



HAL
open science

Inferring the Localization of White-Matter Tracts using Diffusion Driven Label Fusion

Guillermo Gallardo, Gaston Zanitti, Mat Higger, Sylvain Bouix, Demian
Wassermann

► **To cite this version:**

Guillermo Gallardo, Gaston Zanitti, Mat Higger, Sylvain Bouix, Demian Wassermann. Inferring the Localization of White-Matter Tracts using Diffusion Driven Label Fusion. 2021. hal-03360544v1

HAL Id: hal-03360544

<https://hal.science/hal-03360544v1>

Preprint submitted on 30 Sep 2021 (v1), last revised 1 Oct 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Inferring the Localization of White-Matter Tracts using Diffusion Driven Label Fusion

G Gallardo, G Zanitti, M Higger, S Bouix, D Wassermann

Abstract

Inferring which pathways are affected by a brain lesion is key for both pre and post-treatment planning. However, many disruptive lesions cause changes in the tissue that interrupt tractography algorithms. In such cases, aggregating information from healthy subjects can provide a solution to inferring the affected pathways. In this paper, we introduce a novel label fusion technique that leverages diffusion information to locate brain pathways. Through simulations and experiments in publicly available data we show that our method is able to correctly reconstruct brain pathways, even if they are affected by a focal lesion.

1 Introduction

Pathologies such as infiltrative tumors or focal lesions disrupt the structure of white matter, resulting in cognitive deficits. Depending on the type and severity of the pathology, fiber bundles can be displaced, infiltrated or directly interrupted [1, 2, 3]. Inferring which pathways are closely located to the lesion or being directly affected by it is key for palliative care, as well as pre and post-treatment planning. With this knowledge, neurologists and neurosurgeons can decide if a lesion should be treated more aggressively or conservatively [2, 4]. However, the most common method to reconstruct brain pathways, tractography, fails in the presence of brain lesions [5]. Therefore, having a non-tracking method to locate brain pathways which is robust to lesions is key to successfully treat patients with disruptive brain lesions.

In the presence of a disruptive brain pathology, the directionality commonly present in a Diffusion Weighted Image (DWI) can be severely reduced[5]. This is directly reflected as a decrease in the fractional anisotropy (FA) in the affected regions. Since tractography algorithms rely on directionality information in order to reconstruct tracts, areas with significantly decreased FA will hamper or completely stop the tracking. This impedes reconstructing tracts in the brain, and therefore detecting those being affected by the lesion [1, 5]. In such cases where tracking is not feasible, one possible solution is to rely on information from healthy subjects. Assuming we identified major bundles in a group of healthy subjects, we can register them to our patient's brain and aggregate them to infer the affected pathways.

Label fusion is a family of techniques to infer a region's location in a target subject, based on its characteristics in a group of control subjects[6]. One well known label fusion technique is

Majority Voting [7]. Given a voxel on a brain image, each template subject "votes" for a label. The resulting label for the voxel will be that with the most votes. Majority Voting is simple to implement and has been shown to generate accurate brain segmentations even when using few template subjects [6]. Majority Voting weights the vote of each template subject equally, despite how (dis)similar they are from the target subject. However, among the templates there will be subjects more accurate in predicting a region's location, and those less accurate [8]. Because of this, it has been proposed to weigh the vote of each template based on some similarity function with the target [9]. The underlying intuition is that the label choice should be driven by those subjects who resemble the most to the one being labeled. The practical advantages of various strategies based on this idea have been demonstrated by Artaechevarria et al. [10].

In this work, we introduce the first label fusion technique that leverages diffusion data, as represented in DWI, to weights the vote of each template subject. Since the presence of a path constrains the diffusion of water particles, adding diffusion information can help to better delineate fiber bundles [11]. Particularly, we have subjects voting for the presence of a specific tract in a region. Our technique then weighs the votes based on how much the voted pathway is supported by the target's diffusion data (Fig. 1). In this way, we give higher weights to the voted pathways which directionality better agrees with the target's diffusion.

We validate our technique through simulations and experiments in publicly available data. First, we validate that our technique correctly weights votes using synthetic data. For this we generate three types of synthetic DWIs (phantoms) representing: a single fiber phantom, a 90 degrees fiber crossing phantom, and a phantom with no fibers. We generate tracts from the DWIs and test how our technique weighs them and their planar rotations on each phantom. We show that our technique assigns higher weights to the votes for tracts aligned with the DWI's diffusion. After, we randomly select 50 subjects from the Human Connectome Project (HCP) and reconstruct 6 major brain pathways on each of them. We use the reconstructed tracts as ground truth to benchmark the inferences made by our technique and Majority Voting on a leave-one-out cross-validation experiment. Our results show that our proposed technique has a lower sensitivity than Majority Voting, but a higher specificity. Meaning, our technique obtains less false positives at the cost of obtaining more true negatives. Finally, we simulate focal lesions with different degrees of severity in the white matter of one HCP subject. Particularly, we target a spherical region where the Superior Longitudinal Fasciculus passes by. We do so by iteratively decreasing the FA of each voxel in the spherical region, until obtaining isotropic diffusion. We show that, while a tractography-based reconstruction technique fails to reconstruct, our technique is able to identify the affected tract. Furthermore, our technique is always able to correctly label the voxels surrounding the lesion, while labeling less voxels within the lesion as it worsen. This helps not only to detect the affected tract, but also to know how much it is affected.

2 Methods

2.1 Majority Voting.

Let $labels = \{l_i\}$ be the set of labels representing tracts in the brain. Let $L_s, s \in S$ represent the labeling of a set of template subjects S , where each $L_s \in labels^{v_x \times v_y \times v_z}$ is a 3D volume with dimension (v_x, v_y, v_z) representing the labeling of template subject s . Majority Voting [8] infers the label of each voxel x in a target subject by computing:

$$L^*(x) = \operatorname{argmax}_{l \in labels} \sum_{s \in S} p(L(x) = l | L_s(x)),$$

where

$$p(L(x) = l | L_s(x)) = \begin{cases} 1, & \text{if } L_s(x) = l \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

In this case, each subject votes for a label, and the label with the most votes is assigned to the target voxel. It is important to notice that no information from the target subject is being used to infer the label $L^*(x)$.

2.2 Diffusion Based Voting

Majority Voting (Eq. 1) decides the label of a target voxel based on the “votes” of template subjects, without using any information from the target subject. However, when inferring the location of white matter pathways, we can profit from the fact that water molecules tend to diffuse in the direction of the pathways. This means that we can use diffusion information to weigh the voting process and help locate the bundles. In particular, votes for tracts aligned with the diffusion of the target subject should get higher weights.

One way to characterize the directionality on a diffusion weighted image (DWI) is by means of fiber orientation distribution functions (fODFs) [12]. Now, we will first explain how to estimate a fODF from both DWI and tracts. Then, we will show how to leverage them in order to compute weights for each vote.

2.2.1 Fiber Orientation Density Function from dMRI Data.

By fitting the diffusion information into a Constrained Spherical Deconvolution (CSD) model, it is possible to estimate a fiber orientation density function [13] (fODF). The fiber ODF $F_x(\theta, \phi)$ represents the estimated fraction of fibers within the voxel x that are aligned along the direction (θ, ϕ) , expressed in spherical coordinates.

2.2.2 Fiber Orientation Density Function from Tractography.

A tract can be described as a set of streamlines, where a streamline is a discretized 3-dimensional curve. Assuming that a streamline doesn't have sharp turns within a voxel, we can estimate its directionality by looking at the entry and exit points (Fig. 1 A). Repeating this for each streamline on a tract, we obtain a set of directional vectors, representing the directionality of the tract within the voxel. As with the diffusion data, we can once more use the CSD model to estimate an fODF representing the directionality of a tract in the voxel.

2.2.3 Label Fusion Weighted by Diffusion

Majority Voting (Eq. 1) decides the label of a voxel based on how many subjects "vote" for it. Given that we are inferring brain pathways, we want to introduce a weight that denotes how much the voted tract resembles the target's diffusion data:

$$L^*(x) = \operatorname{argmax}_{l \in \text{labels}} \sum_{s \in S} p(L(x) = l | L_s(x)) p(D(x) | D_{sl}(x)). \quad (2)$$

In our segmentation scheme, the term $p(L(x) = l | L_s(x))$ is modeled as in the voting scheme (eq. 1). Our second term, $p(D(x) | D_{sl}(x))$ express the probability of seeing the diffusion of our target subject, $D(x)$, on voxel x , given the diffusion of subject s generated by tract l on the same voxel, $D_{sl}(x)$. Since registering DWIs is a highly time consuming task [14], we want to avoid it. Instead, we can register tracts, for which efficient algorithms exist, and use them as an estimator of the diffusion of each template subject. Knowing that water particles in the brain diffuse along tracts, we can estimate $D_{sl}(x)$ by computing the fODF of the registered tract l as explained in 2.2.2 . Simultaneously, we can characterize $D(x)$ with the fODF computed from the DWI of the target subject as explained in 2.2.3. In order to reflect how much the fODF of our target subject's diffusion resembles the fODF of the voted tract on a voxel, we model $p(D(x) | D_{sl}(x))$ as:

$$p(D(x) | D_{sl}(x)) = \begin{cases} \langle F(x), F_{sl}(x) \rangle, & \text{if } L_s(x) = l, \\ & \text{and } l \neq 0 \\ \langle F(x), U \rangle, & \text{if } L_s(x) = 0 \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where $F(x)$ is the fiber ODF on voxel x estimated from the target's DWI, and is normalized such that $\langle F(x), F(x) \rangle = 1$. $F_{sl}(x)$ is the fiber ODF of the tract l registered from subject s and normalized as $F(x)$. U is a uniformly distributed fiber ODF, this is the diffusion assumed for either the label 'no-tract', representing the background or a gray matter structure. By computing the inner product between normalized ODFs, we can estimate how much they look alike. In this way, we weight the vote of each subject accounting for the white-matter structure of both the voting and target subjects.

3 Experiments and Results

In section 2 we presented how to add diffusion information to Majority Voting (Eq. 2). This allows us to weigh the vote for a tract by how much the diffusion of the target supports it. Now, we validate our technique using both synthetic data and subjects from the Human Connectome Project (HCP). We start by assessing our techniques correctly weighs votes using DWI phantoms. Then, we proceed to infer the location of white-matter pathways in subjects of the HCP, and compare them with their ground-truth as reconstructed with the white-matter query language (WMQL)[15]. Finally, we simulate a lesions in the white matter and test how our method behaves in its presence.

3.1 Data and Preprocessing

We created three types of diffusion weighted image phantoms using Phantomas [16]. The first phantom possess only one tract, traveling from one side to the other of the image horizontally (Fig. 2 A). The second possess two crossing tracts, forming a 90 degrees angle between them (Fig. 2 B). The last has no fibers on it, representing isotropic diffusion (Fig. 2 C). From each phantom we generated 30 diffusion weighted images (DWIs). All the DWIs were generated using a signal to noise ratio (SNR) of 20, and a resolution of $1mm$ per voxel. The final images are 3-dimensional matrices with 10 voxels in each dimension. Having such small images allows us to test our label fusion on a controlled environment. We estimated fiber orientation distributions (FODs) in each voxels by means of constrained spherical deconvolution (CSD)[17] using Dipy [18]. The fODFs were discretized on a sphere with $n = 100$ vertices. A uniform fODF, U , was created by assigning to each vertex of the sphere the value $\sqrt{(n)}/n$, making $\langle U, U \rangle = 1$.

To test our technique in more realistic scenarios, we randomly selected 50 subjects from the HCP500 dataset from the Human Connectome Project. For each subject, we first estimated FODs by means of CSD from they diffusion data. Then, we leveraged the FODs to perform whole-brain probabilistic tractography. For the tractography we used each voxel in the white-matter as a seed, and simulated 8 particles per seed [18]. Finally, we extracted four tracts from the left hemisphere using the implementation of the white-matter query language (WMQL) [15]. For each subject we computed non linear registrations to the rest using as reference their T1w images [19]. Using the resulting warp transformations, we registered the tracts between every pair of subjects.

3.2 Assessing the Correctness of Voting Weights in Synthetic Data

In order to study how the tract’s directionality influences its vote weight, we started by reconstructing the tract present in the first phantom (Fig. 2 A). For this, we took one of the 30 generated DWIs, and computed 1000 streamlines by means of probabilistic tracking from the voxels through which the tract passes. If our technique behaves correctly, the tract we reconstructed should obtain a high vote weight in the DWIs derived from phantom A. At the same time, any change in its directionality, should decrease the received weight. To assess this, we computed weights in 30 DWIs for the reconstructed tract, and for planar rotations of it around the central voxel. Figure 3 A shows the obtained weights on the first phantom. Effectively,

the weight starts to rapidly decrease as the angle increments and the directionality of the tract moves away from that of the diffusion.

Figure 3 B shows the weights obtained when computing weights of the reconstructed tract in 30 DWIs derived from the second phantom (Fig. 2 B), which have a fiber crossing in the central voxel. In this case, the weight is higher when the reconstructed tract aligns with one of the crossing fibers (at 0 degrees or 90 degrees), while rapidly decaying in between them. Finally, figure 3 C shows the weights when using 30 DWIs with isotropic diffusion (Fig. 2 C). In this case, the weight is always low, driven by the discrepancy between the directionality of the tract and the free diffusion present in the DWIs.

To assess that the proposed model is not over-weighting tracts, we also computed the weight that a “non-tract” label would receive in each of the phantoms. This is, when the user votes that there is no tract present in the voxel. As explained in section 2, equation 3, when a subject is voting for a non-tract label, a uniform fiber ODF is compared against the diffusion fODF. Figure 2 A, B, and C show the weight obtained in the central voxel when a subject is voting for the label “non-tract”. In figure 2 A, we can see that the weight of “non-tract” is low, specially when compared with the high weight of the correctly aligned tracts (low angle rotations). This is driven by the highly directional underlying diffusion data of the phantom. Figure 2 B shows that the weight of a “non-tract” vote is similar to that of an aligned tract. Finally, figure 2 C shows always a higher weight for the “non-tract” than for any tract, consistent with the isotropic diffusion of our third phantom.

3.2.1 Inferring Tracts in Human Connectom Project Subjects

To validate our technique in a more realistic but still controlled scenario, we used our technique to infer the location of specific white-matter tracts in 50 HCP subjects. We selected the following tracts to work with: Inferior Fronto-Occipital Fascicle (IFOF), Corticospinal Tract (CST), Inferior Longitudinal Fasciculus (ILF), and Superior Longitudinal Fasciculus (SLF I, II, and III). These 6 tracts provide a fair diversity of directionality, shape, and position in the brain.

For each tract we performed a leave-one-out cross-validation experiment. At each step, we inferred the tract of one subject from the registered tracts of the others using both Majority Voting and our technique. Then, we quantified the performance of both techniques using as ground truth the target’s bundle. In particular, we computed their confusion matrix. A confusion matrix is a matrix $M \in \mathbb{R}^{labels \times labels}$, where each entry M_{ij} represents the number of times the label in the ground truth was i and the technique labeled j . In this experiment, since we are only inferring one tract, we have two labels: tract and non-tract. We computed the sensitivity, and specificity of each confusion matrix [20]. Sensitivity measures the proportion of voxels in the ground-truth tract that were “discovered”. Specificity measures the proportion of voxels that were correctly labeled, over all the labeled voxels. Table 1 shows the results obtained for each technique and tract. In all of the tracts, our technique shows on average a 5% a lower sensitivity than Majority Voting. This means that we label a smaller portion of the ground-truth bundle. On the other hand, our diffusion weighted label-fusion achieves on average a 13% higher specificity. Therefore, our technique is discovering less voxels, but those

which are labeled can be trusted more.

3.2.2 Inferring Tracts in the Presence of Simulated Lesions

To test how our technique behaves on an injured brain, we simulated lesions at different degrees of severity in the white matter of one of our subjects. Given that some brain lesions directly affect Fractional Anisotropy (FA) [1, 2], we simulated lesions by adding isotropic signal to a set of voxels, therefore lowering their FA. We targeted the SLF bundle, in order to compare how the labeling changes with lesion of different degrees. We did so by selecting a spherical region of $4mm$ where the SLF passes, and mixing the diffusion signal there with signal from the ventricles. Since the ventricles are regions filled with cerebrospinal fluid (CSF), their diffusion is approximately isotropic. In particular, for each voxel x in the affected region, we chose a voxel v in the ventricle and mixed their DWI signals, $S(\cdot)$, as follows:

$$S(x) = S(x)(1 - \alpha) + S(v)\alpha, \alpha \in [0, 1], \quad (4)$$

where α manages the severity of the lesion. In this case, $\alpha = 0$ represents healthy tissue, and $\alpha = 1$ represents a total disruption of the white-matter, resulting in pure isotropic diffusion.

We created four datasets by setting α to 0.25, 0.5, 0.75 and 1. On each dataset we performed a whole-brain probabilistic tractography, and filtered the resulting streamlines using WMQL to reconstruct the SLF. We found that for the values of $\alpha = 0.75$ and $\alpha = 1$ WMQL failed to recover the SLF, since it was interrupted by the simulated lesion. Then, we used both majority voting and our technique to aggregate the SLF from the rest of the subjects. Since Majority Voting for all datasets. Moreover, both Majority Voting and our technique obtained a high overlap in the result for $\alpha = 0.25$. However, our technique labeled less voxels within the lesion as the alpha value increased (making FA to decrease). This is a good behaviour, since by lowering the FA we make the diffusion more isotropic, losing the underlying tract. In particular, for $\alpha = 1$, our technique results in two disconnected sets of voxels being labeled as SLF. This is correct, because for $\alpha = 1$ the diffusion is completely isotropic, meaning that the tract was interrupted. This behaviour allows not only to know which tract was affected, but also at which degree it is affected, informing if the tract is interrupted or not.

4 Discussion

In this work we presented a label-fusion technique that leverages diffusion information to better infer the location of brain pathways. Our technique allows to correctly locate white-matter pathways even in the presence of a lesion that disrupts tractography algorithms. Furthermore, being based on a label-fusion algorithm, our technique can achieve accurate segmentations even when the inference is made from few subjects [6].

Our technique adds diffusion information in the process of label-fusion. Given that fiber bundles constrains the diffusion of water particles in the brain, our technique uses diffusion

information to improve Majority Voting [7]. More specifically, we weight each vote based on how the voted pathway is supported by the target’s diffusion data. In this way, voted pathways that better resemble the white matter of the target subject obtain a higher weight. The weights come from comparing how much the diffusion fODF of our target subject’s resembles the fODF of the voted tract on a voxel. In this way, we can compare the white-matter structure of our target subject with that of the voting subject, without having to register DWIs. Adding diffusion weights to Majority Voting, allowed us to profit from its robustness while improving the labeling of white-matter bundles, as shown by our results in synthetic data and subjects from the Human Connectome Project.

4.1 Our Technique Creates Weights Consistent With the Underlying Diffusion Data.

To assess that our technique correctly weighed votes, we created three different phantoms, and derived DWIs from them. The phantoms had: (A) a single bundle running straight, (B) a 90 degrees crossing bundles, and (C) no bundles. We reconstructed the tract in the phantom A, and computed the weights obtained by the tract and planar rotations of it in the 3 phantoms. Our results show that our technique assigns high weights when the tract aligns with the bundles represented in the underlying diffusion data. Furthermore, when the directions differ by more than 10 degrees, the weight starts to drop rapidly. Our results also show that in the phantom C, where there is no bundle, the ‘no-tract’ vote receives a higher weight than any rotation. These results show that our technique is able to correctly weight each label based on diffusion directionality.

4.2 Our Technique Shows Lower Sensitivity but Higher Specificity than Majority Voting.

To test our technique in realistic data, we registered tracts between different subjects of the HCP. Using these registered tracts, we inferred the position of individual tracts in each subject. Table 1 shows that in each inferred tract, our technique achieved on average a 13% higher specificity at the trade off of obtaining a 5% lower sensitivity. Hence, our techniques presents less false positives at the cost of more false negatives, making it more conservative than Majority Voting. Furthermore, our specificity is always higher than 0.7. This means that our technique was able to discover less voxels belonging to the tracts, but at least 70% of those labeled are correct.

4.3 Our Technique Allows to Infer Which Tract is Affected by a Focal Lesion, and How Affected the Area Is

We further characterized the behaviour of our technique by simulating non-deforming lesions in the white-matter of a HCP subject. Particularly, we defined a spherical region on the path of the Superior Longitudinal Fasciculus, and increased its Fractional Anisotropy until achieving isotropic diffusion. In doing so, we effectively simulated a focal lesions ranging from an infiltration to the interruption of a tract. Under these common scenarios in which the diffusion

data is affected, reconstruction techniques such as WMQL are not able to recover the tract. This is because, as the diffusion data loses directionality, the tracking becomes more and more hampered. Indeed, in our experiment WMQL was not able to dissect the affected SLF, since the process of tractography was stopped by the lesion. On the other hand, both majority voting and our technique were able to correctly label the affected tract. However, as shown by our results (Figure 4), only our technique was able to label less voxels as the tract became more interrupted. This is the expected behaviour in the presence of a lesion, since the more damaged an area is, the less it is expected for the tract to be present. When the tract is completely interrupted, our technique is still able to label the voxels surrounding the lesions, thus enabling to detect the affected tract. Therefore, our technique not only allows to detect the affected tracts, but also gives information on how altered the tract is.

5 Conclusions

We presented a label fusion technique that leverages diffusion data to infer the localization of white-matter tracts. The results show that our technique is well suited to create subject specific segmentation of the brain. Our method could be used in specificity medicine to interrogate diffuse pathologies like TBI, even in the presence of small lesions (as shown in fig. 4). Our technique could be used in the study of infiltrative brain tumors, where registration is less of a problem, and in the future, could be extended to larger tumors, due to the growing advances in registration based on tumor growth models.

Table 1: Sensitivity and specificity of our proposed method (Weighted) and a non-weighted voted (Majority) when inferring single bundles from 9 subjects. The inferred bundles are: Superior Longitudinal Fasciculus (SLF) I, II and III, Inferior Longitudinal Fasciculus (ILF), Cortico Spinal Tract (CST), and Inferior Occipito-Frontal Fascicle (IOFF).

	Sensitivity		Specificity	
	Weighted	Majority	Weighted	Majority
IOFF	0.10 (0.04)	0.13 (0.04)	0.71 (0.15)	0.55 (0.15)
CST	0.38 (0.03)	0.42 (0.03)	0.80 (0.04)	0.70 (0.04)
ILF	0.22 (0.01)	0.27 (0.02)	0.77 (0.04)	0.63 (0.04)
SLF I	0.19 (0.01)	0.27 (0.01)	0.80 (0.04)	0.66 (0.05)
SLF II	0.23 (0.01)	0.31 (0.01)	0.82 (0.04)	0.68 (0.04)
SLF III	0.11 (0.01)	0.17 (0.02)	0.71 (0.04)	0.58 (0.05)

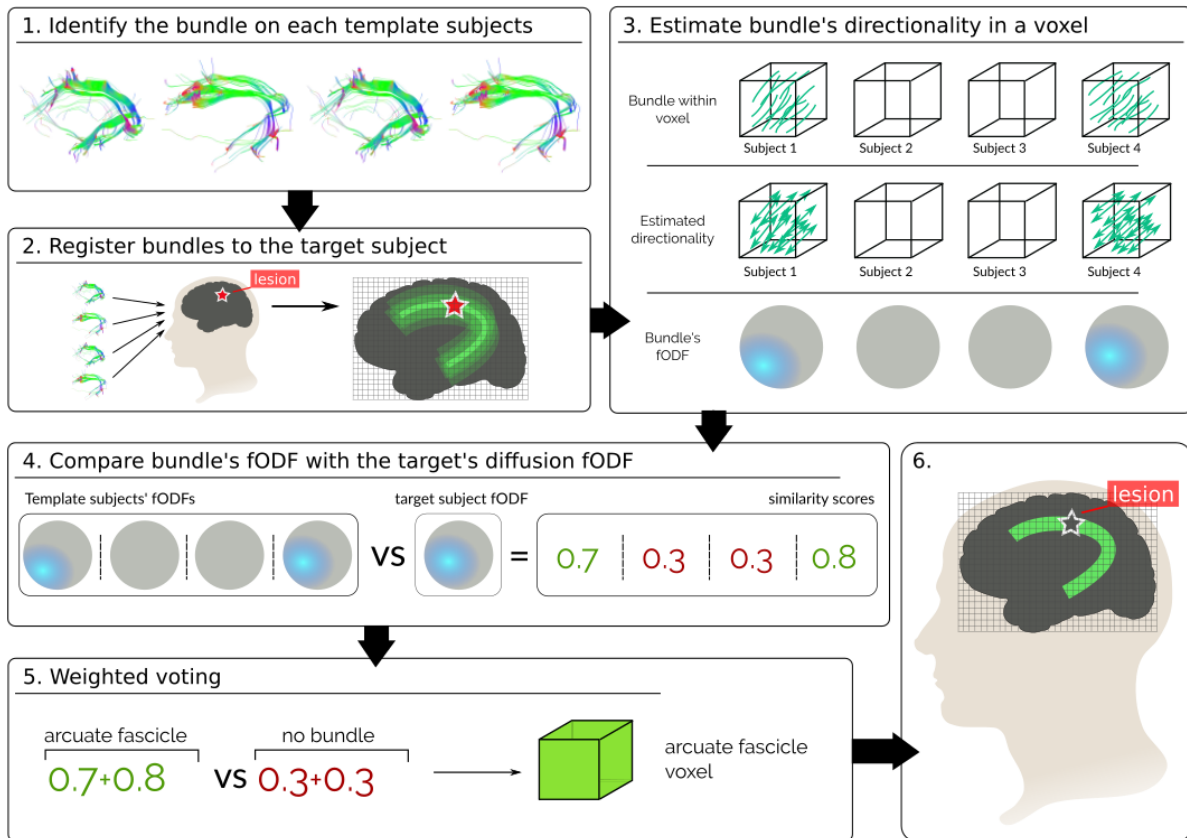


Figure 1: We use a label fusion technique to infer the location of white matter bundles by aggregating information of healthy subjects. After registering the tracts, each subject "votes" for either a tract or a non-tract structure (gray matter or background). Our technique adds diffusion based weights to each one of these votes. The weights are computed based on the similarity between the fiber Orientation Distribution Function of the structure being voted and the DWI of the target subject.




Figure 2: In order to test our technique, we created three types of synthetic DWIs, known as phantoms. (A) A phantom with only one tract in the white matter. (B) A phantom with two fibers crossing. (C) A phantom with no tracts, representing isotropic diffusion.

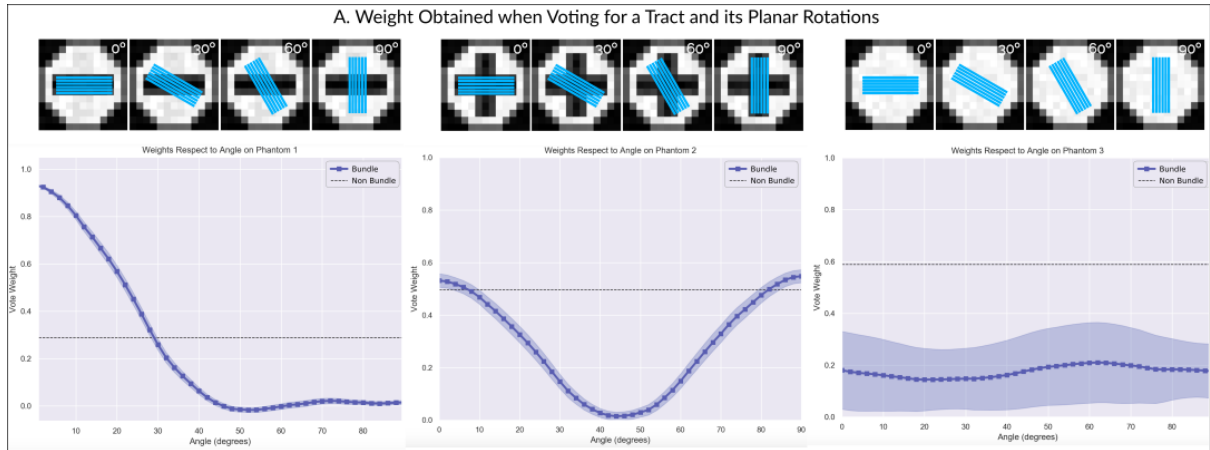


Figure 3: In order to study how a tract’s directionality influences its vote weight, we estimated the fiber bundle by means of tractography in the Phantom A. Then, we computed the weights obtained by the tract and planar rotations of it in: (A) Phantom A. (B) Phantom B. (C) Phantom C. The figures show that our technique gives the highest weights to structures that are aligned with the underlying diffusion.




Figure 4: We tested how our technique behaves in the presence of simulated lesions. Lesions were simulated in a specific region (red circle) by following eq. 4 to lower the FA of the region. The SLF passes through such a region. The figures show the result of applying our technique in order to determine the SLF at different values of the parameter α in eq. 4: (A) $\alpha = 0.2$, (B) $\alpha = 0.5$, (C) $\alpha = 0.75$, and (D) $\alpha = 1$. Results show that while the value of α increases, the amount of voxels labeled within the affected region decreases.

References

- [1] T. Schonberg, P. Pianka, T. Hendler, O. Pasternak, and Y. Assaf, “Characterization of displaced white matter by brain tumors using combined DTI and fMRI,” *Neuroimage*, vol. 30, pp. 1100–1111, 2006.
- [2] T. A. G. M. Huisman, “Tumor-like lesions of the brain,” *Cancer Imaging*, pp. 10–13, 2009.
- [3] Y. I. Won, C. K. Chung, C. H. Kim, C.-k. Park, B.-b. Koo, J.-m. Lee, and H.-w. Jung, “White Matter Change Revealed by Diffusion Tensor Imaging in Gliomas,” *Brain Tumor Res Treat*, vol. 4, no. 2, pp. 100–106, 2016.
- [4] M. J. McGirt, K. L. Chaichana, M. Gathinji, F. J. Attenello, K. Than, A. Olivi, J. D. Weingart, H. Brem, and A. redo Quiñones-Hinojosa, “Independent association of extent of resection with survival in patients with malignant brain astrocytoma,” *J. Neurosurg.*, vol. 110, pp. 156–162, jan 2009.
- [5] B. J. Jellison, A. S. Field, J. Medow, M. Lazar, M. S. Salamat, and A. L. Alexander, “Diffusion Tensor Imaging of Cerebral White Matter: A Pictorial Review of Physics, Fiber Tract Anatomy, and Tumor Imaging Patterns,” 2004.
- [6] A. J. Asman and B. A. Landman, “Non-local statistical label fusion for multi-atlas segmentation,” *Med. Image Anal.*, vol. 17, no. 2, pp. 194–208, 2013.
- [7] L. Xu, A. Krzyżak, and C. Y. Suen, “Methods of Combining Multiple Classifiers and Their Applications to Handwriting Recognition,” *IEEE Trans. Syst. Man Cybern.*, vol. 22, no. 3, pp. 418–435, 1992.
- [8] T. Rohlfing, D. B. Russakoff, and C. R. Maurer, “Performance-based classifier combination in atlas-based image segmentation using expectation-maximization parameter estimation,” *IEEE Trans. Med. Imaging*, vol. 23, no. 8, pp. 983–994, 2004.
- [9] M. R. Sabuncu, B. T. T. Yeo, K. V. Leemput, B. Fischl, and P. Golland, “A Generative Model for Image Segmentation Based on Label Fusion,” *IEEE Trans. Med. Imaging*, vol. 29, no. 10, pp. 1714–1729, 2010.
- [10] X. Artaechevarria, A. Muñoz-Barrutia, and C. Ortiz-de Solórzano, “Combination strategies in multi-atlas image segmentation: Application to brain MR data,” *IEEE Trans. Med. Imaging*, vol. 28, no. 8, pp. 1266–1277, 2009.
- [11] G. Girard, A. Daducci, L. Petit, J. P. Thiran, K. Whittingstall, R. Deriche, D. Wassermann, and M. Descoteaux, “AxTract: Toward microstructure informed tractography,” *Hum. Brain Mapp.*, vol. 38, no. 11, pp. 5485–5500, 2017.
- [12] D. S. Tuch, “Q-ball imaging,” *Magn. Reson. Med.*, vol. 52, pp. 1358–1372, dec 2004.
- [13] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly, “Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution,” *Neuroimage*, vol. 23, pp. 1176–1185, nov 2004.

- [14] L. J. O’Donnell, A. Daducci, D. Wassermann, and C. Lenglet, “Advances in computational and statistical diffusion MRI,” *NMR Biomed.*, no. July, p. e3805, 2017.
- [15] D. Wassermann, N. Makris, Y. Rathi, M. Shenton, R. Kikinis, M. Kubicki, and C.-f. Westin, “The white matter query language: a novel approach for describing human white matter anatomy,” *Brain Struct. Funct.*, vol. 221, pp. 4705–4721, dec 2016.
- [16] E. Caruyer, A. Daducci, M. Descoteaux, J.-C. Houde, J.-P. Thiran, and R. Verma, “Phantomas: a flexible software library to simulate diffusion MR phantoms,” *Ismrm*, vol. 17, p. 20013, 2014.
- [17] J.-d. Tournier, F. Calamante, and A. Connelly, “Robust determination of the fibre orientation distribution in diffusion MRI : Non-negativity constrained super-resolved spherical deconvolution,” *Neuroimage*, vol. 35, pp. 1459–1472, 2007.
- [18] E. Garyfallidis, M. Brett, B. Amirbekian, A. Rokem, S. van der Walt, M. Descoteaux, and I. Nimmo-Smith, “Dipy, a library for the analysis of diffusion MRI data,” *Front. Neuroinform.*, vol. 8, p. 8, feb 2014.
- [19] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith, “Fsl,” *Neuroimage*, vol. 62, no. 2, pp. 782–790, 2012.
- [20] M. Kuhn and K. Johnson, *Applied Predictive Modeling*. Springer, 2013.