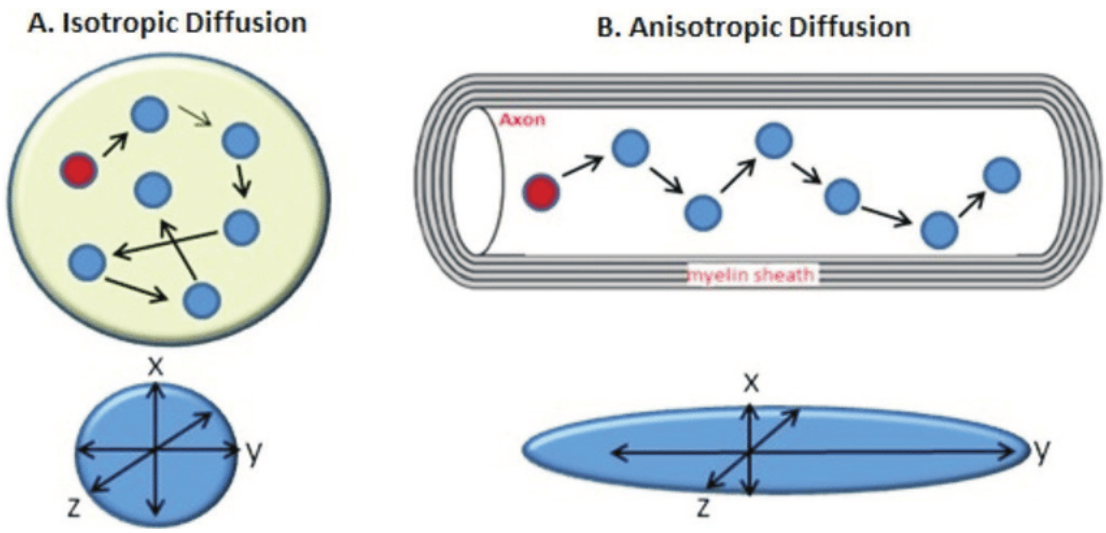

Figure 1.1: Molecular diffusion, also known as brownian motion, is the phenomenon tof thermally induced molecular displacement in a fluid. Water molecules are constantly moving. A: Diffusion is isotropic when motion is unconstrained, which means that motion occurs equally and randomly in all directions. B: When molecular motion is restricted, such as in white-matter tracts, diffusion is anisotropic, which means that motion occurs predominantly in one direction [11].

Unrestricted isotropic diffusion, such as that found in CSF-filled regions, is represented by a sphere with a radius close to $D_w = 3 \times 10^{-3} mm^2$; diffusion is still isotropic in randomly organized tissue, but it is limited by D<Dw (smaller sphere); in highly organized structures like WM, diffusion is represented by an elongated ellipsoid aligned with the orientation of the underlying fibre.

1.2.2 Diffusion Tensor Imaging

Diffusion-Weighted Magnetic Resonance Imaging (dwMRI) is a type of conventional MRI in which the diffusion of the water molecules is exploited to visualize internal microstruture. The image contrast in DWI reflects the difference in rate of diffusion between tissues [12].

Diffusion is the foundation of dwMRI. Because dwMRI is performed at millisecond intervals, it collects information at the micrometre scale, making it sensitive to microstructural changes in the brain. The average displacement of water molecules along a specified direction can be determined by applying two magnetic field gradient pulses that impose a phase on each molecule based on its displacement. A phase distribution is generated for a population of water molecules in a single voxel, which reflects the displacement distribution. Figure 1.2 ilustrates a spin-echo diffusion weighted image sequence:

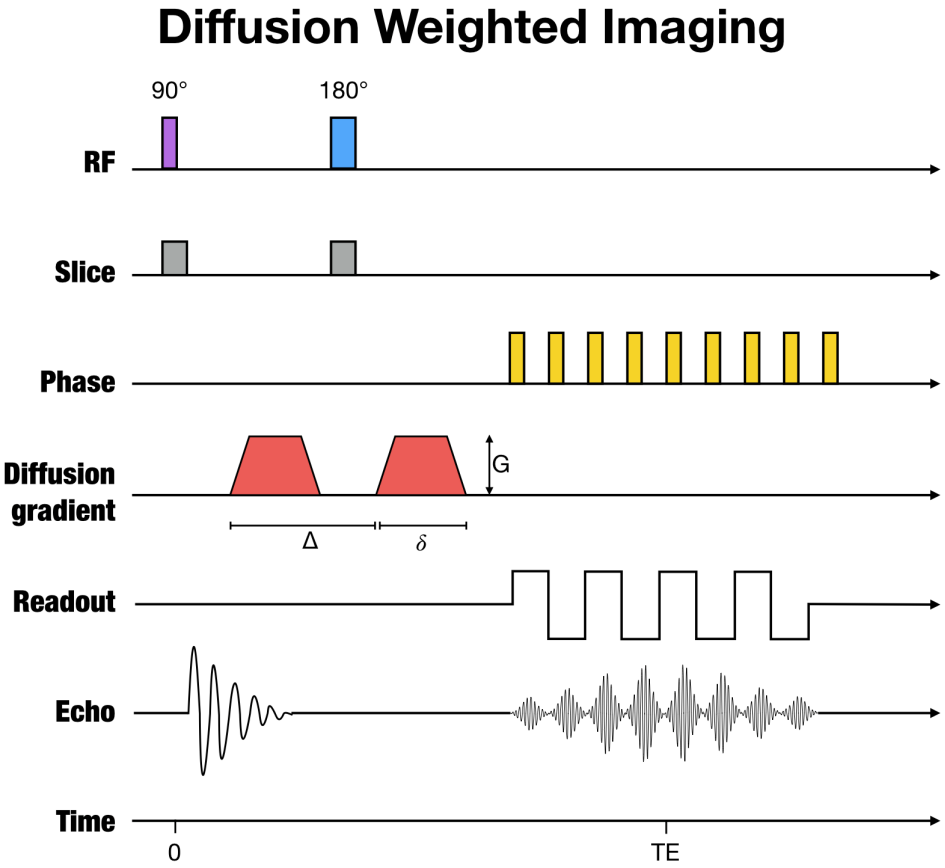

Figure 1.2: Simplified pulse diagram of a spin-echo diffusion-weighted image sequence. The diffusion gradients are shaded orange (note that they are both positive since the 180-degree pulse between them reverses the phase). The addition of equal, paired diffusion gradients to the standard spin-echo sequence causes moving protons to dephase. The degree of diffusion weighting depends on the strength of the gradient (amplitude G and duration δ) as well as the time spacing between them (Δ) - this is referred to as the b-value, further discussed. (Adapted from Assoc Prof Frank Gaillard, Radiopaedia.org, rID: 21753).

When water molecules are free to move, fast diffusion processes lead to a broader range of displacements - hence a wider distribution of phases, which results in higher loss of signal coherence and the radio frequency signal detected by the scanner is weaker, which creates a darker image. Conversely, if there is biological tissue constraining the water molecules, diffusion becomes slower, phase spread is narrower and less Gaussian, resulting in less signal attenuation and a brighter image. The images acquired are DW: each voxel has a signal attenuation that reflects the diffusion of water molecules in that voxel, which in turn is dependent on tissue properties and micro-structure. The contrast is determined by association of characteristics ilustrated in Fig. 1.2: the magnitude, G, the duration, δ, distance between the gradient pulses, Δ, and the proton gyromagnetic ratio, γ, into a single main parameter of acquisition, the b-value:

$$b = \gamma^2 \delta^2 G^2 (\Delta - \frac{\delta}{3}) \tag{1.1}$$

The attenuation of the signal can be predicted using the b-value:

$$S = S_0 * exp(-b * ADC) \tag{1.2}$$

Here, apparent diffusion coefficient (ADC) replaces D_w given that the measured diffusion depends on the acquisition parameters and on tissue barriers - it is not the true diffusion coefficient of water. S_0 is the signal with no diffusion sensitization (i.e signal for b = 0 s/mm^2).

1.2.3 Diffusion Tensor

Diffusion tensor imaging (DTI) is a modification of DWI that enables the depiction of vectors corresponding to the strength and direction of water movement in tissues, usually with a second-order tensor[3], in order to capture the anisotropic effects of diffusion in biological tissue, which is not possible with a single diffusion-weighted (DW) image. DTI has the potential to increase our understanding of neural connectivity and how specific brain regions interconnect [13].

The partial volume effect (PVE) arises in volumetric images, such as DTI, when there are crossed nerve fibres, presence of inhibited extracellular water or more than one tissue type occurs in a voxel. In such cases, the voxel intensity depends not only on the imaging sequence and tissue properties, but also on the proportions of each tissue type present in the voxel. As a result, DTI-derived scalar measurements are not always easy to interpret, which can lead to questionable results [14].

To overcome these problems, DTI requires a minimum of six dwMRI images in linearly independent directions to compute six separate diffusion components for each voxel, using a symmetric 3 x 3 tensor to model diffusion. The diffusion tensor can be broken down into several metrics, each of which contains more data about brain microstructure than a single dwMRI image.

1.2.4 The effective Diffusion Tensor and Scalar invariants

After model fitting, meaningful quantities can be estimated, including diffusion tensor and fractional anisotropy in DTI [15].

In the work of Basser et. al [16], the signal decay (S_k) and effective diffusion tensor are defined as follows:

$$S_k = S_0 * exp(-b * q_k^T * D * q) \tag{1.3}$$

where S_0 is the non diffusion-weighted signal, q_k is the normalized direction (column vector) of the k^{th} applied gradient (b = 0 s/mm^2), b is the weighting factor that controls the contrast and D is the effective diffusion tensor:

$$[D = \begin{matrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{matrix}] \tag{1.4}$$

which can be visually represented by an ellipsoid. For isotropic diffusion, the diffusion tensor can be imagined as a sphere; conversely, when the diffusion tensor is strongly anisotropic due to complex tissue confinements, it can be visualized as an ellipsoid with its major axis parallel to the direction of highest diffusivity.

Another measure obtained from DTI data that has been used to investigate differences in brain structural integrity is mean diffusivity (MD), which describes the rotationally-invariant magnitude of water diffusion within brain tissue:

$$MD = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3} \tag{1.5}$$

The eigenvalues of D are represented by $\lambda_1 > \lambda_2 > \lambda_3$. MD is generally similar between different subjects across WM and GM - at values close to $0.8 \times 10^{-3} mm^2 s^{-1}$, and close to D_w ($D_w = 3 \times 10^{-3} mm^2 s^{-1}$) in CSF. An increase in overall diffusion typically consists of increased water content (i.e., edema and inflammation) and thus has relatively less resistance, and therefore, higher diffusion rates.

[3] A second-order tensor describes a mapping that takes one vector as input, and gives one vector as output

The degree of anisotropy of a diffusion process is described by fractional anisotropy (FA), which is a scalar value between zero and one. A value of zero indicates that diffusion is isotropic, which means that it is unrestricted (or equally restricted) in all directions. The commonly used fractional anisotropy index can be regarded as the standard deviation of eigenvalues to the mean diffusivity:

$$FA = \sqrt{\frac{3}{2}} * \sqrt{\frac{(\lambda_1 - MD)^2 + (\lambda_2 - MD)^2 + (\lambda_3 - MD)^2}{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \tag{1.6}$$

1.3 Free Water Elimination DTI (FW-DTI)

The goal of FWE-DTI is to separate the water and tissue components of the signal, which is done by adding a second compartment to equation 1.3, as proposed by Pierpaoli et al., 2004 [17]:

$$S_k = S_0 * [(1 - f_w)\, exp(-b * q_k^T * D_t * q) + f_w * exp(-b * D_w)] \tag{1.7}$$

Where $0 < fw < 1$ quantifies the FW volume fraction and D_t is the corrected diffusion tensor. For single-shell data, fitting equation 1.7 to estimate the seven model parameters (six diffusion tensor elements and the volume fraction) is an ill posed problem because for every fixed value of free water, a viable D_t can be obtained with linear regression. As a result, selecting the optimal pair of estimations (FW, D_t) (global minimum) is not easy, and many believe it is impossible. In his work, Pasternak et al. [18] used local spatial constraints that enforce smooth variation of the diffusion tensor resulting in stable FW-DTI fits in single-shell data. In the case of multi-shell data, however, there is enough information to fit D_t and FW at the same time. While anisotropic tissue effects are more proeminent at high b-values, isotropic FW effects are more evident at low b-values.

1.4 State-of-the-Art

FWE-DTI has attracted a huge interest to characterize diffusion alterations in the context of Parkinson's disease [19], Alzheimer's disease [20], age-related differences in white matter [21], white matter alterations following severe stroke [22], schizophrenia [23], peritumoral tissue [24], among others. These studies suggest that FWE-DTI is fundamental to decouple changes associated with microstructural alterations from increases in partial volume near interfaces between tissue and CSF in particular, partial volume increases associated with gross tissue atrophy and ventricular enlargement. Moreover FWE-DTI has the advantage of providing a FW fraction estimate, a potential surrogate marker for edema.

The enthusiasm in developing FW estimation/elimination algorithms in order to correct the biased diffusion measures obtained with standard DTI is increasing. An easy way to approach it is to modify the standard DTI model and add a second compartment that accounts for the isotropic component - the FW fraction becomes an additional parameter to be estimated. Hoy et al., 2014 [25], performed a modified version of Weighted Linear Least Squares (WLLS) on a multi-shell (at least two non-zero different diffusion weightings) dwMRI acquisition to fit the two compartment model. Convergence was stabilized and diffusion measures corrected for FW were obtained due to the extra information provided by the multi-shell data. An open source implementation of this algorithm was later optimized by Henriques et al., 2017 [26]. However, multi-shell datasets are less common in a clinical environment, due to longer acquisition times (since the protocols are typically optimized for DTI, if the number of directions per shell remains the same, a two shell acquisition will take twice as long). Therefore, single-shell acquisitions

(typically acquired with $b = 1000s/mm^2$) are preferred. Given that, Pasternak et al., 2009 [18] proposed and implemented the single-shell strategies: this algorithm consists of adding spatial constraints to the data and performing a Gradient Descent (GD) method in order to help convergence of the ill-posed problem. In his work, Golub et. al [9] made modifications to each step of the algorithm as an attempt to improve the robustness technique. All developments performed were tested in both synthetic and *in vivo* data and then integrated into the open-source platform Dipy. The performance of the algorithm was also compared to the performance of the open source algorithm for multi-shell data implemented by Henriques et al., 2017 [25]. In [9], the mathematical framework behind the implemented algorithm is presented.

1.5 Single-shell VS Multi-shell

In DWI, signal measurements are a limited number of samples of the diffusion signal in 3D q-space. Reconstruction in dMRI entails fitting the signal samples by using a well-designed model, including DTI [15].

The optimal sampling scheme yields the best reconstruction quality. A single spherical shell is a uniformly sampled spherical manifold with a single diffusion weighting (b-value) where a multi shell includes two or more spherical shells. A typical DTI sampling scheme is a single-shell scheme with about 30 samples and a low b-value ($\leq 1200s/mm^2$), while multi-shell schemes are preferred (or necessarily needed in some cases) with a second shell of b-value ($\leq 2000s/mm^2$).

Because the FWE-DTI fitting is only well suited for multi-shell acquisitions, a regularised gradient descent (RGD) method was previously implemented to allow it to be used to single-shell data. However, the reliability of the RGD method has been poorly assessed. Golub et. al [27] conducted a simulation study that aimed to quantify the specificity of FWE-DTI procedures on single and multi-shell data.

1.6 Regularized Gradient Descent fitting procedure

Previous studies suggested utilising RGD techniques with precise parameter initializations to fit the FWE-DTI model to single-shell data [18][28]. Golub et. al [9] set tissue MD and FA estimates to zero when the refined tissue is below 0.1, since FWE-DTI estimates are not well defined for voxels containing only FW.

Three different strategies were tested to initialize f_w (free water fraction, $(0 < f_w \leq 1)$) in the RGD algorithm: a) initialization based on the T2-weighted images, b) initialization based on a tissue's Mean Diffusivity prior and c) hybrid initialization, which are further explained in [9]. Particularly, a) initialization based on T2-weighted information uses priors on the typical pure FW and tissue signals, b) initialization based on MD assumes a constant prior for MD_t (MD_t is the fixed tissue's MD prior), while c) the hybrid initialization is just a log interpolation between the former techniques. It was shown that the plausibility of FWE-DTI initializations is the main determinant factor for the plausibility of FWE-DTI single-shell estimates. For the case of fixed ground truth $MD_t = 0.6\mu m^2 ms^{-1}$ and varying f_w, the initialization based on MD_t shows the smaltest deviations to the ground truth line, and was therefore used in the present work. The mathematical context of this work and the theory behind how the implemented algorithm works are detailed in [9].

Chapter 2

Methods

2.1 Participants

For the current study, a dataset of multi-shell DWI data of 77 healthy controls was used, previously acquired as part of a longitudinal traumatic brain injury study [8] led by the NTNU (Norwegian University of Science and Technology) in Trondheim. This dataset included 47 females and 30 males, ages ranging from 16 to 60 years old with education status that ranges from below GCSE (*General Certificate of Secondary Education*) to university degrees. Exclusion criteria were: (1) non-fluency in the Norwegian language; (2) having a pre-existing severe neurological, psychiatric, somatic, or substance use determined to be severe enough to interfere with outcome assessment; (3) having a prior history of a mild (e.g. prior traumatic lesions on CT[1]), moderate, or severe TBI; or (4) other major trauma. All participants, or their legally authorized representative, gave consent prior to participation.

2.2 Image Acquisition

MRI was acquired with a 3T Siemens Skyra system (Siemens Healthcare, Erlangen, Germany) with a 32-channel head coil. The image protocol consisted of a series of clinical MRI sequences, including: three dimensional (3D) T1-weighted magnetization-prepared rapid acquisition with gradient echo (MPRAGE) and two dimensional (2D) diffusion weighted imaging (DWI) and others not used in the current project, but formerly specified [29]. The T1w MPRAGE scans were obtained with echo time, TE= 4.21ms, repetition time, TR= 2300ms, time to inversion, TI= 996ms, and a flip angle of 9 degrees. T1w images consisted of 176 slices with isotropic $1mm^3$ voxels covering a FOV (*field of view*) of 256x256. Diffusion weighted scans included four base line volumes (b= $0s/mm^2$) and 60 diffusion sensitised volumes with two different b-values (b= $1000,2000s/mm^2$, same 30 directions each). To acquire DWI images (60 axial slices; FOV =96x96) with isotropic $2.5x2.5x2.5mm^3$ voxels, scanner parameters were set to TR= 8800ms and TE= 95ms. Additional b0 volumes with opposite phase encoding directions were acquired (same parameters).

2.3 Image processing

T1w images were N4 bias corrected and spatially normalised to T1w MNI template via affine and deformable registration, both using ANTs - Advanced Normalization Tools [30] and brain masks were computed, using a brain extraction tool (HD-BET[2]) [31]. Diffusion images were denoised via MPPCA[3] (MRtrix3) [32] and corrected for Gibbs' ringing artefacts (MRtrix[4])[33]. Susceptibility distortions were estimated with topup[5] (from FSL[6]) [34] by using an additional b0 volume with opposite phase encoding direction. Head motion and eddy current artefacts were corrected for via eddy[7] (FSL) [35]. Lastly, B0 magnetic field inhomogeneities[8] were corrected (MRtrix3) [30].

The preprocessed, artefact corrected DWI images were used to compute anisotropic power maps, which have a similar contrast to T1w MPRAGE images. For each subject, anisotropic powermaps were aligned with T1w using the neuroimaging tool ANTs. The found transformation was then used to project FA, MD maps and other DWI derived images to T1w space.

Brain masking was achieved by registering diffusion data to T1w scans of the same scan session (ANTs) [36] and backprojecting of the previously computed brain mask from T1w to diffusion space. Diffusion parameter (FA & MD) and free water (FW) maps were estimated both with single (b=0, 1000s/mm^2 volumes) [37][38] and multi-shell data (Dipy) [39][40].

2.4 Statistical analysis

For both datasets (single and multi) and all parameters we used one fixed pipeline.

2.4.1 Registration

The first step was to do the normalization of the controls DWI images - which are in the native space - into standard stereotactic space (Montreal Neurological Institute, MNI) - in order to allow the performance of statistical tests.

Tract-based spatial statistics (TBSS) [41] was used for registration and alignment but not for statistical analysis. Furthermore, the process of building the common white skeleton characteristics of TBSS was not performed either. TBSS uses a template based on FA maps obtained with the DTI fit (e.g. the FA images obtained using the function dtifit in FSL, which are not corrected for free-water), so in order to get the same contrast, the first input to TBSS registration steps were the DTI fit FA maps.

First, the images were slightly eroded and the end slices were set to zero (to remove likely outliers from the diffusion tensor fitting). Next, with nonlinear registration (FNIRT[9]), all FA images were aligned to a 1mm^3 standard space: the FMRIB58 FA[10]image. Each subject's FA image had the nonlinear transform to the target and the affine transform to MNI152 space applied, resulting in a transformation of the

[2]Automated brain extraction of multisequence MRI using artificial neural networks

[3]Marcenko-Pastur PCA algorithm

[4]Advanced tools for the analysis of diffusion MRI data

[5]A tool for estimating and correcting susceptibility induced distortions

[6]Functional Magnetic Resonance Imaging of the Brain Software Library is a software library containing image analysis and statistical tools for functional, structural and diffusion MRI brain imaging data

[7]A tool for correcting eddy currents and movements in diffusion data

[8]Normal magnets contain various magnetic field inhomogeneities, which is one of the sources effecting T2*. This results in protons precessing at different Larmor frequencies and subsequently dephasing.

[9]Non linear registration intensity based for MRI brain exams.

[10]A high-resolution average of 58 well-aligned good quality FA images from healthy male and female subjects aged between 20-50

original FA image into MNI152 space - which will be refered to as template space. Next, these were all merged into a single 4D image file (all_{FA}) and the mean of all FA images was created, called $mean_{FA}$. Here, the traditional pipeline was interrupted and a spatial filter was applied to these images in order to exclude from the mask voxels that have a value of zero for at least one of the subjects - null voxels are excluded because they do not represent the correct value in that part of the brain, and by including them in the statistical analysis, they will artificially lower the mean for the most affected group.

It is straightforward to apply TBSS to data diffusion-derived other than FA images - there was an interest in how MD or FW vary between the two different algorithms in this study. At this point, the images for all parameters (FA, MD, FW) and for both algorithms (single and multi-shell) were aligned to MNI152 space, the original nonlinear registration (using the original DTI-fit FA data) was applied to the new data and a 4D file statistics resulting of the merging of all subjects' warped data was created (for each parameter and algorithm). Finally, all data were prepared for voxel-wise statistical analysis over the original (DTI-fit) FA space.

2.4.2 Voxelwise statistics

Voxel-wise analysis across the single and multi-shell images was performed using permutation-based, voxel-wise nonparametric testing [42] implemented as the "*randomise*" function in the FSL package, in order to investigate which voxels are significantly different between the two groups. The number of permutations in all tests was equal to 2,000. Statistical threshold $p < 0.05$, after correction for multiple comparisons with TFCE(Threshold-Free Cluster Enhancement, identical to cluster-based thresholding, but more robust in general and the initial random cluster-forming threshold is unnecessary) was used during the analysis [43].

To run *randomise*, a design matrix file (design.mat) and contrasts file (design.con) are needed. In this case, each subject has exactly 2 measures (from the two different algorithms, for each parameter) and the average difference between these two measures across subjects is of interest. The Glm GUI[11][44], which allows the specification of designs, was used to generate a single-group paired difference (Paired T-Test) matrix for the 77 controls.

A similar voxelwise analysis procedure was applied to the MD and FW images.

2.4.3 Analysis of regional values

In order to understand how the algorithms behave at a regional level, white matter, cerebral cortex and subcortical masks were obtained by mapping the Harvard-Oxford cortical and subcortical structural atlases [45] 21 regions (1 mm resolution) to the p-value maps produced by *randomise* in MNI1521 space. The ventricular mask was excluded.

[11] General linear model graphical user interface

Table 2.1: 21 regions of the Harvard-Oxford cortical and subcortical structural atlas and the corresponding masks created: white matter (regions 1 + 12); cerebral cortex (regions 2 + 13) and subcortical (regions 4 to 11 + 15 to 21)

Left hemisphere		Right hemisphere		Mask
1	Cerebral white matter	12	Cerebral white matter	White matter
2	Cerebral cortex	13	Cerebral cortex	Cerebral cortex
3	Lateral ventricle	14	Lateral ventricle	Ventricular
4	Thalamus	15	Thalamus	Subcortical
5	Caudate	16	Caudate	
6	Putamen	17	Putamen	
7	Pallidum	18	Pallidum	
8	Brain-stem	19	Hippocampus	
9	Hippocampus	20	Amygdala	
10	Amygdala	21	Accumbens	
11	Accumbens			

Chapter 3

Results and discussion

3.1 Visual comparison of data run with the multi-shell algorithm and the single-shell algorithm

In order to better understand the difference between images after the implementation of multi-shell algorithm and the single-shell algorithm, Figures 3.1-3.3 are a side-by-side comparison of the two algorihms applied to the same randomnly chosen subject, for all parameters studied (FA, MD, FW), respectively.

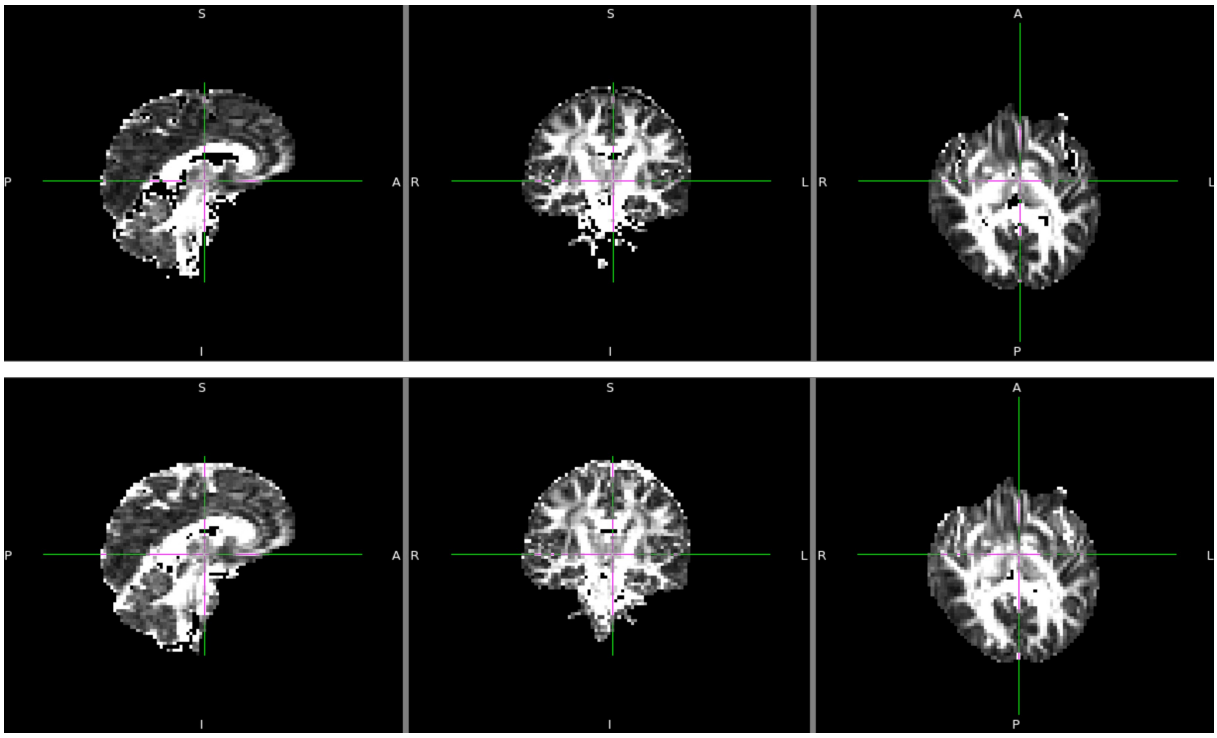

Figure 3.1: Comparison of DTI images for FA values of a random subject, after implementation of the multi-shell (top section of the figure) and single-shell algorithm (bottom of the figure). For this specific case, FA values range between 0-0.637 when processed with the multi-shell algorithm and between 0-0.613 when processed with the single-shell algorithm.

3. RESULTS AND DISCUSSION

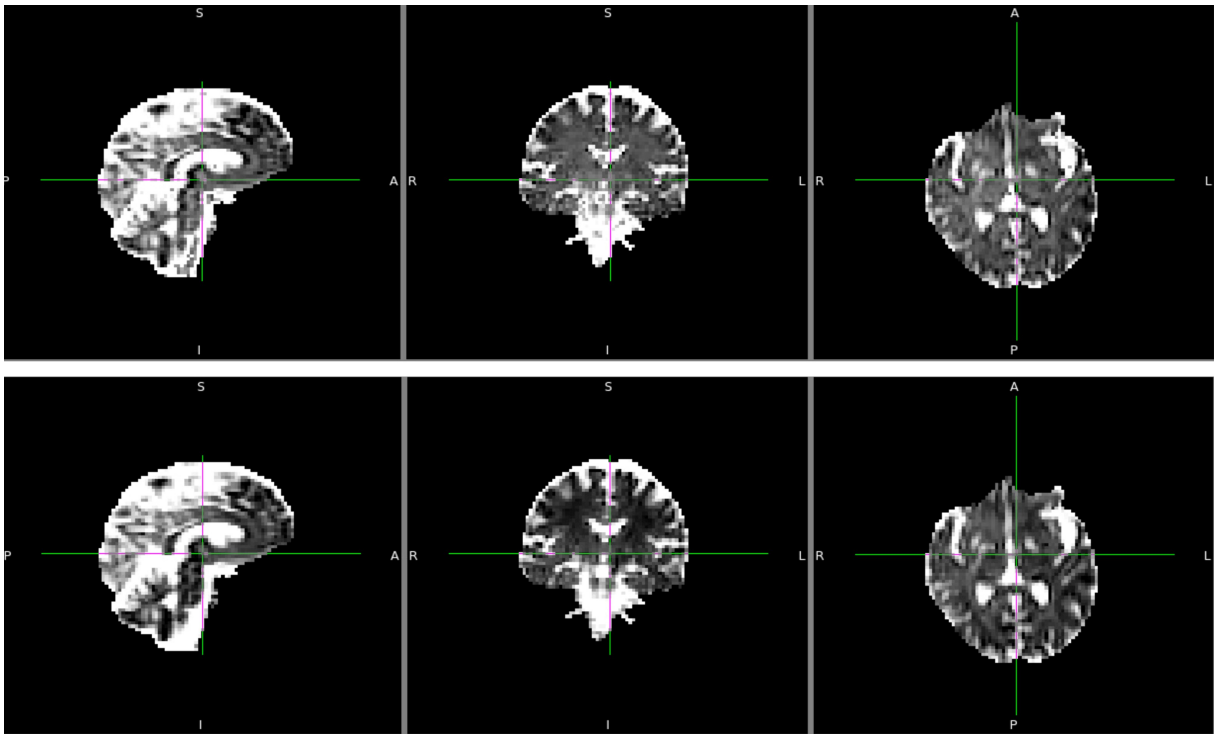

Figure 3.2: Comparison of DTI images for FW values of a random subject, after implementation of the multi-shell (top section of the figure) and single-shell algorithm (bottom of the figure). For this specific case, FW values range between 0-0.808 when processed with the multi-shell algorithm and between 0-0.817 when processed with the single-shell algorithm.

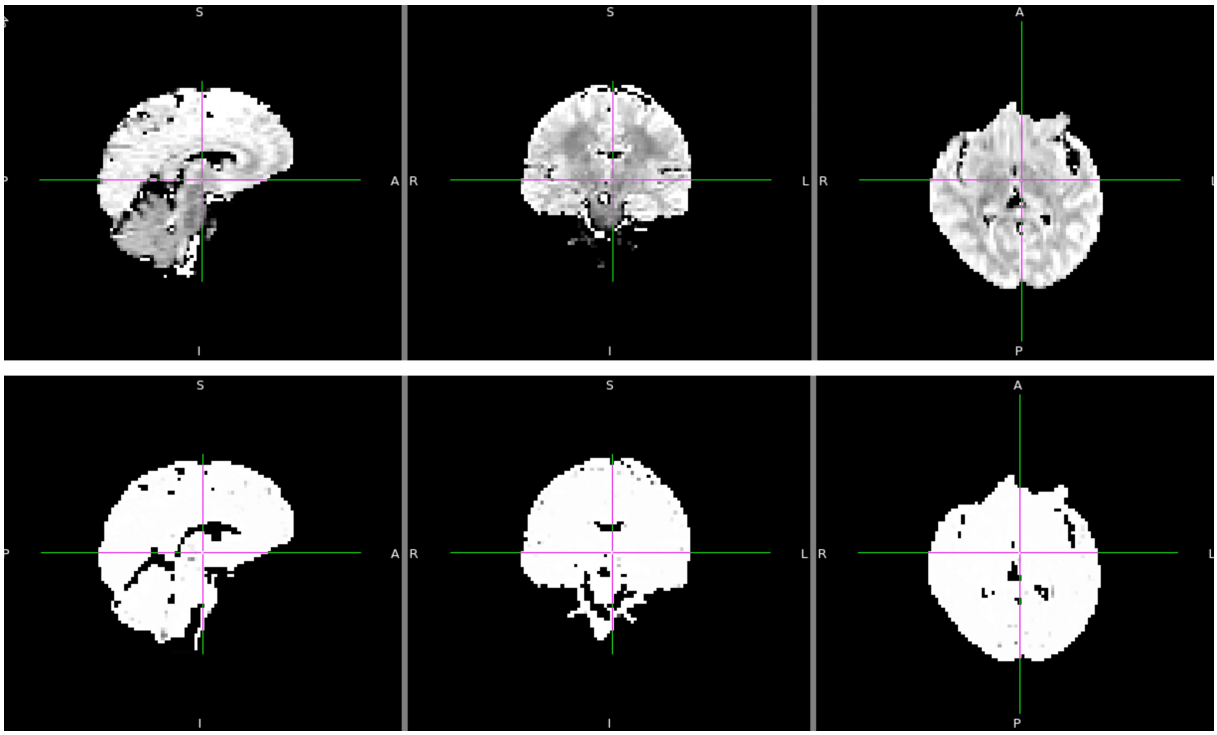

Figure 3.3: Comparison of DTI images for MD values of a random subject, after implementation of the multi-shell (top section of the figure) and single-shell algorithm (bottom of the figure). For this specific case, MD values range between $0 - 0.660\mu m^2 ms^{-1}$ when processed with the multi-shell algorithm and between $0 - 0.605\mu m^2 ms^{-1}$ when processed with the single-shell algorithm.

The ortho view (Figures 3.1-3.3) comprises three canvases, which display the overlays along three

orthogonal planes. In the case of a NIFTI image which is oriented acording to the MNI152 template, those canvases correspond to the sagittal, coronal, and axial planes. Figures 3.1-3.3 are a random sample that allows a visualization of the original data used, and should not be generalized or used to draw conclusions. However, it is possible to see that there are obvious differences in the images run with the two algorithms - the maximum intensity value of the voxels, the location of the most and least intense zones, how the voxels are perceived in the CSF zones (ventricles and perimeter of brain parenchyma). Regarding figure 3.3, we can say that using initialization with $MD_t = 0.6\mu m^2 ms^{-1}$ conditioned the MD images run with the single-shell algorithm and limited the range of values for this parameter to values very close to this prior.

Upon obtaining the results in which the ventricles appear significantly different between both algorithms (very clear in Figure 3.1), an anomaly was detected: if they are assigned the value 1 in the FW parameter for both algorithms, one can assume that there are no significant differences in this region. However, the differences in the ventricles may be due to the effect of normalization. The reason the FA maps obtained with the DTI fit were used to do the normalization was because the corrected FA maps (with the FWE) had zeros in certain areas (as explained in 1.6, MD and FA were considered zero for values below 0.1), which does not happen with the DTI fit FA maps.

When comparing the FA maps with both algorithms (multi and single-shell) for the same subject, the parenchyma presents many differences - which can be seen through statistics. The intensity of the FA values in the ventricle (scale from 0 to 1) was supposed to be close to zero and that is not the case - some voxels have unexpected high values. This can be explained as a PVE, in this case due to ”leakage” from the subcortical region to the ventricles.

To make sure that the voxels set to zero from both algorithms are comparable (and that the differences are not a result of corregistration), subjects maps both before and after registration were analysed. The difference of the number of null voxels between both algorithms (for the same subject) is $< 1\%$ for FA and MD, in both cases: images before corregistration (native space) and after corregistration (template space). As to FW, comparison of the number of voxels with values > 0.99 for these parameters led to the same conclusions: the difference between algorithms is $< 1\%$, and was considered negligible. As expected, the difference between algorithms is smaller after registration.

3.2 Statistical results

The results of the voxelwise analysis procedure ”*randomise*” are shown in the form of a corrected p-value image for each parameter:

3. RESULTS AND DISCUSSION

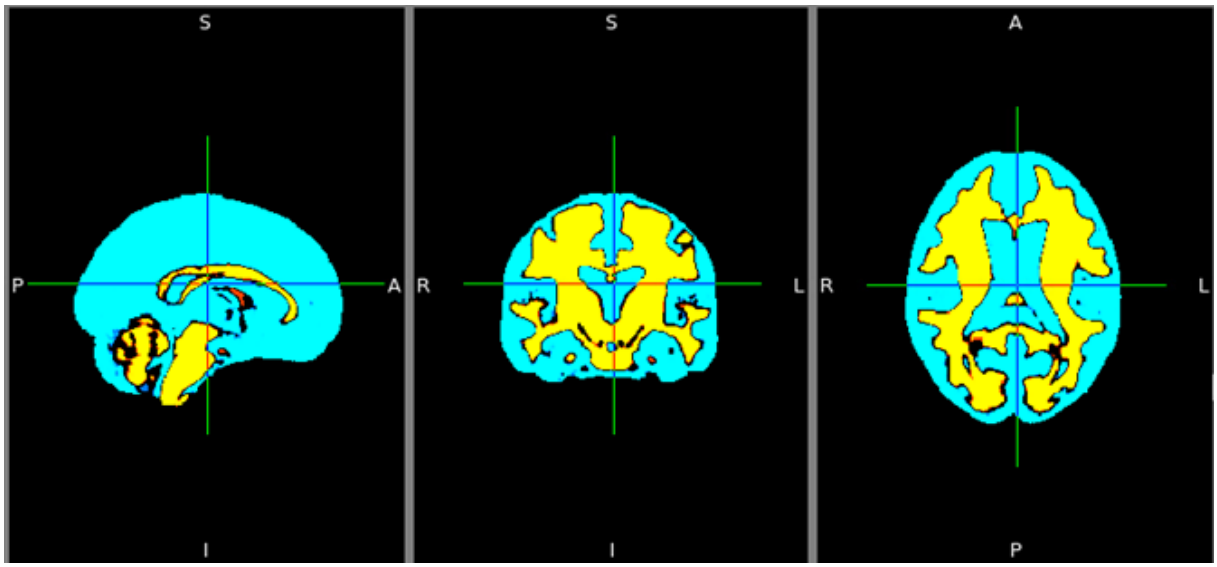

Figure 3.4: Corrected p-value image for FA values: red-yellow corresponds to the multi>single test and blue-light blue gives the multi<single test, both thresholded at .95 which shows significant clusters

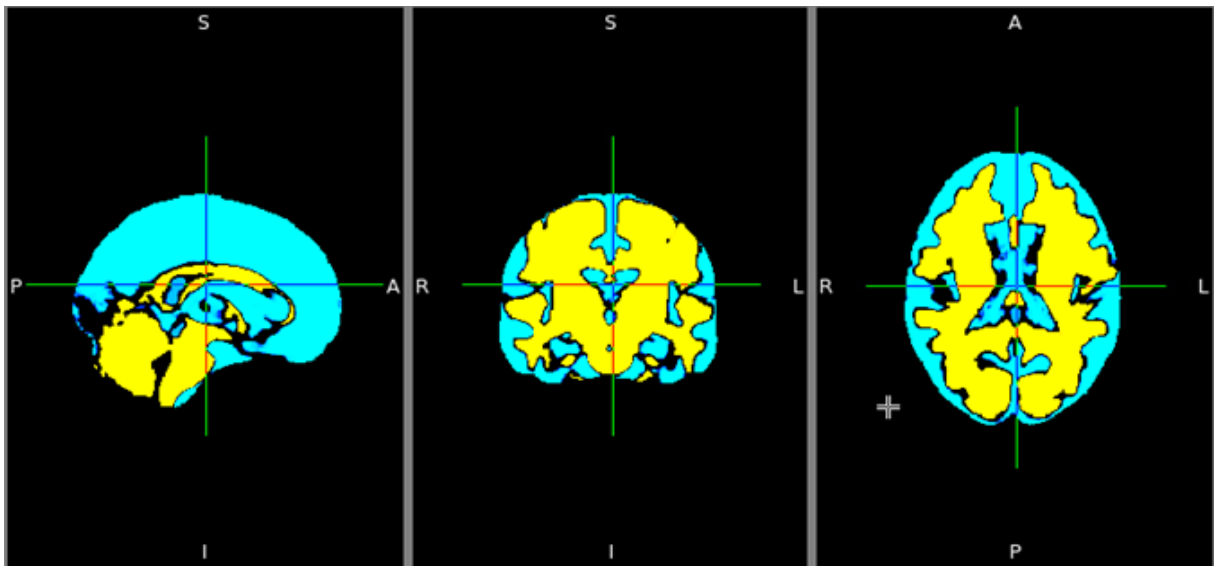

Figure 3.5: Corrected p-value image for FW values: red-yellow corresponds to the multi>single test and blue-light blue gives the multi<single test, both thresholded at .95 which shows significant clusters

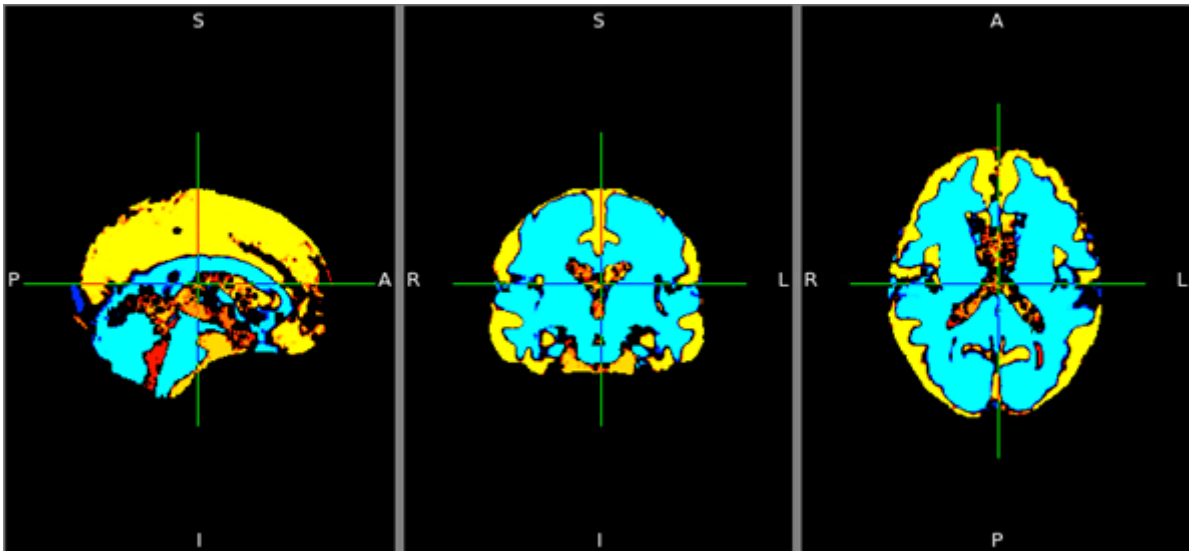

Figure 3.6: Corrected p-value image for MD values: red-yellow corresponds to the multi>single test and blue-light blue gives the multi<single test, both thresholded at .95 which shows significant c lusters. The red parts represented here correspond to the voxels where the p-values are lower and closer to the 0.95 threshold for the case of the multi>single test.

The difference is widespread throughout the brain and the only systematic difference visually is gray matter vs white matter.

Considering the multishell algorithm as gold standard, for both FA and FW the singleshell algorithm seems to underestimate the white matter values and overestimate the gray matter values. For MD, it is the opposite: the singleshell algorithm seems to overestimate the white matter values and underestimate the gray matter values. Given that we are using a prior based on MD, this can explain the overestimation of the singleshell algorithm for this parameter (further analysed ahead).

3.3 Further analysis

Each subject images for all parameters (FA, MD, FW) and for both algorithms (single and multi-shell) was used to get an average value of the voxels (excluding null values from this average), in order to understand how these values differ according to the algorithm used. The following boxplots display the distribution of data based on a five number summary (minimum, first quartile (Q1), median, third quartile (Q3), and maximum). In each box, the central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually using dots. A second graphic was created depicting lines that match values of the same subject for the two algorithms, making it possible to see if there is a common tendency between algorithms.

3. RESULTS AND DISCUSSION

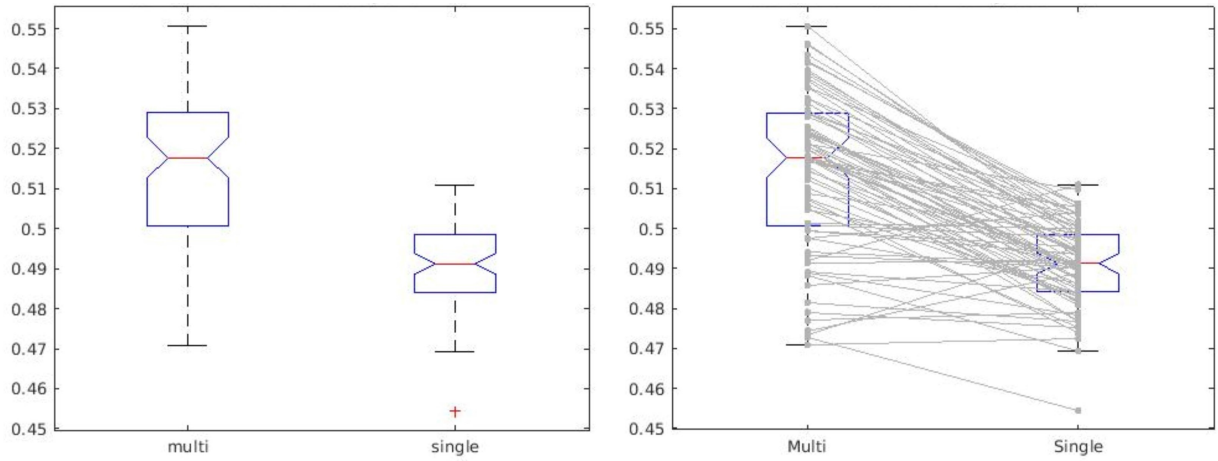

Figure 3.7: Distribuition of the average FA values of each subject for both algorithms in white matter. The second graphic depicts lines that match FA values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

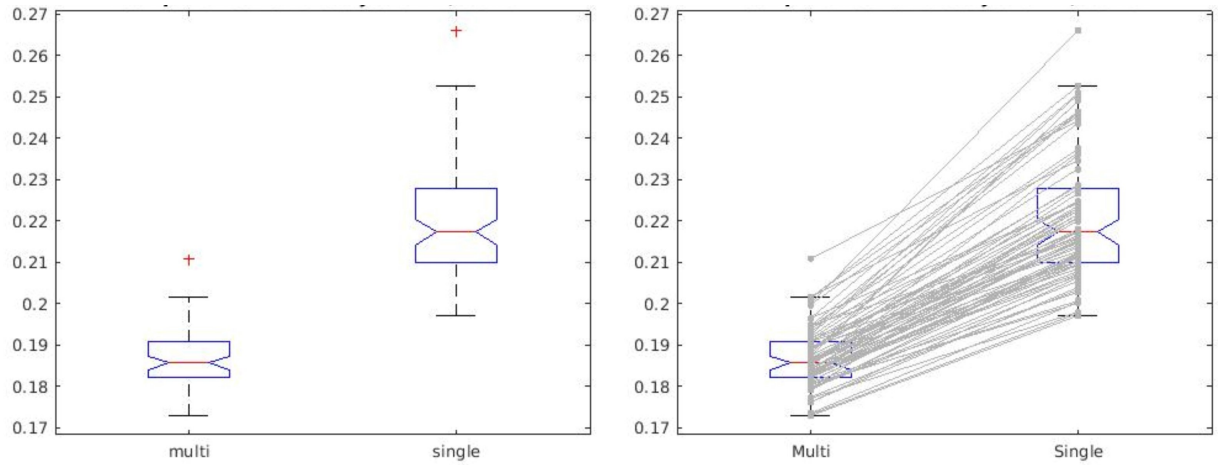

Figure 3.8: Distribuition of the average FA values of each subject for both algorithms in the cerebral cortex. The second graphic depicts lines that match FA values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

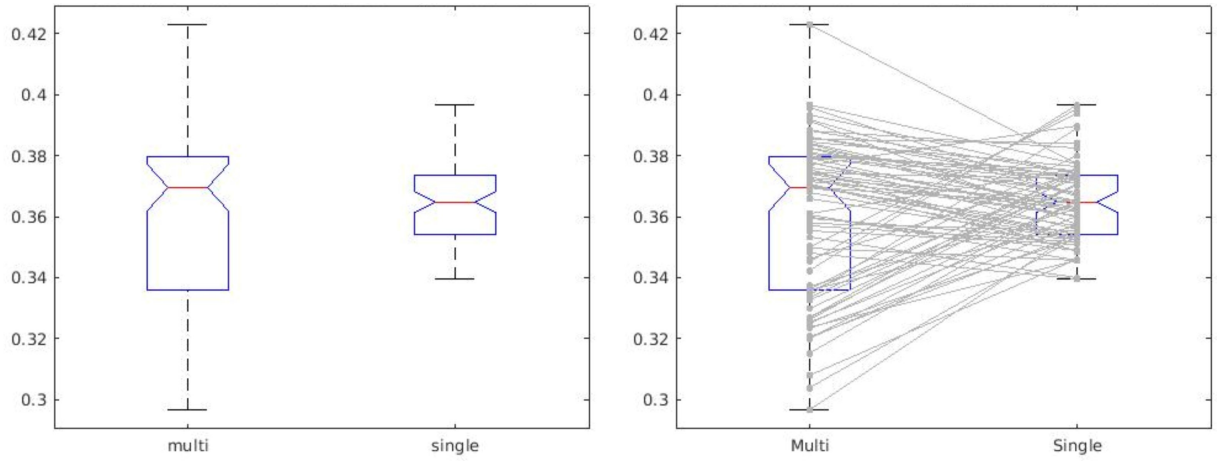

Figure 3.9: Distribuition of the average FA values of each subject for both algorithms in the subcortical region. The second graphic depicts lines that match FA values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

Regarding the distribution of the average FA values of each subject, Figures 3.7 and 3.8 corroborate the results of the voxelwise analysis as they show that the single-shell algorithm is underestimating the white matter values, and overestimating the grey matter (cerebral cortex) values. For the subcortical region, the median is similiar for both algorithms but the range of the distribution is much smaller in the single-shell algorithm. Regarding the second graphic, it is possible to see a tendency of subjects with higher averages in multi-shell algorithm corresponding to subjects with higher values in the single-shell algorithm for WM and for the cerebral cortex. However, for subcortical region there seems to be a number of subjects that present lower averages on the multi-shell algorithm and correspond to higher averages on the single-shell algorithm (which can be depicted by ”crossing” lines in the second graphic of Figure 3.9). One of the things that is noticeable is that for both WM and the cerebral cortex, the median of the single-shell boxplot is outside the range of values for the 25th-75th percentile of the multi-shell boxplot. This can indicate that they are signicantly different, which meets the results of the statistical analysis.

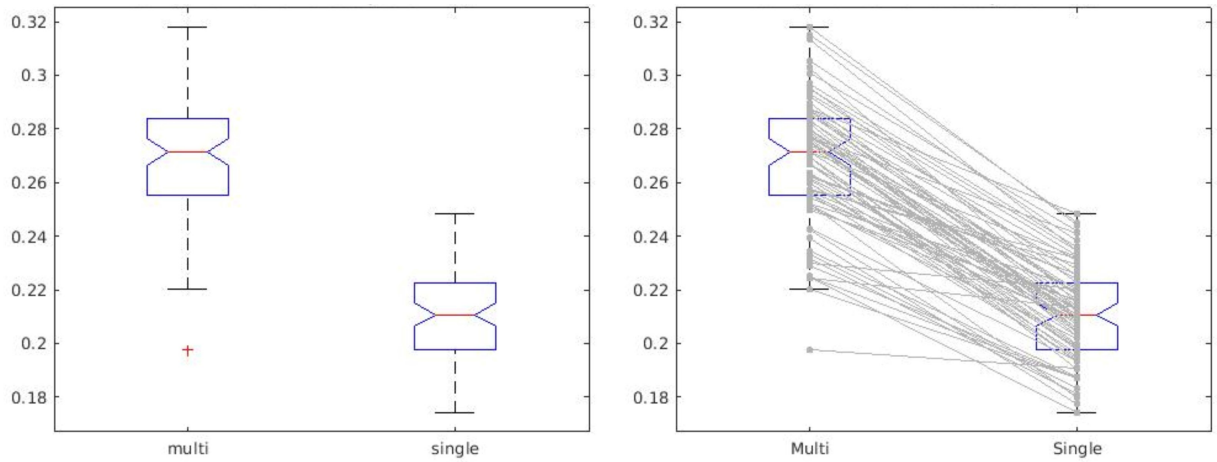

Figure 3.10: Distribuition of the average FW values of each subject for both algorithms in white matter. The second graphic depicts lines that match FW values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

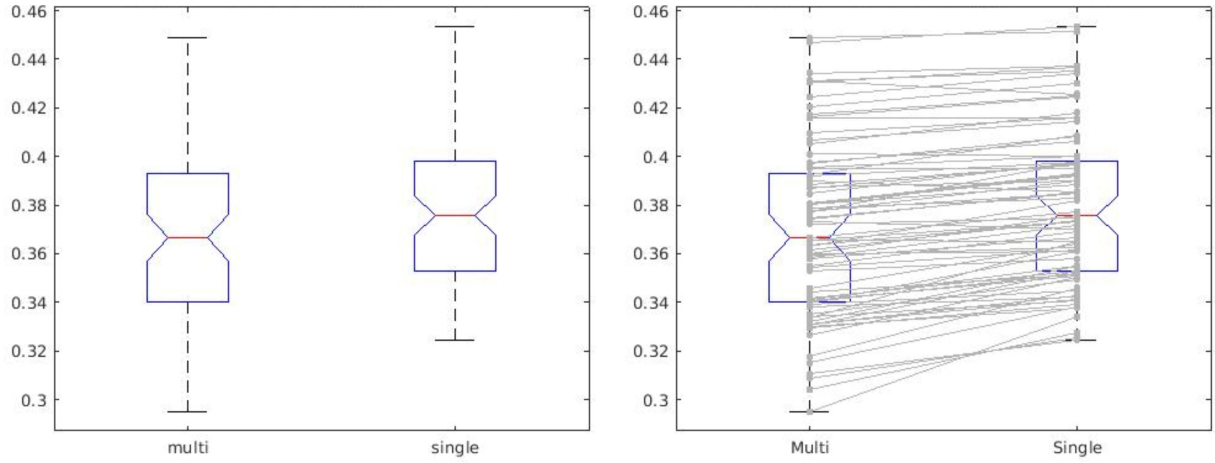

Figure 3.11: Distribuition of the average FW values of each subject for both algorithms in the cerebral cortex. The second graphic depicts lines that match FW values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

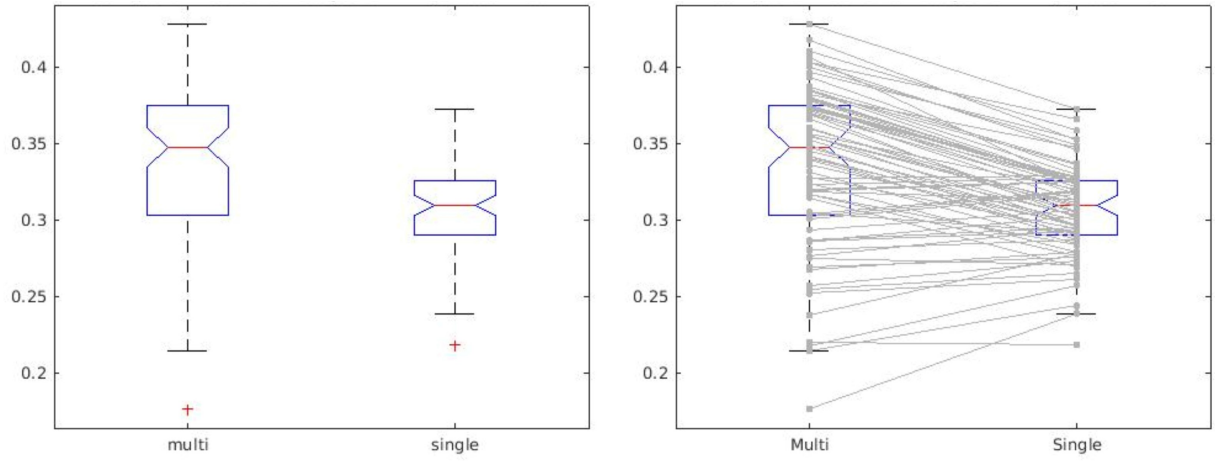

Figure 3.12: Distribuition of the average FW values of each subject for both algorithms in the subcortical region. The second graphic depicts lines that match FW values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

When analysing the distribution of the average FW values of each subject, Figures 3.10 is similar to Figure 3.7 and shows that the single-shell algorithm is underestimating the white matter values, which can be observed in the statistical results. For the cerebral cortex, there is a slight overestimation of the grey matter FW values by the single-shell algorithm, but the boxplots for both algorithms look very similar. When analysing the second graphic of Figure 3.11 the lines that match the average of the same subject for both algorithms tend to be horizontal and without ”crossing lines”, meaning that the average value for the same subject is quite similar. For the subcortical region, the range of the distribution is much smaller in the single-shell algorithm and in the second graphic, this region seems to be the one that presents more ”crossing lines”. Again, for WM the median of the single-shell boxplot is outside of the range of values of the 25th-75th percentile of the multi-shell boxplot, which can indicate that they are signicantly different.

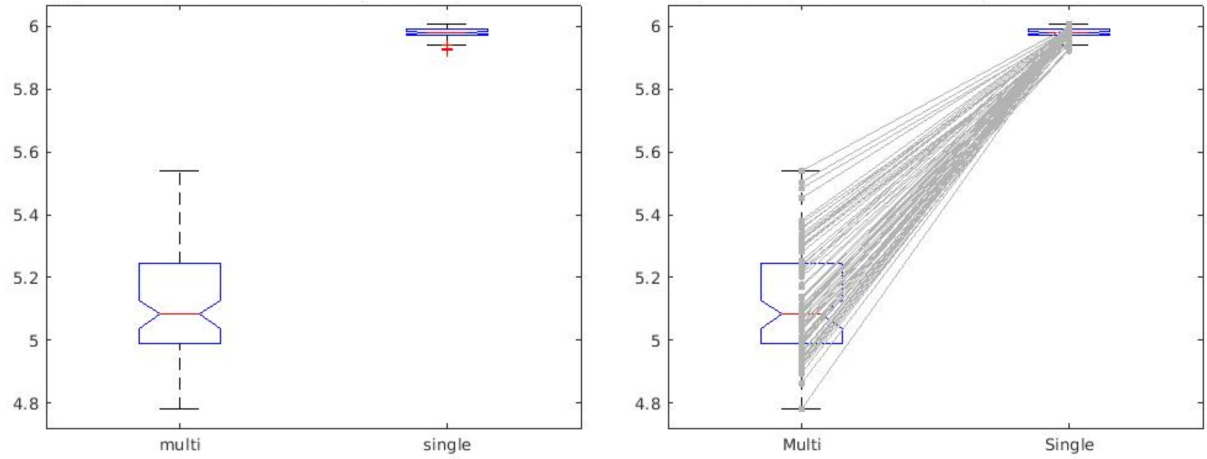

Figure 3.13: Distribuition of the average MD values of each subject for both algorithms in white matter. The second graphic depicts lines that match MD values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

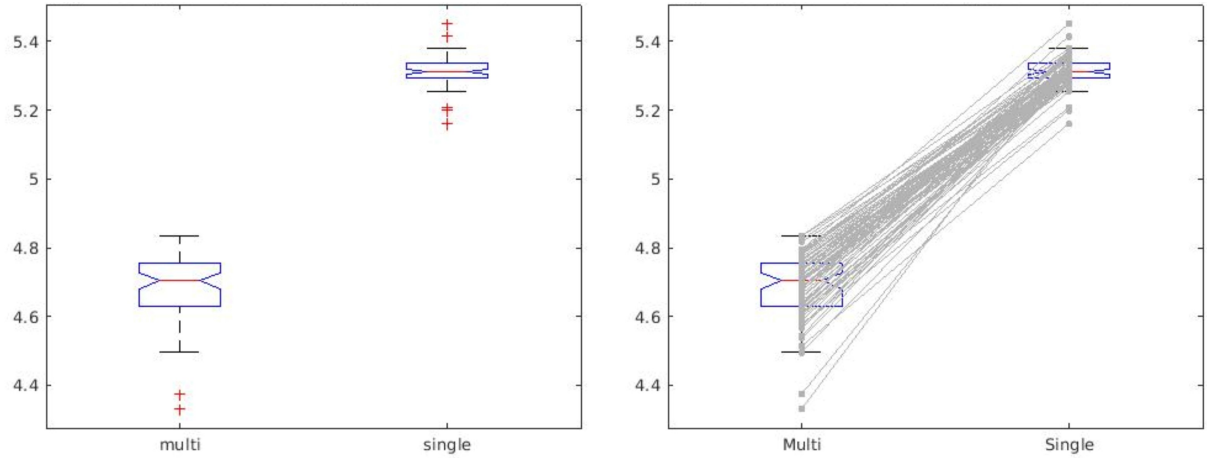

Figure 3.14: Distribuition of the average MD values of each subject for both algorithms in the cerebral cortex. The second graphic depicts lines that match MD values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

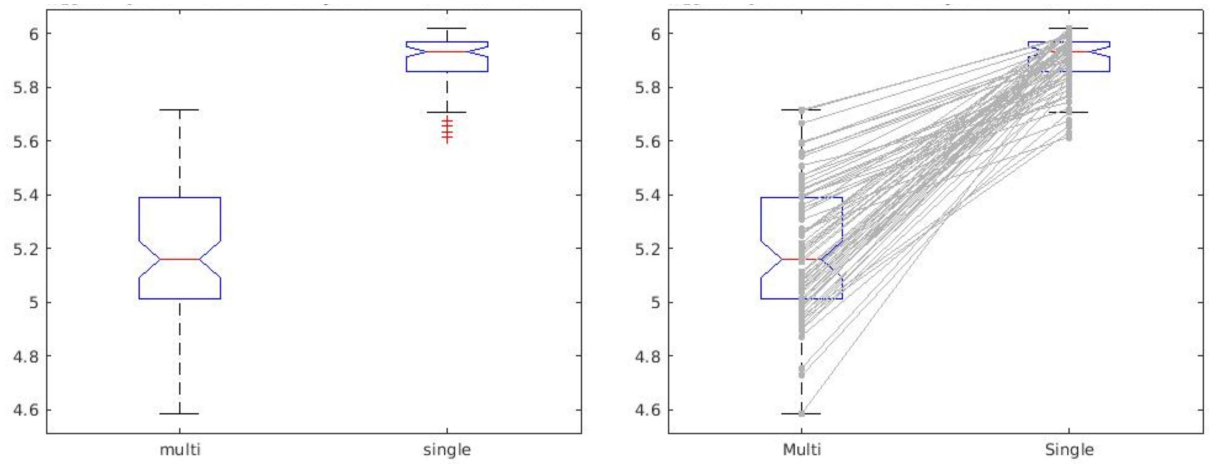

Figure 3.15: Distribuition of the average MD values of each subject for both algorithms in the subcortical region. The second graphic depicts lines that match MD values of the same subject for the two algorithms in order to evaluate tendencies between algorithms.

The comparison of WM values for both algorithms was the most difficult from the beginning. As we can observe in Figures 3.13-3.15, for the regions under study (WM, cerebral cortex and subcortical region), the single-shell boxplot is outside the range of values of the 25th-75th percentile of the multi-shell boxplot, which can indicate that they are signicantly different, meeting the statistical analysis results. The single-shell boxplot seems to have a very thin distribution of values around the $MD_t = 0.6\mu m^2 ms^{-1}$ value used as a fixed ground truth for the initialization of the RGD for WM and the subcortical region, and between $0.52 - 0.54\mu m^2 ms^{-1}$ for the cerebral cortex. According to this, the single-shell algorithm is overestimating the WM, cerebral cortex and subcortical region MD values. When comparing to the statistical results (Figure 3.6), the coronal and axial planes are mainly blue (multi<single test), which agrees with Figures 3.13-3.15 but the sagital view depicts a lot of yellow (multi>single test). This view shows the outtermost layer of the brain, which is mainly grey matter. However, in the coronal plane of Figure 3.6, the yellow (multi>single) layer corresponds to a thin section delimitating the brain. Given that the mask used for the cerebral cortex region will include much more voxels in blue than in yellow, both these images lead to the same conclusion. Besides, when studying diffusion in gray matter, results must be critically analysed: in gray matter there is not a defined ”fiber” so tensor assumptions do not exist: the tensor can find something more isotropic, an average is calculated and the closest approximation is given as result.

Chapter 4

Conclusion

4.1 Conclusion

In 2021, Golub, et al, [38] stressed that results from single-shell FWE-DTI in previous and future studies should be interpreted with care, as based solely on priors, FWE-DTI estimates are not able to distinguish changes in FW content from changes in the tissue's apparent diffusion tensor for single-shell data acquisitions. The results of this study are in line with that, as there was an obvious effect of using tissue MD prior on the parameters values.

When studying MD, the impact of the MD prior initialization method is clear: the range of values of MD is limited to a small scope. And even though the effects are not as straight-forward on FA and FW, they were affected. Plus, when compared to the gold-standard for FWE-DTI - the multi-shell algorithm - this study found that the values obtained with the single-shell algorithm are considered significantly different for both WM and GM, for all parameters except the FA values of the subcortical region and the FW values for the cerebral cortex. For FA and FW, the single-shell algorithm underestimates WM values and overestimates GM values. For MD, these values are highly conditioned by the prior which result in an overstimation regardless of the tissue type and region in study.

4.2 Limitations of the current study

One limitation of the current study was the use of TBSS for the registration of the subjects' images, as studies have showed that methods such as groupwise registration based on ANTS perform better than TBSS[46].

4.3 Future work

In the future, it may be interesting to replicate the results on dwMRI data from real lesions or even on physical phantoms as proposed by Farrher and colleagues [47].

Bibliography

[1] MC Papadopoulos et al. “Molecular mechanisms of brain tumor edema”. In: *Neuroscience* 129.4 (2004), pp. 1009–1018.

[2] AW Unterberg et al. “Edema and brain trauma”. In: *Neuroscience* 129.4 (2004), pp. 1019–1027.

[3] Nikos G Papadakis et al. “Study of the effect of CSF suppression on white matter diffusion anisotropy mapping of healthy human brain”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 48.2 (2002), pp. 394–398.

[4] Andrew L Alexander et al. “Analysis of partial volume effects in diffusion-tensor MRI”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 45.5 (2001), pp. 770–780.

[5] Luis Concha, Donald W Gross, and Christian Beaulieu. “Diffusion tensor tractography of the limbic system”. In: *American Journal of Neuroradiology* 26.9 (2005), pp. 2267–2274.

[6] Ming-Chung Chou et al. “FLAIR diffusion-tensor MR tractography: comparison of fiber tracking with conventional imaging”. In: *American journal of neuroradiology* 26.3 (2005), pp. 591–597.

[7] Mara Cercignani et al. “Mean diffusivity and fractional anisotropy histograms of patients with multiple sclerosis”. In: *American Journal of Neuroradiology* 22.5 (2001), pp. 952–958.

[8] Toril Skandsen et al. “The epidemiology of mild traumatic brain injury: the Trondheim MTBI follow-up study”. In: *Scandinavian journal of trauma, resuscitation and emergency medicine* 26.1 (2018), pp. 1–9.

[9] Marc Golub. “Implementation of an algorithm for estimating free-water fraction in Diffusion-Weighted Magnetic Resonance Imaging”. Master's thesis (2018).

[10] Anthea Maton et al. "Cells: Building blocks of life". In: *Pearson Prentice Hall* (1997).

[11] Ahmet Onur Keskin et al. “Diffusion tensor imaging of visual pathways in patients with acute and chronic optic neuritis”. In: *International Eye Science* 18.9 (2018), pp. 1559–1566.

[12] Geetha Soujanya Chilla et al. “Diffusion weighted magnetic resonance imaging and its recent trend—a survey”. In: *Quantitative imaging in medicine and surgery* 5.3 (2015), p. 407.

[13] José M. Soares et al. “A hitchhiker’s guide to diffusion tensor imaging”. In: *Frontiers in Neuroscience* 7 (2013).

[14] Silvia De Santis et al. “Why diffusion tensor MRI does well only some of the time: Variance and covariance of white matter tissue microstructure attributes in the living human brain”. In: *NeuroImage* 89 (Apr. 2014), pp. 35–44.

[15] Peter J Basser, James Mattiello, and Denis LeBihan. “MR diffusion tensor spectroscopy and imaging”. In: *Biophysical journal* 66.1 (1994), pp. 259–267.

[16] P.J. Basser, J. Mattiello, and D. LeBihan. “MR diffusion tensor spectroscopy and imaging”. In: *Biophysical Journal* 66.1 (Jan. 1994), pp. 259–267.

[17] C Pierpaoli and DK Jones. “Removing CSF contamination in brain DT-MRIs by using a two-compartment tensor model”. In: *International Society for Magnetic Resonance in Medicine Meeting*. 2004, p. 1215.

[18] Ofer Pasternak et al. “Free water elimination and mapping from diffusion MRI”. In: *Magnetic Resonance in Medicine* 62.3 (Sept. 2009), pp. 717–730.

[19] Takashi Ogawa et al. “White matter and nigral alterations in multiple system atrophy-parkinsonian type”. In: *NPJ Parkinson’s disease* 7.1 (2021), pp. 1–12.

[20] Andrew R Hoy et al. “Microstructural white matter alterations in preclinical Alzheimer’s disease detected using free water elimination diffusion tensor imaging”. In: *PloS one* 12.3 (2017), e0173982.

[21] Jordan A Chad et al. “Re-examining age-related differences in white matter microstructure with free-water corrected diffusion tensor imaging”. In: *Neurobiology of aging* 71 (2018), pp. 161–170.

[22] Roza M Umarova et al. “Distinct white matter alterations following severe stroke: longitudinal DTI study in neglect”. In: *Neurology* 88.16 (2017), pp. 1546–1555.

[23] Konasale M Prasad et al. “Differential susceptibility of white matter tracts to inflammatory mediators in schizophrenia: an integrated DTI study”. In: *Schizophrenia research* 161.1 (2015), pp. 119–125.

[24] Abdol Aziz Ould Ismail et al. “Characterizing Peritumoral Tissue Using Free Water Elimination in Clinical DTI”. In: *MICCAI 2018-21st International Conference on Medical Image Computing and Computer Assisted Intervention; Workshop: Brain Lesion*. 2018, pp. 1–9.

[25] Andrew R. Hoy et al. “Optimization of a free water elimination two-compartment model for diffusion tensor imaging”. In: *NeuroImage* 103 (Dec. 2014), pp. 323–333.

[26] Rafael Neto Henriques et al. “[Re] Optimization of a free water elimination two-compartment model for diffusion tensor imaging”. In: (Feb. 2017).

[27] Marc Golub, Rafael Neto Henriques, and Rita Gouveia Nunes. “Free-water DTI estimates from single b-value data might seem plausible but must be interpreted with care”. In: *Magnetic Resonance in Medicine* (Dec. 2020).

[28] Ofer Pasternak et al. “Hockey Concussion Education Project, Part 2. Microstructural white matter alterations in acutely concussed ice hockey players: a longitudinal free-water MRI study”. In: *Journal of Neurosurgery* 120.4 (Apr. 2014), pp. 873–881.

[29] Cathrine Elisabeth Einarsen et al. “Patients with mild traumatic brain injury recruited from both hospital and primary care settings: a controlled longitudinal magnetic resonance imaging study”. In: *Journal of neurotrauma* 36.22 (2019), pp. 3172–3182.

[30] Nicholas J Tustison et al. “N4ITK: improved N3 bias correction”. In: *IEEE transactions on medical imaging* 29.6 (2010), pp. 1310–1320.

[31] Fabian Isensee et al. “Automated brain extraction of multisequence MRI using artificial neural networks”. In: *Human brain mapping* 40.17 (2019), pp. 4952–4964.

[32] Jelle Veraart et al. “Denoising of diffusion MRI using random matrix theory”. In: *Neuroimage* 142 (2016), pp. 394–406.

[33] Elias Kellner et al. “Gibbs-ringing artifact removal based on local subvoxel-shifts”. In: *Magnetic resonance in medicine* 76.5 (2016), pp. 1574–1581.

[34] Jesper LR Andersson, Stefan Skare, and John Ashburner. “How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging”. In: *Neuroimage* 20.2 (2003), pp. 870–888.

[35] Jesper LR Andersson and Stamatios N Sotiropoulos. “An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging”. In: *Neuroimage* 125 (2016), pp. 1063–1078.

[36] Brian B Avants, Nick Tustison, Gang Song, et al. “Advanced normalization tools (ANTS)”. In: *Insight j* 2.365 (2009), pp. 1–35.

[37] Ofer Pasternak et al. “The estimation of free-water corrected diffusion tensors”. In: *Visualization and processing of tensors and higher order descriptors for multi-valued data*. Springer, 2014, pp. 249–270.

[38] Marc Golub, Rafael Neto Henriques, and Rita Gouveia Nunes. “Free-water DTI estimates from single b-value data might seem plausible but must be interpreted with care”. In: *Magnetic Resonance in Medicine* 85.5 (2021), pp. 2537–2551.

[39] Andrew R Hoy et al. “Optimization of a free water elimination two-compartment model for diffusion tensor imaging”. In: *Neuroimage* 103 (2014), pp. 323–333.

[40] Rafael Neto Henriques et al. “[Re] Optimization of a free water elimination two-compartment model for diffusion tensor imaging”. In: *bioRxiv* (2017), p. 108795.

[41] Stephen M Smith et al. “Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data”. In: *Neuroimage* 31.4 (2006), pp. 1487–1505.

[42] Thomas E Nichols and Andrew P Holmes. “Nonparametric permutation tests for functional neuroimaging: a primer with examples”. In: *Human brain mapping* 15.1 (2002), pp. 1–25.

[43] Anderson M Winkler et al. “Permutation inference for the general linear model”. In: *Neuroimage* 92 (2014), pp. 381–397.

[44] Russell A Poldrack, Jeanette A Mumford, and Thomas E Nichols. *Handbook of functional MRI data analysis*. Cambridge University Press, 2011.

[45] Rahul S Desikan et al. “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest”. In: *Neuroimage* 31.3 (2006), pp. 968–980.

[46] Christopher G Schwarz et al. “Improved DTI registration allows voxel-based analysis that outperforms tract-based spatial statistics”. In: *Neuroimage* 94 (2014), pp. 65–78.

[47] Ezequiel Farrher et al. “Dedicated diffusion phantoms for the investigation of free water elimination and mapping: Insights into the influence of T2 relaxation properties”. In: *NMR in Biomedicine* 33.4 (2020), e4210.