

Rochester Institute of Technology

RIT Digital Institutional Repository

Theses

7-31-2019

Registration of Diffusion Tensor Images in Log-Euclidean and Euclidean Space

Kevin Tuttle

Follow this and additional works at: <https://repository.rit.edu/theses>

Recommended Citation

Tuttle, Kevin, "Registration of Diffusion Tensor Images in Log-Euclidean and Euclidean Space" (2019). Thesis. Rochester Institute of Technology. Accessed from

This Thesis is brought to you for free and open access by the RIT Libraries. For more information, please contact repository@rit.edu.

Registration of Diffusion Tensor Images in Log-Euclidean and Euclidean Space

by

KEVIN TUTTLE

A Thesis Submitted in Partial Fulfillment of the Requirements
for the Degree of Master of Science in Applied Mathematics
School of Mathematical Sciences, College of Science

Rochester Institute of Technology

Rochester, NY

July 31, 2019

Committee Approval:

Dr. Nathan Cahill Date
School of Mathematical Sciences
Thesis Advisor

Dr. Nathaniel Barlow Date
School of Mathematical Sciences
Committee Member

Dr. Laura Muñoz Date
School of Mathematical Sciences
Committee Member

Dr. Matthew Hoffman Date
School of Mathematical Sciences
Director of Graduate Programs

Abstract

Diffusion Tensor Imaging is a type of Magnetic Resonance Imaging that allows for the examination of brain connectivity and axonal integrity. Diffusion Tensor Images are created by capturing Diffusion-Weighted MRI images with specific RF pulses, inputting the images and the RF pulse gradient vectors into a set of equations, and solving the equations with linear algebra. To compare one DTI image with another, the images can be aligned using Image Registration. Image Registration works by defining a metric that describes the similarity between two images and iteratively transforming one of the images until the similarity measure is minimized. Existing methods of DTI comparison fit tensors to DW-MRI images, compute matrix logarithms to transform the tensors into a vector-space, register the vector-space structures, and then matrix exponentiate the results to transform them back to tensors. Logging and exponentiating the tensors introduces biases and noise so this registration framework is not ideal. Additionally, the information encoded in a diffusion tensor is a subset of the information present in the original DW-MRI images. This thesis proposes a new registration framework which avoids these shortcomings by registering the underlying DW-MRI images and then fitting the diffusion tensors to the registered DW-MRI images and transformed gradient vectors. The existing DTI registration framework and the new DW-MRI registration framework are applied to a small image set and their results are compared along a number of qualitative and quantitative attributes.

CONTENTS

I	Introduction	1
I.1	Physiological Structure of the Brain	1
I.2	Diffusion-Weighted Magnetic Resonance Imaging	2
I.3	Modeling Diffusion	4
I.4	Diffusion Tensors	6
II	Image Registration	7
II.1	Image Representation	8
II.2	Transformations	8
II.3	Similarity Measure	9
II.4	Multilevel Representation	10
II.5	Registration of DTI data	11
III	Registration of Diffusion Tensors in Log-Euclidean Space	12
III.1	Matrix Logarithm and Exponential	12
III.2	The DTI Registration Framework	13
III.3	Dealing with Anomalous Eigenvalues	13
IV	Registration of DW-MRI Image Sets In Euclidean Space	15
IV.1	Stack Registration	17
V	Comparison of Methods	18
V.1	Materials	18
V.2	Registration Results	18
V.3	Angular Distance Comparison	19
V.4	FA Comparison	19
V.5	MD Comparison	21
V.6	Atlas Comparison	23
V.7	Conclusion	23
VI	Acknowledgments	27

I. INTRODUCTION

On November 8th, 1895, Wilhelm Conrad Röntgen created the first X-ray image, showing the skeletal structure of his wife's hand. From those beginnings, radiography has grown to be a core tool used by medical experts to examine internal injuries, diagnose disease, and even monitor pregnancies. In the decades since Röntgen's initial discovery, numerous other techniques for the imaging of the human body have been invented. While some modalities are best able to capture the structural details of bones or soft tissue, others can be used to provide complementary information about the physiological activity within an organ or body. To examine a bone fracture, an orthopedic surgeon will often turn to X-rays; scrutinize a tumor in a soft tissue such as a muscle or a gland, an oncologist will utilize sonography. In many cases, the diagnosis and investigation of brain injuries and disorders requires examining the integrity of the connections between brain cells and the integrity of the brain cells' axons. A technology that is particularly well suited to imaging the brain is Magnetic Resonance Imaging (MRI). An MRI machine creates an image of the brain by exploiting the electro-magnetic properties of hydrogen atoms, which are abundant in water and fat, and thus the brain. A hydrogen atom is constantly spinning on a random axis. To generate an image, an MRI machine first creates a strong magnetic field. When placed in the magnetic field, all of the hydrogen atoms line up their rotational axes parallel with the magnetic field. The MRI machine then applies a radio frequency (RF) pulse, which knocks the atoms off alignment with the base magnetic field and into alignment with the pulsed gradient. The MRI detects those shifts of alignment, and uses them to generate an image. Different types of MRIs, utilizing different RF pulses, can be used to image the overall structure of the brain (sMRI), neural activity in different regions of the brain (fMRI), and the direction and integrity of white matter pathways in the brain (DW-MRI).

I.1 Physiological Structure of the Brain

Signals are carried from one area of the brain to another by tracts of white matter. White matter is made up of neurons with long axons. The core of an axon contains long, cylindrical structures called microtubules and neurofilaments. The axon body is surrounded by an axonal membrane. The axonal membrane is then surrounded by a sheath of fatty material called myelin. This myelination acts as insulation and facilitates the fast transfer of impulses along the length of the axons. As shown in Figure 1, the microtubules, neurofilaments, axonal membrane, and myelin all hinder water diffusion perpendicular to the length of the axon [1]. Additionally, the axons in some white matter tracts are more densely packed, which also reduces perpendicular diffusion. Axial diffusivity (λ_{\parallel})




Figure 1: A simplified longitudinal view of the structure of an axon. $D(\parallel)$ denotes the parallel diffusion coefficient and $D(\perp)$ denotes the perpendicular diffusion coefficient. Taken from [1].

refers to water movement parallel to the axon tract and radial diffusivity (λ_{\perp}) refers to water movement perpendicular to the axon tract. The fraction of the total diffusion that can be attributed to the orientation of the axon body is described by fractional anisotropy (FA), which ranges from 0 (isotropic, unrestricted multi-directional diffusion) to 1 (diffusion completely restricted to one directional axis).

I.2 Diffusion-Weighted Magnetic Resonance Imaging

The relationship between diffusion and axonal structure leads to a method of imaging and assessing axons in the brain. Diffusion-weighted Magnetic Resonance Imaging (DW-MRI) is a form of MRI whose RF pulse generates images where the brightness of each voxel reflects the rate of diffusion at that location with respect to the RF pulse's gradient vector (b-vector). After the capture of multiple images with different diffusion gradients, a diffusion tensor can be computed, which encodes diffusion direction and intensity data at every voxel in a symmetric, positive definite matrix. Once derived, the diffusion tensor for each voxel can be used to compute the axial diffusivity, radial diffusivity, and fractional anisotropy at each voxel's location. By comparing the diffusion tensors and derived metrics of images taken before and after a possible injury, changes in brain connectivity and axonal integrity can be determined. As shown in Figure 2, different changes in axial and radial

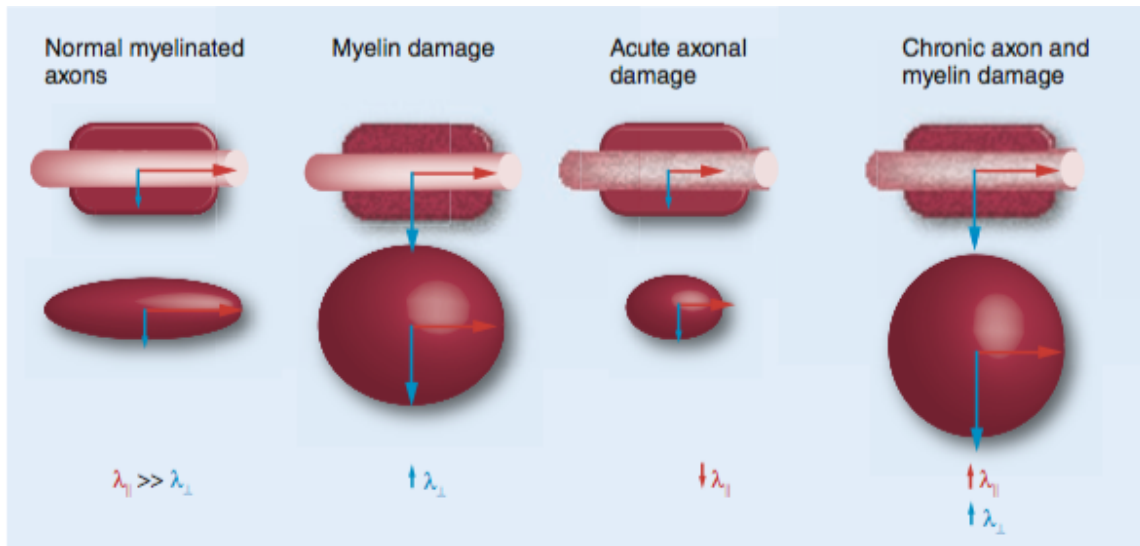


Figure 2: Top row: longitudinal diagram of an axon. Bottom row: associated diffusion ellipsoids and description of the relative values and changes to axial diffusivity ($\lambda_{||}$) and radial diffusivity (λ_{\perp}). In a normal myelinated axon, axial diffusivity is much greater than radial diffusivity. Damage to myelin causes an increase in radial diffusivity. Damage to the internal structures of the axon can impair axonal transport and reduce axial diffusivity. With chronic axonal injury, gliosis and inflammation increase axial diffusivity. Taken from [2].

diffusivity suggest different types of axonal injury. Often, a pre-injury image is not available to make comparisons with. In these cases it can be useful to compare the post-injury images with an atlas, or standard model. The simplest form of atlas is an average of many images from a representational population.

I.3 Modeling Diffusion

Diffusion can be represented as a Brownian random process [3]. The probability that a molecule starting at x_0 at $t = 0$ reaches position x (where $x \in \mathbb{R}^3$) at time $t = \tau$ is given by the Gaussian conditional probability density function:

$$\rho(x|x_0, \tau) = \frac{1}{\sqrt{|D(\tau)|(4\pi\tau)^3}} \exp\left[\frac{-(x - x_0)^T D^{-1}(\tau)(x - x_0)}{4\tau}\right] \quad (\text{I.1})$$

The anisotropy of the medium in which the diffusion occurs is described by the covariance matrix, D . The equidensity contours, or confidence area boundaries, of the Gaussian multivariate distribution are ellipsoids. To find an equidensity contour, or confidence interval, we would set the conditional pdf equal to a constant. Since the x and x_0 only appear in the exponent term, solving for the set of values which hold the condition pdf equal to a constant is equivalent to solving for the set of values which hold the exponent term equal to some constant c . If we take $c = \frac{1}{2}$, we are left with a quadratic form:

$$\frac{(x - x_0)^T D^{-1}(\tau)(x - x_0)}{2\tau} = 1 \quad (\text{I.2})$$

Here we can note that, D , being a covariance matrix, is symmetric and positive semi-definite. Assuming for a moment that D is not just positive semi-definite, but positive definite, D then has positive eigenvalues λ_1 , λ_2 , and λ_3 , and their associated eigenvectors v_1 , v_2 , and v_3 are orthogonal. These eigenvectors can be written in matrix form as $E = (v_1|v_2|v_3)$. Then using the change of basis

$$x' = E^T(x - x_0) \quad (\text{I.3})$$

and expanding out, we get

$$\left(\frac{x'_1}{\sqrt{2\lambda_1\tau}}\right)^2 + \left(\frac{x'_2}{\sqrt{2\lambda_2\tau}}\right)^2 + \left(\frac{x'_3}{\sqrt{2\lambda_3\tau}}\right)^2 = 1 \quad (\text{I.4})$$

In this new basis, (I.4) describes the mean displacement of water particles at time τ for the diffusion process modeled by the tensor D . We see in (I.4) that the eigenvalues λ_1 , λ_2 , and λ_3 are proportional to the lengths of the effective diffusion ellipsoid. Looking at (I.4) we can also see that negative eigenvalues would lead to imaginary principal axis lengths. Similarly, an eigenvalue of zero would




Figure 3: Diffusion ellipsoids on a sagittal slice through the ventricles of a brain. The color of the ellipsoid corresponds to the orientation of the largest principal axis. The ventricles are cavities in the brain containing free-moving cerebrospinal fluid. We would expect voxels in the interior of the ventricles to display completely isotropic diffusion, with all three eigenvalues being equal. Noise in the image has caused 1 of the eigenvalues at each of these voxel to be just slightly larger than the others, leading to the random rainbow of colors. Image generated by FSL with data from the MGH-USC dataset.

necessitate an associated zero vector, which would collapse the ellipsoid. Therefore in order for the diffusion tensor to describe a physically realistic diffusion process, it must be positive definite. Rendering the diffusion tensors as ellipsoids is an effective way to visualize the diffusion being modeled by the tensors, as shown in Figure 3.

I.4 Diffusion Tensors

After capturing a multitude of diffusion weighted MRI images, a diffusion tensor D can be estimated by solving the Stejskal-Tanner equation set [5]:

$$\begin{aligned} S_1 &= S_0 e^{-bg_1^T D g_1} \\ &\quad \vdots \\ S_k &= S_0 e^{-bg_k^T D g_k} \end{aligned} \tag{I.5}$$

where $k =$ the number of diffusion weighted images captured, S_n ($n \in [1 \dots k]$) is the diffusion weighted image captured with the diffusion gradient g_n , and S_0 is the baseline image captured without a diffusion gradient. Dividing by S_0 and taking the logarithm of both sides of (I.5) yields

$$\log\left(\frac{S_k}{S_0}\right) = -bg_k^T D g_k \tag{I.6}$$

Then moving all of the terms to one side gives us

$$\log\left(\frac{S_k}{S_0}\right) + bg_k^T D g_k = 0 \tag{I.7}$$

Linear regression can then be used to solve for D by minimizing the following sum of squared differences error:

$$D = \arg \min_{T \in S} \sum_k \left(\log\left(\frac{S_k}{S_0}\right) + bg_k^T T g_k \right)^2 \tag{I.8}$$

Where S is the space of symmetric tensors. As established previously, the matrices making up the diffusion tensor D are symmetric and positive definite. Thus, six or more diffusion gradients must be used in order to solve for the 6 degrees of freedom of D .

II. IMAGE REGISTRATION

Comparison of pre- and post-trauma images and comparison against (and creation of) an atlas requires that all the images be properly aligned with one another. Done manually, this alignment work would be time consuming and error-prone. In image registration, one image is designated the reference, or fixed, image and the other the template, or moving, image. The registration process transforms the template image to align it with the reference. An example can be seen in Figure 4.

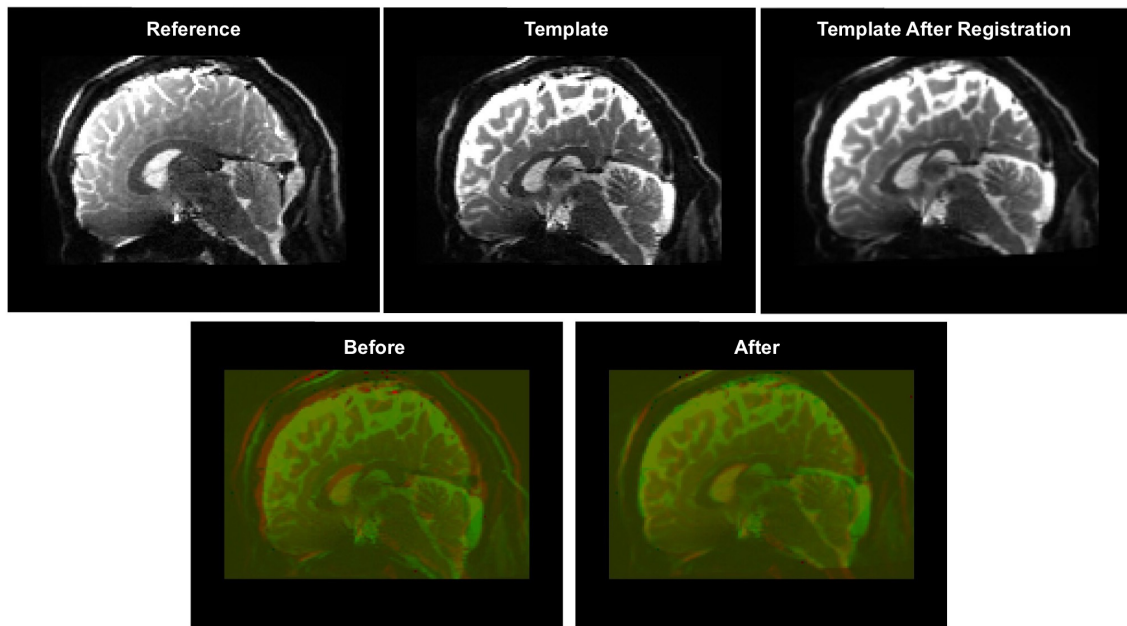


Figure 4: Sagittal slices from sMRI images of 2 brains. The images were registered in 3D with affine transformation. Before registration the brain in the template image (green overlay) is smaller and has less anterior tilt than the reference. After registration, the template image has been scaled up and rotated to align with the reference. The difference can be seen most acutely at the outer contours of the brain, the ventricles, and the orientation of the brainstem.

Image registration algorithms are designed to compute the transformation that optimizes the similarity between images. An iterative approach to maximizing the similarity proceeds as follows:

- Step 1. Apply a transformation to the template image (often starting with the identity transformation).
- Step 2. Calculate the similarity between the reference image and the transformed template image.

- Step 3. Repeat steps 1 & 2 until the similarity has been maximized.

Designing an image registration appropriate to a particular task involves defining a number of things: the image representation, the transformation, the similarity measure, and the optimization procedure, each of which will be described herein.

II.1 Image Representation

Given a reference image R and a template image T , we must first decide upon how to represent them. While computer science commonly represents an image as array of color values, from a mathematical standpoint we are better served by defining it as a function. Thus letting $d \in \mathbb{N}$ denote the number of spatial dimensions of the images we have $R, T : \mathbb{R}^d \rightarrow \mathbb{R}$, where $T(x)$ gives the template image's grey value at spatial position x . The images have finite size, and so are contained within a bounded domain Ω . We will assume R and T to be constant for all $x \notin \Omega$. All images are discrete reproductions of their subjects, with the granularity determined by the imaging device's resolution. Thus a naïve function representation which solely returns the values gathered by the imaging device can only ever be defined over a finite subset of the image domain. Since the optimization process will be greatly aided by having a continuously differentiable objective function, we will use interpolation (trilinear in the 3-dimensional case) to determine the image values at all spatial locations for which data was not provided by the imaging device.

II.2 Transformations

In order to find the transformation that maximizes the similarity measure, we must define a consistent method of transforming the image. The simplest transformation method commonly used in medical image registration is rigid transformation, which allows for translation, rotation, and reflection. In \mathfrak{R}^3 :

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ 1 \end{bmatrix} = \begin{bmatrix} \cos w_3 \cos w_2 - \sin w_3 \sin w_1 \sin w_2 & -\sin w_3 \cos w_1 & \cos w_3 \sin w_2 + \sin w_3 \sin w_1 \cos w_2 & w_4 \\ \sin w_3 \cos w_2 + \cos w_3 \sin w_1 \sin w_2 & \cos w_3 \cos w_1 & \sin w_3 \sin w_2 - \cos w_3 \sin w_1 \cos w_2 & w_5 \\ -\cos w_1 \sin w_2 & \sin w_1 & \cos w_1 \cos w_2 & w_6 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{bmatrix} \quad (\text{II.1})$$

Where \vec{x} are the original coordinates, \vec{y} are the transformed coordinates, and \vec{w} are the parameters of the transformation. Rigid transformation preserves the distance between every point in an image.

There are circumstances in which rigid transformation is useful, such as aligning the 3D MRI scans of a patient’s head taken at two different times. The skull is a rigid object, so assuming no deformations due to tumor growth, injury, etc., the differences in the MRI scans taken at separate capture times will be due to varying location and orientation of the skull within the imaging device, and therefore should be well-described by a 3D rigid transformation. The registration of MRIs from one subject to a different subject would not be, however. If we drop the point distance-preserving property of the rigid transformation, we get the affine transformation. Affine transformations are a superset of rigid transformations; In addition to all the transformations afforded by rigid transformation, affine transformation also allows for shearing and scaling. In \mathbb{R}^3 :

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ 1 \end{bmatrix} = \begin{bmatrix} w_1 & w_2 & w_3 & w_4 \\ w_5 & w_6 & w_7 & w_8 \\ w_9 & w_{10} & w_{11} & w_{12} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{bmatrix} \quad (\text{II.2})$$

This additional freedom allows for the registration of images across patients, where the size and geometry of the captured structures may vary to some degree. Furthermore, from a computational point of view, affine transformations are easy to utilize in optimization problems since they are linear with respect to their parameters. Since rigid transformations are not linear with respect to their parameters, rigid registration requires special care in computing gradients of the similarity measure with respect to the transformation parameters.

II.3 Similarity Measure

One of the most straight-forward and computationally efficient similarity measures is the Sum of Squared Differences (SSD). Defined continuously,

$$\text{SSD}(T, R) = \frac{1}{2} \int_{\Omega} (T(x) - R(x))^2 dx \quad (\text{II.3})$$

And thus the discretized objective function for registration with SSD is defined as

$$\text{SSD}(T, R) = \frac{1}{2} \sum_{\vec{x} \in \Omega} (T(\vec{x}) - R(\vec{x}))^2 \quad (\text{II.4})$$

SSD is based on the assumption that when two images are perfectly aligned, the color value at corresponding pixels (or voxels) will be the same. This principle does not hold for multimodal medical image registration (for instance, registering a CT scan to an MRI), but it performs well for intra-modal registration. The $\frac{1}{2}$ coefficient at the front of the SSD definition is a convention which

makes the computation of the the similarity measure's derivative cleaner (the 2 in the square term cancels it out). There is a problem that can arise when using SSD for image registration: as the transformation being evaluated changes, so to does the area of overlap between the images. Including values outside of the area of overlap means biasing the overall SSD value with information which is not pertinent. And restricting the sum to be over the overlap region also has potential pitfalls. If you define T and R to have the same constant value outside of Ω , then a transformation which pushes T entirely off of R so that they do not overlap at all would have an SSD value of zero, giving us a solution which is unhelpful. To address this issue, it's preferred to use the Mean of Squared Differences (MSD), where the SSD measure is divided by the size of the overlap region. There are a whole host of other, more complex similarity measures that have been devised for a variety of different purposes. For multi-modal registration, which assumes that the same structures are visible but potentially at different intensities, Mutual Information (MI) is a standard similarity measure which borrows from information theory to quantify the amount of information shared between R and T [6, 7]. There are many variants of MI including Normalized Mutual Information and Adaptive Local Mutual Information[8]. Other measures, such as Normalized Gradient Field (NGF), compare the image derivatives [9]. This listing is by no means exhaustive, as such an undertaking is beyond the scope of this thesis.

II.4 Multilevel Representation

Image registration can be optimized with the use of a multilevel representation scheme [4]. With a multilevel scheme, registration begins on smaller versions of the images, and the resulting transformation is utilized as the starting point when the process is repeated on larger, higher resolution versions of the images. The set of images at different resolutions is known as an image pyramid. The base level of the pyramid is the original full resolution image. Subsequent levels can be generated by applying a gaussian blur filter and down-scaling. With the gaussian blur filter, each voxel in an image corresponds to the local average of the voxel's neighborhood on the lower level of the pyramid. The benefits of a multilevel scheme are two-fold. First, by beginning the registration process on coarser images, only the global features of the images will be matched to one another, with more and more local features being taken under consideration by the similarity measure as the levels are traversed. This helps protect against the optimization of the similarity measure getting stuck at a local minima and returning a sub-optimal solution. The second is computational efficiency. It is cheaper to perform the registration on smaller images, and the registration will converge on a final solution at the lower levels more quickly when the process is seeded with a better initial

guess.

II.5 Registration of DTI data

To compare or make an atlas from DTI images, it is common to register the sets of diffusion tensors constructed from the diffusion weighted images. When applying geometric transformations to a diffusion tensor, the directional information encoded in the tensor must also be transformed to match. This often introduces physically unrealistic defects, as shown in Figure 5. To avoid these defects, tensors can be registered with the Log-Euclidean Framework [10]. Before registration, the tensors are transformed into the Log-Euclidean space and the results of the registration are then translated back to the Euclidean Space. Registering the diffusion tensors allows for more information to be taken into consideration than simply registering a single derived value such as FA. However, a great deal of information is lost at the step of fitting diffusion tensors to the diffusion-weighted MRI images. The focus of this thesis is to present a method of DTI registration that registers the original diffusion-weighted MRI images and then fits tensors to the transformed registration results.



Figure 5: Interpolation between two tensors, which can be thought of as a template tensor (far left) and a reference tensor (far right). Top: linear interpolation. Bottom: Log-Euclidean interpolation. Note the swelling that occurs in the case of linear interpolation. Taken from [10].

III. REGISTRATION OF DIFFUSION TENSORS IN LOG-EUCLIDEAN SPACE

III.1 Matrix Logarithm and Exponential

In order to register diffusion tensors, we need to be able to compute a similarity measure. Thus we must define a distance measure on symmetric matrices. This can be done many ways. One common such metric is the Euclidean distance $\text{DistEuc}(S_1, S_2) = \|S_1 - S_2\|_F$ where $\|\cdot\|_F$ denotes the Frobenius norm. Although relatively easy to compute, the distance metric and associates of the Euclidean family have a few key drawbacks [10]. Firstly, they do not guarantee positive definite matrices. Since all the eigenvalues of a diffusion tensor must be positive, this means that Euclidean metrics lead to matrices which can not be interpreted as valid diffusion tensors. Secondly, the Euclidean metrics do not preserve the matrix determinant. The determinant is the product of the eigenvalues, and the eigenvalues are the lengths of the principal axes of the confidence area ellipsoid. Looking again at (I.4), the volume of the ellipsoid is proportional to the square root of the determinant. Carried back to what the ellipsoid represents, the determinant is a measure of the diffusion's dispersion, or total amount of diffusion (the volume of the ellipsoid). So when the Euclidean metrics inflate the determinant, they introduce more dispersion, which is physically unrealistic. These problems with the Euclidean metrics arise because the set of symmetric positive definite matrices do not form a vector space, but rather a Riemannian manifold. Another measure related to the dispersion is Mean Diffusivity (MD), defined as the average of the three eigenvalues:

$$MD = (\lambda_1 + \lambda_2 + \lambda_3)/3 \tag{III.1}$$

To avoid the tensor swelling caused by Euclidean metrics, Arsigny et al. introduced the Log-Euclidean metrics. The Log-Euclidean distance is:

$$\text{dist}_{LE}(S_1, S_2) = \|\log(S_1) - \log(S_2)\|_F \tag{III.2}$$

where \log refers to the matrix log. The matrix logarithm and exponential are the key. For a matrix S , the matrix exponential is given by $\exp(S) = \sum_{k=0}^{\infty} \frac{S^k}{k!}$. The matrix logarithm is defined as the inverse of the matrix exponential. Importantly, the matrix logarithm of a tensor matrix is symmetric, as can be seen by examining the domain and codomain of the natural logarithm. Since the matrix exponential is the inverse of the matrix logarithm, the matrix exponential of any symmetric matrix is also guaranteed to yield a unique positive definite symmetric matrix. The computation of the matrix logarithm of a tensor matrix S can be obtained in three steps [10]:

1. Diagonalize S into a rotation matrix R and a diagonal matrix D such that the eigenvalues of S lie along the diagonal of D . $S = R^T D R$
2. Replace each eigenvalue (which is guaranteed to be positive) with its natural logarithm to construct a new matrix \tilde{D}
3. Recompose R with \tilde{D} to get L , the logarithm of S : $L = \log(S) = R^T \tilde{D} R$

After transforming the positive semidefinite tensor matrix S into a symmetric matrix L , we can simplify the computation of the Log-Euclidean distance further by converting L into a vector \vec{c} such that $\|L\|_F = \|\vec{c}\|_2$, where $\|\cdot\|_2$ denotes the L_2 norm: $\vec{c} = (L_{11}, L_{22}, L_{33}, \sqrt{2}L_{12}, \sqrt{2}L_{13}, \sqrt{2}L_{23})^T$. We now have a procedure whereby the computation of the Log-Euclidean distance between two tensor matrices A and B can be simplified to the computation of Euclidean distance between vectors \vec{a} and \vec{b} .

III.2 The DTI Registration Framework

Having established all the necessary tools, we can now examine the entire flow of registering Diffusion-Weighted MRI images in the DTI Registration Framework, shown at the top Figure 6.

The process begins by fitting diffusion tensors to the individual stacks of DW-MRI images. Each diffusion tensor is then transformed into a multi-dimensional array of vectors in the logarithmic domain. These arrays can be registered by modifying the standard SSD distance measure to accommodate vectors rather than scalars (more discussion in section IV). After registration, the arrays are transformed back into diffusion tensors using matrix exponentiation. If we are constructing an atlas by registering and averaging several DW-MRI stacks, we will average the registered multidimensional arrays prior to exponentiation.

III.3 Dealing with Anomalous Eigenvalues

We previously established that in order for a tensor matrix to describe a physically realistic diffusion process, and in order for the matrix to have a matrix logarithm, the matrix must have positive eigenvalues. In practice however, commonly used methods of solving the Stejskal-Tanner equations do not enforce the positive eigenvalue condition. In this case, noise in the DW-MRI images can cause the calculated diffusion tensor matrices to have negative eigenvalues. This often occurs in areas of high diffusion anisotropy such as the corpus callosum, as shown in Figure 7 [11]. To enable

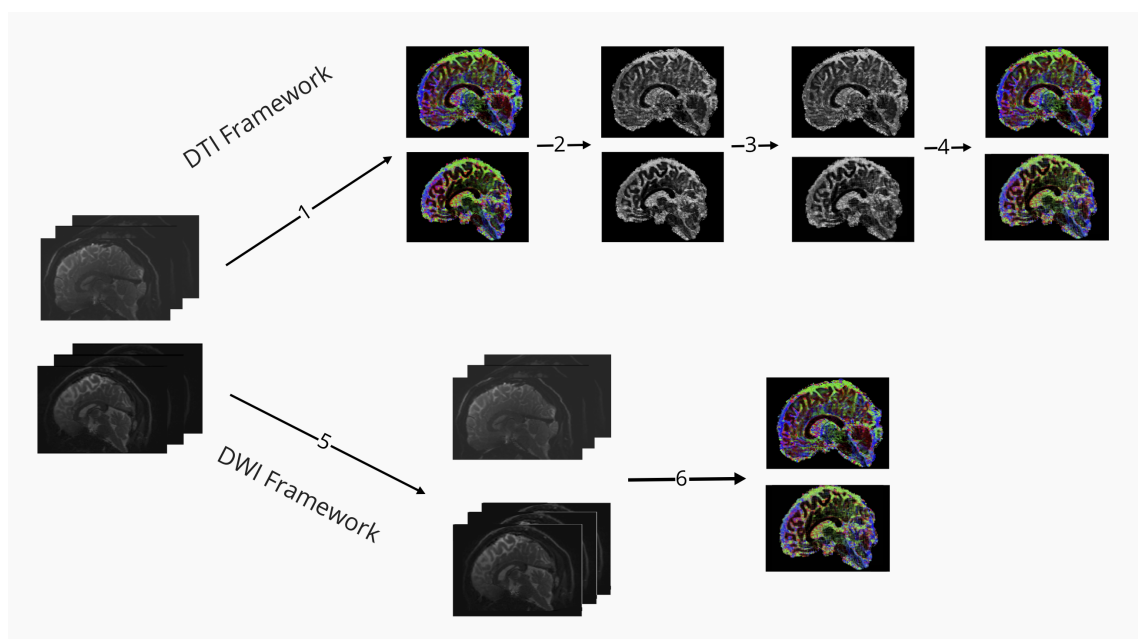


Figure 6: The DTI Framework: 1) Fit diffusion tensors to the unregistered DW-MRI images. 2) Take the log matrix of the diffusion tensor matrices. 3) Register the tensor log matrices. 4) Exponentiate the registered tensor log matrices to obtain registered diffusion tensors. The DW-MRI (Stack) Registration Framework: 5) Register the DW-MRI image stacks. 6) Fit diffusion tensors to the registered DW-MRI images to obtain registered diffusion tensors.

the usage of the DTI Framework, negative eigenvalues in the tensor matrices can be replaced with a fraction (such as 10%) of the smallest positive eigenvalue present in the matrices before taking the logarithm.




Figure 7: λ_1 map of an image overlaid with a mask indicating the voxels with at least one negative eigenvalue (red). The largest concentration of anomalous voxels occurs in the corpus callosum.

IV. REGISTRATION OF DW-MRI IMAGE SETS IN EUCLIDEAN SPACE

When fitting diffusion tensors to a set, or stack, of DW-MRI images, the vast array of information (and noise) present in the the images is averaged and simplified. Calculating a measure such a Fractional Anisotropy further simplifies and reduces the amount of encoded information. Registering images with lower information levels is faster, but the maximum level of accuracy is also reduced, like comparing the black and white versions of color photos. Examining more information-heavy images allows the similarity measure to make a more nuanced and detailed analysis of the differences between images, just like registering on lower levels of a multilevel scheme allow for more micro aspects of the images to be utilized. So while registering tensors leads to better results than registering a derived measure such as FA, registering the initial DW-MRI image stacks offers even more performance gains. Furthermore, transforming the matrix logs of a diffusion tensor and then exponentiating does not result in exactly the same final tensor as applying the same transformation to the original DWI-MRI image stack and fitting tensors to the transformed image stack. As an example, I took a random Affine transformation and applied it to the matrix logs of a diffusion tensor and the original DWI-MRI image stack. At first glance, the final tensors derived via each method are very similar, (see the colorized v_1 maps in Figure 8) and indeed at a macro-level they are perfectly aligned. But a closer examination shows that they are subtly different. We can quantify the difference by calculating the angular distance between the largest eigenvectors of the two tensors. Angular distance is a similarity measure commonly used to compare vectors in higher-dimensional




Figure 8: Colorized v_1 maps of diffusion tensors derived after having the same random Affine transformation applied to (top row) the tensor matrix logs and (bottom row) the original DWI-MRI image stack.

spaces defined as follows:

$$\text{Angular distance} = \frac{\cos^{-1}\left(\frac{A \cdot B}{\|A\| \|B\|}\right)}{\pi} \quad (\text{IV.1})$$

where A and B are vectors. Angular difference ranges from 0 (the two vectors are the same) to 1 (the vectors are oriented totally opposite). The mean angular difference of the largest eigenvectors of the two tensors was 0.197845091544. So although the morphological elements of the brain are aligned regardless of when the transformation is applied, the eigenvectors of the resulting diffusion tensors are slightly different.

IV.1 Stack Registration

The traditional image registration process requires only a few modifications in order to accommodate registering stacks of images against one another. The DWI-MRI (Stack) Registration Framework is shown at the bottom of Figure 6. We will first assume that the two sets of images we are setting out to register were acquired in accordance with the same scanning procedure, using the same b-values and gradient directions. We will also assume for the sake of simplicity that the image sets do not require any intra-set registration. Each image stack can then be represented by a multi-dimensional array of vectors, where each i -th element of a vector is the value of the i -th image at that voxel. The SSD similarity measure can be modified as follows:

$$SSD = \frac{1}{2} \sum_{i=0}^n \sum_{\vec{x} \in \Omega} (T_i(\vec{x}) - R_i(\vec{x}))^2 \quad (\text{IV.2})$$

During the transformation step of the registration process, we will apply the transformation to each image in the moving set individually. Once we have registered the image stacks, we must transform the image capture gradient b-vectors. The b-vector is constant throughout the image capture space, so translation does not need to be accounted for. We obtain the b-vectors of the registered image stack by multiplying the original b-vectors by the un-augmented matrix for the registration's resulting transformation. After multiplication, we re-normalize the resulting vectors. Tensor fitting can then proceed as usual with the resulting transformed image stack and matching b-vectors. If creating an atlas, we will average the registered image sets and b-vectors, and then fit tensors to the average image and b-vector set.

V. COMPARISON OF METHODS

V.1 Materials

To register DW-MRI images in the DTI framework and register DW-MRI image sets in the new DWI framework, I used and modified Dipy, an open source python package for diffusion magnetic resonance imaging analysis [12]. To fit diffusion tensors to DW-MRI image sets I used FSL, a software library created by the Analysis Group at the Oxford Centre for Functional MRI of the Brain (FMRIB) [13, 14]. I ran comparisons on nine DW-MRI image sets from the MGH–USC Adult Diffusion Dataset of the Human Connectome Project[15]. The MGH–USC Adult Diffusion Dataset contains MRIs for 35 healthy adults, both male and female, ranging in age from 20 to 59. The images I worked with were pre-processed with FSL to remove Eddy current distortion. For the sake of computational expedience, I worked solely with the images’ 1K b-value shells. I computed the DTI framework registrations locally on a MacBookPro with 16GB of RAM and a dual-core 2.5 GHz Intel Core i7 processor. To carry out the DW-MRI framework registrations I used AWS Fargate, so the exact hardware specs are unknown, but likely comparable. Each registration under the DTI framework took about 20 minutes, and each registration under the DW-MRI framework took about 7.5 hours. The runtime for the DW-MRI framework could be brought down significantly by parallelizing some of the computations, as my implementation was very naive and did everything sequentially.

V.2 Registration Results

I designated one DW-MRI image set (MGH_1007) the reference and registered eight other DW-MRI image sets to it using both the traditional DTI registration framework and the new Stack registration method. I chose the reference image at random. In a clinical setting, you would typically use the earliest acquired image as the reference. In the case of pre- and post-injury image registration you would use the pre-injury image as the reference. I also constructed atlases out of both sets of registered images. With both registration methods I utilized a multilevel representation scheme with 3 levels. The three levels were created with a Gaussian blur kernel with sigmas 0, 1, and 3, and downsampling factors of 1, 2, and 4 (the lowest level of the image pyramid was 1/4 the size of the original image). At each level, I registered the images in three rounds. The first registration round used translation transformations, the second used rigid transformations, and the third used affine transformations, with the result of the previous round used as the starting point for the next.

The similarity optimization was capped at 10,000 iterations at the top level, 1000 iterations at the second level, and 100 iterations at the base level.

V.3 Angular Distance Comparison

Based on a visual inspection, the DWI registration method appears to provide better results in many cases. For example, Figure 9 shows one case where the DTI registration method resulted in a transformed Template image where the brain is pitched back further than the brain in the Reference image to the extent that the brain stem and cerebellum are distinctly mis-aligned. We can compare the registration results quantitatively by examining the Angular Distance between the v_1 maps of the Reference image and the Template images after being registered via each method (Table 1). The Stack registration method produced a smaller average Angular Distance in all cases.

Table 1: The average Angular Difference between the v_1 maps of the Reference image and the Template images after DTI registration and Stack registration

Patient Number	DTI	Stack
1001	0.49389660299861	0.43174498561832
1002	0.49016711190169	0.42681112994145
1003	0.47995090144764	0.40748124494888
1004	0.49222712102017	0.42857044411107
1010	0.48490892897584	0.40967580454875
1016	0.48482177482763	0.41747987514091
1019	0.48328977110275	0.41013182636671
1031	0.48346228590888	0.41390373274335

V.4 FA Comparison

Fractional Anisotropy varies in any given region of the brain from patient to patient, even among the population of healthy adults. So no inter-patient registration will ever result in tensors whose FAs perfectly match. But the shared underlying structural architecture of all healthy brains do provide a means to compare the results of inter-subject image registration using FA. Regions with high structural order (such as the corpus callosum) will always have larger FA than regions with a moderate amount of structural orientation (white matter throughout the cerebrum and




Figure 9: Corresponding slices of the FA map of the Reference image (top row), Template image overlaid on Reference image after DTI registration (center row), Template image overlaid on Reference image after Stack registration (bottom row)

cerebellum) which in turn will have larger FA than regions with little to no structural orientation (grey matter and cerebrospinal fluid). Therefore, a better inter-subject registration result, which aligns characteristically high-FA regions in the Reference with characteristically high-FA regions in the Template and characteristically low-FA regions in the Reference with characteristically low-FA regions in the Template, will have a smaller FA SSD between the Reference and the Template than a less-accurate registration result. With that in mind, I calculated the SSD of the FA between the reference image and each of the template images after both registration methods (Table 3). Based on this metric, the Stack registration framework yielded a better result for all of the the template images; on average the FA SSD obtained after DTI registration was 16% larger than the FA SSD obtained after Stack registration. Taking the sets of FA values of the result images generated by the two methods and applying Wilcoxon signed-rank tests showed that the registration methods generated results whose median FAs were statistically significantly different (Table 2).

Table 2: Wilcoxon signed-rank results for FA value sets

Patient Number	Statistic	P-Value
1001	51052864248.0	0.0
1002	57399475821.0	0.0
1003	77007570033.0	1.25329738407e-84
1004	70955730244.0	0.0
1010	66633502479.0	0.0
1016	77338675894.0	5.1892702647e-280
1019	62600478970.0	0.0
1031	83149661861.0	7.84663674525e-39

V.5 MD Comparison

Comparison of the MD produced by the two registration methods is complicated by the presence of negative eigenvalues. FSL does not enforce the positive-eigenvalue constraint, and all the image tensors I worked with had some degree of negative eigenvalues, ranging from .45% of the voxels to .79% of the voxels with a mean of .57% of the voxels (Table 4). In order to register the images with the DTI framework, I replaced all the negative eigenvalues with a value of ten percent of the smallest positive eigenvalue present in the image. The Stack registration method does not require any corrective pre-processing of the negative eigenvalues. In fact, the process of transforming

Table 3: The Sum of Squared Differences in FA between the reference and template images before registration and after registration via the two methods.

Patient Number	Original	LogEuc	Stack
1001	0.01941060176512011	0.01769212491769468	0.01533814929780506
1002	0.01862004078975340	0.01620056082917935	0.01522300564512915
1003	0.01838521114012011	0.01839719073955072	0.01484173936908748
1004	0.01944132486979167	0.01772994636649026	0.01543133171237245
1010	0.01848105709568984	0.01661073085043165	0.01406338438695791
1016	0.01872617552880527	0.01803354691124977	0.01618971428903593
1019	0.02098667066924426	0.01634817653770409	0.01403901080695950
1031	0.02023369198753720	0.01815474131646657	0.01491951767279177

Table 4: The percentage of voxels in each image tensor with at least one negative eigenvalue.

Patient Number	Percentage of Voxels With a Negative Eigenvalue
1001	0.78613945578
1002	0.59954294218
1003	0.54262329932
1004	0.67144982993
1007	0.53656462585
1010	0.47108843537
1016	0.54730017007
1019	0.45307185374
1031	0.51642219388

and interpolating the image stacks with the registration’s resulting transformation reduced the number of voxels in the derived tensor in every case, even when the transformation included positive scaling. The DTI framework magnified the problem of the anomalous eigenvalues. Although the eigenvalues of the registration results are all positive, given that they are restricted to the co-domain of the exponential function, the registration introduced anomalously large eigenvalues. Furthermore, the number of voxels with extremely large eigenvalues is more than you would expect given the estimation method of applying the registration transformation to a mask of the negative eigenvalues in the original image. The MD of the original images and the Stack registered images ranges from -0.0072 to 0.0045 , and as a result of the amplified noise the MD of the DTI registered images ranges

from 0 all the way up to 1 , so no meaningful comparison of the final MD SSD can be drawn.

V.6 Atlas Comparison

Lastly, I generated diffusion tensor atlases with both registration methods (Figure 11). The atlas generated via the Stack registration has lower Fractional Anisotropy than the average across the original, unregistered images and the DTI generated atlas (Table 5). The DTI-derived atlas' FA is larger than the average across the original images. Looking at the FA maps, fewer features are visible in the Stack-derived atlas, but those that are visible are much smoother. This smoothness is also apparent when viewing the v_1 maps of the two atlases (Figure 12). The Stack registration method leads to a less noisy, more internally coherent tensor atlas. There are multiple, competing instances of macro-level features present in the DTI atlas. In particular, there is a second brain border visible especially towards the posterior of the image. This is made very apparent when examining the Laplacian (edge detecting filter) of the DTI atlas (Figure 13).

V.7 Conclusion

The Stack registration method performed better according to a comparison of final Angular Distance and Fractional Anisotropy SSD. Additionally, there is qualitative evidence to suggest that the DTI framework leads to images that are mis-aligned at a macro-level. The main drawback of the Stack registration framework is that it is more computationally intensive. Further work is possible to examine ways in which the computational time can be brought down with parallelization and hyperparameter optimization with the number of levels in the multilevel scheme. Another avenue of further inquiry is an examination of the interplay between Stack registration accuracy and efficiency when using more b-value shells or a smaller subset of images from the DW-MRI image set.




Figure 10: Corresponding slices of the MD map of MGH_1004 after Stack registration (top row), MD map of MGH_1004 after DTI registration (center row), MD map of MGH_1004 after DTI registration with the same brightness and contrast settings as the top row after a rough anomalous-voxel mask (red) of all values greater than $e^{-5.5}$ is applied (bottom row).




Figure 11: Corresponding slices of the FA maps of the DTI-derived atlas (top row), and the Stack-derived atlas (bottom row).




Figure 12: Corresponding slices of the colorized v_1 maps of the DTI-derived atlas (top row), and the Stack-derived atlas (bottom row).



Figure 13: Laplacian of the DTI-derived atlas

Table 5: The average Fractional Anisotropy of each unregistered image and the atlases

Patient Number	Mean FA
1001	0.06026357
1002	0.06242111
1003	0.07432495
1004	0.07149965
1007	0.07866789
1010	0.07040649
1016	0.07341061
1019	0.05969153
1031	0.07435881
DTI Atlas	0.08006171
Stack Atlas	0.04718261

VI. ACKNOWLEDGMENTS

Data were provided by the Human Connectome Project, MGH-USC Consortium (Principal Investigators: Bruce R. Rosen, Arthur W. Toga and Van Wedeen; U01MH093765) funded by the NIH Blueprint Initiative for Neuroscience Research grant; the National Institutes of Health grant P41EB015896; and the Instrumentation Grants S10RR023043, 1S10RR023401, 1S10RR019307.

REFERENCES

- [1] Beaulieu, C 2002. The Basis of Anisotropic Water Diffusion in the Nervous System - a Technical Review. *NMR in Biomedicine* 15(7-8): 435-455. doi: 10.1002/nbm.782
- [2] Aung, W. Y., S. Mar, and T. L. S. Benzinger 2013. Diffusion Tensor MRI as a Biomarker in Axonal and Myelin Damage. *Imaging Med.* 5(5): 427-440. doi: 10.2217/IIM.13.49
- [3] Basser, P.J., D. LeBihan, and J. Mattiello 1994. MR Diffusion Tensor Spectroscopy and Imaging. *Biophysical Journal* 66(1): 259-267. doi: 10.1016/S0006-3495(94)80775-1
- [4] Haber, E., J. Modersitzki 2006. A Multilevel Method for Image Registration. *SIAM Journal on Scientific Computing* 27(5): 1594-1607. doi: 10.1137/040608106
- [5] Niethammer, M., R. S. Jose Estepar, S. Bouix, M. Shenton and C. F. Westin 2006. On Diffusion Tensor Estimation *Engineering in Medicine and Biology Society 28th Annual International Conference of the IEEE, New York, NY, 2006*: 6707-6710. doi: 10.1109/IEMBS.2006.260927
- [6] Viola, P., W. M. Wells III 1997 Alignment by Maximization of Mutual Information *International Journal of Computer Vision* 24: 137-154. doi: 10.1023/A:1007958904918
- [7] Maes, F., A. Collignon, V. Dirk, G. Marchal, and P. Suetens 1997. Multimodality Image Registration by Maximization of Mutual Information. *IEEE Trans. Med. Imaging* 16: 187-198 doi: 10.1.1.85.1157
- [8] Nigris, D. D., L. Mercier, R. D. Maestro, D. L. Collins, and T. Arbel 2010. Hierarchical Multimodal Image Registration Based on Adaptive Local Mutual Information *Med Image Comput Comput Assist Interv.* 13(Pt 2): 643-651. doi: 10.1007/978-3-642-15745-5_79
- [9] Haber, E., J. Modersitzki 2006. Intensity Gradient Based Registration and Fusion of Multimodal Images. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006* 726-733. doi: 10.1007/11866763_89
- [10] Arsigny, V., P. Fillard, X. Pennec, and N. Ayache. 2006. Log-Euclidean Metrics for Fast and Simple Calculus on Diffusion Tensors. *Magn Reson Med* 56: 411–421. doi: 10.1002/mrm.20965
- [11] Koay, C. G., J. D. Carew, A. L. Alexander, P. J. Basser, M. E. Meyerand 2006. Investigation of Anomalous Estimates of Tensor-Derived Quantities in Diffusion Tensor Imaging. *Magnetic Resonance in Medicine* 55:930–936. doi: 10.1002/mrm.20832

- [12] Garyfallidis, E., M. Brett, B. Amirbekian, A. Rokem, S. van der Walt, M. Descoteaux, I. Nimmo-Smith, and Dipy Contributors. 2014. DIPY, a Library for the Analysis of Diffusion MRI Data. *Frontiers in Neuroinformatics* 8(8):1-17. doi: 10.3389/fninf.2014.00008
- [13] Jenkinson, M., C.F. Beckmann, T.E. Behrens, M.W. Woolrich, S.M. Smith. 2012. FSL. *NeuroImage* 62:782-90. doi: 10.1016/j.neuroimage.2011.09.015
- [14] Jesper, L., R. Andersson, and S.N. Sotiropoulos. 2016. An Integrated Approach to Correction for Off-Resonance Effects and Subject Movement in Diffusion MR Imaging. *NeuroImage* 125:1063-1078. doi:10.1016/j.neuroimage.2015.10.019
- [15] Fan, Q., T. Witzel, A. Nummenmaa, K.R. Van Dijk, J.D. Van Horn, M.K. Drews, L.H. Somerville, M.A. Sheridan, R.M. Santillana, J. Snyder, T. Hedden, E.E. Shaw, M.O. Hollinshead, V. Renvall, R. Zanzonico, B. Keil, S. Cauley, J.R. Polimeni, D. Tisdall, R.L. Buckner, V.J. Wedeen, L.L. Wald, A.W. Toga, and B.R. Rosen 2016. MGH-USC Human Connectome Project Datasets with Ultra-High B-Value Diffusion MRI. *Neuroimage* 124(Pt B):1108-1114. doi: 10.1016/j.neuroimage.2015.08.075