Published in final edited form as:

Med Image Anal. 2014 January ; 18(1): 161–175. doi:10.1016/j.media.2013.10.007.

Improved Segmentation of White Matter Tracts with Adaptive Riemannian Metrics

Xiang Hao, Kristen Zygmunt, Ross T. Whitaker, and P. Thomas Fletcher

School of Computing, University of Utah, Salt Lake City, UT; Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT

Abstract

We present a novel geodesic approach to segmentation of white matter tracts from diffusion tensor imaging (DTI). Compared to deterministic and stochastic tractography, geodesic approaches treat the geometry of the brain white matter as a manifold, often using the inverse tensor field as a Riemannian metric. The white matter pathways are then inferred from the resulting geodesics, which have the desirable property that they tend to follow the main eigenvectors of the tensors, yet still have the flexibility to deviate from these directions when it results in lower costs. While this makes such methods more robust to noise, the choice of Riemannian metric in these methods is ad hoc. A serious drawback of current geodesic methods is that geodesics tend to deviate from the major eigenvectors in high-curvature areas in order to achieve the shortest path. In this paper we propose a method for learning an adaptive Riemannian metric from the DTI data, where the resulting geodesics more closely follow the principal eigenvector of the diffusion tensors even in high-curvature regions. We also develop a way to automatically segment the white matter tracts based on the computed geodesics. We show the robustness of our method on simulated data with different noise levels. We also compare our method with tractography methods and geodesic approaches using other Riemannian metrics and demonstrate that the proposed method results in improved geodesics and segmentations using both synthetic and real DTI data.

Keywords

Diffusion tensor imaging; Riemannian manifold; Conformal factor; Geodesic; Front-propagation

1. Introduction

In order to study normal brain development, as well as neuropsychiatric disorders such as autism, it is crucial to understand how different functional regions of the brain are connected by white matter pathways. One approach to studying white matter in vivo is diffusion tensor imaging (DTI), a magnetic resonance imaging (MRI) modality that measures the diffusion of water in tissue. These diffusion measurements provide a means for inferring the microstructural properties of the white matter and analyzing fiber tracts. Three approaches to DTI analysis are: whole-brain connectivity analysis; localizing white matter regions by registration to an atlas; and segmenting individual white matter tracts from specified regions of interest (ROI). In whole-brain connectivity analysis, the goal is to explore the connectivity among many anatomical regions over the whole brain, typically using tractography and graph statistics (Hagmann et al., 2007). In atlas-based methods, the white matter is analyzed at the voxel level (Barnea-Goraly et al., 2005) or the atlas is used to segment the white matter into several anatomical tracts (Bazin et al., 2011). In this paper, we focus on segmentation of individual white matter tracts connecting two specified ROIs.

Several works have developed segmentation methods for white matter tracts from DTI data. Zhukov et al. (2003) employ level-sets to create geometric models of brain structures. Rousson et al. (2004) extend region-based surface evolution to DTI. Lenglet et al. (2005) model DTI data as multivariate Gaussian distributions and employ a level variational approach to segment the white matter structures. Wang and Vemuri (2005) use the square root of the J-divergence as the distance of tensors in a region-based active contour model for DTI segmentation. Ziyan et al. (2006) propose a modified spectral clustering method to segment thalamic nuclei. Awate et al. (2007) use a nonparametric model to get a fuzzy segmentation of the white matter tracts. Melonakos et al. (2007b) propose a locally constrained Bayesian region growing approach based on a pre-computed anchor path inside the white matter tract. Niethammer et al. (2009) develop a segmentation framework for near-tubular white matter tracts through global statistical modeling and local reorienting of the diffusion orientation. These methods focus on segmenting the white matter tracts of interest from the tensor field and they do not compute parameterized fiber pathways connecting the two end regions of the tracts. Tractography (Mori et al., 1999b; Conturo et al., 1999; Basser et al., 2000; Koch et al., 2002; Behrens et al., 2003; Parker et al., 2003; Lazar and Alexander, 2005; Friman et al., 2006; Jones, 2008) and front-propagation (Parker et al., 2002; O'Donnell et al., 2002; Jackowski et al., 2005; Melonakos et al., 2007a; Pichon et al., 2005; Fletcher et al., 2007; Jbabdi et al., 2008; Hao et al., 2011) approaches, however, provide both a volumetric segmentation of the tract suitable for region-based analysis and a parameterization suitable for along-tract statistics (Corouge et al., 2006). In this paper, we extend the front-propagation approaches and compare our approach to previous methods for front propagation and tractography.

1.1. Properties of Tractography

Deterministic tractography (Mori et al., 1999b; Conturo et al., 1999; Basser et al., 2000) computes streamlines (sometimes called fibers) by forward integration of the principal eigenvector of the diffusion tensors from one region. One major problem with tractography is that imaging noise causes errors in the principal eigenvector direction, and these errors accumulate in the integration of the streamlines. Another disadvantage to tractography is that it has difficulty in cases where the goal is to find pathways between two regions. In this scenario, streamlines begin in one of the regions and are accepted only if they eventually pass through the desired ending region. However, several factors conspire to often result in only a small fraction of fibers being accepted. For example, accumulated errors in the streamlines can throw them off the final destination. Also, noise and partial volume effects in the tensor field can cause stopping criteria to be prematurely triggered, either by low anisotropy tensors or sudden direction changes. The Brute-Force (BF) approach proposed by Conturo et al. (1999) can increase the number of accepted fibers by initiating fiber tracking from every voxel in the brain. However, this approach still suffers from the same factors mentioned above and can often segment only the main core of the white matter tracts.

Stochastic tractography (Koch et al., 2002; Behrens et al., 2003; Parker et al., 2003; Lazar and Alexander, 2005; Friman et al., 2006; Jones, 2008) is an approach that deals with the problems arising from image noise. In these methods, large numbers of streamlines are initiated from each seed voxel and are integrated along directions determined stochastically at each point. However, this is a computationally-intensive procedure (typically requiring several hours). Also, stochastic tractography suffers from the same problems with streamlines stopping in noisy or low-anisotropy regions, leading to artificially low (or even zero) probabilities of connection. Although Barbieri et al. (2012) combine tensor clustering technique with stochastic tractography in order to improve the accuracy of the segmentation results, this method introduces more parameters and strongly depends on the quality of the connectivity map.

1.2. Properties of Front-Propagation

In the DTI literature, front-propagation approaches are one class of methods to analyze the white matter pathways. These methods infer the pathways of white matter by first evolving a level set representing the time-of-arrival of paths emanating from some starting region. Then the pathways are computed by integrating the characteristics vectors of the level set backward from any target point to the starting region (Jackowski et al., 2005). The direction and speed of this evolving front at each point is determined by some cost function derived from the diffusion tensor data. One such method, first proposed by O'Donnell et al. (2002), is to treat the inverse of the diffusion tensor as a Riemannian metric, and the paths in the propagating front as geodesics, i.e., shortest paths, under this metric. This makes intuitive sense: traveling along the large axis of the diffusion tensor results in shorter distances, while traveling in the direction of the small axes results in longer distances. Therefore, the shortest paths will tend to remain tangential to the principal eigenvector of the diffusion tensor.

Front-propagation approaches for analyzing white matter pathways are attractive for at least three reasons. First, the front-propagation algorithms are more robust to noise than both deterministic tractography and stochastic tractography. This is because front-propagation methods compute fibers by optimizing a global criterion over the whole brain, so the wavefront is not constrained to exactly follow the principal eigenvector of the tensors. Although the principal eigenvector of the tensor is the preferred direction for paths to travel, the minimal-cost paths may deviate from these directions if the deviation decreases the overall cost, and hence are less sensitive to noise or partial voluming. Second, front-propagation methods can compute a large number of fibers using a short computational time. Efficient implementations of front-propagation solvers are much faster (typically requiring several seconds) than stochastic tractography. The graphics processing unit (GPU) implementation by Jeong et al. (2007) even runs at near real-time speeds. Finally, as shown by Fletcher et al. (2007), front-propagation methods can be used to segment white matter tracts by solving the geodesic flow from two ROIs and combining the resulting cost functions. This approach has the advantage that the solution will not get stuck in regions of noisy data or low anisotropy, in contrast to tractography methods. However, it also has the disadvantage that it requires the user to predefine two ROIs at the endpoints of the white matter tract of interest. Consequently, this approach is only appropriate when the anatomy of the white matter pathway is well-known, i.e., its endpoint regions can be reasonably identified, because a white matter path will always be found. Although, if a “false positive” connection is found, this can be detected using heuristic connectivity metrics as introduced by Parker et al. (2002) and Jackowski et al. (2005).

1.3. High Curvature Tract Deviation

While front-propagation is a powerful framework for computing white matter pathways and despite the advantages that front-propagation methods have over tractography, there is one severe drawback. These geodesics have the serious deficiency that in high-curvature tracts they tend to deviate from the eigenvector directions and take straighter trajectories than is desired. That is, in high-curvature regions, the incremental cost of following the tensor field is overcome by the cost associated with the longer (more curved) path. The top image of Figure 1 is a diagram illustrating the problem. In a curved tensor field, one would typically prefer a path that follows, to whatever extent possible, the major eigenvectors of the tensors (shown in blue). The shortest path, using a Euclidean metric (i.e., ignoring the tensors), follows a straight line except at constraints (shown in red). The typical geodesic with a local, anisotropic metric (e.g., using the inverse tensors as metric) will find a compromise between these two (shown in magenta). Although the magenta geodesic is taking infinitesimally higher-cost steps than the blue curve, its overall length under the inverse-tensor metric is shorter.

Fletcher et al. (2007) have addressed this issue previously by “sharpening” the tensor, i.e., increasing the anisotropy by taking the eigenvalues to some power and renormalizing them, which increases the cost of moving in directions other than the principal eigenvector. Actually, the first front-propagation algorithm proposed by Parker et al. (2002) essentially takes this sharpening strategy to its limit, which results in a cost function that is the dot product of the level set velocity with the principal eigenvector, and Jbabdi et al. (2008) show that the geodesics more closely follow the principal eigenvectors as the anisotropy of the noiseless tensor increases. However, first of all, the sharpening is applied equally across the image, rather than taking the curvature of the tract into account, and sharpening that increases with the curvature of the tract could be more effective. Secondly, another downside of sharpening is that it changes the shape of the tensor and reduces the ability to deviate from the principal direction, thus decreasing the desired robustness to noise. Lastly, the amount of sharpening is an ad hoc parameter, and it is not clear how to set the amount of sharpening to find the best balance between robustness to noise versus faithful following of the eigenvectors.

In this paper we develop a new Riemannian metric that is a modulated version of the inverse diffusion tensor field. This metric is able to adaptively correct the geometry of geodesic curves in high-curvature regions so that they more closely follow the principal eigenvectors of the tensors. The resulting algorithm requires solving for an unknown scalar field (one example is shown in the bottom image of Figure 1), which we compute by solving a Poisson equation on the Riemannian manifold—however, it does not require any arbitrary choice of parameters. Based on the computed geodesics from the two end regions of a white matter tract, we develop an automatic segmentation framework resulting in segmentations that better delineate the white matter tracts and does not require parameters tuning or other kinds of user intervention. A preliminary version of this work has been introduced previously in Hao et al. (2011). In this paper, we provide additional details about the derivations of the geodesic equation and Poisson equation (in Section 2), discuss the numerical methods for solving the Poisson equation (in Section 3), and develop a new segmentation framework based on the computed geodesics (in Section 4). In the end, we show that our adaptive metric is sufficient to eliminate the problem with geodesics in high-curvature regions described above and illustrated in Figure 1, and we demonstrate the corrected behavior of geodesics, robustness of our adaptive metric, and the improvements of segmentations of the white matter tracts on both synthetic and real data (in Section 5).

2. Adaptive Riemannian Metrics

In this section we derive a procedure for computing geodesic flows in diffusion tensor data that resolves the major drawback of front-propagation approaches outlined above. Namely, the geodesics generated by our method more closely conform to the principal eigenvector field. Rather than directly using the inverse of the diffusion tensor as the Riemannian metric, as is typically done, we compute a spatially-varying scalar function that modulates the inverse tensor field at each point and use this as our metric. We show that this scalar field can be chosen in such a way that the resulting geodesic flows have the desired property of following the eigenvector directions. This entails solving the classical variational problem for geodesic curves, with the exception that the Riemannian metric is scaled by a positive function. In the resulting Euler-Lagrange equation, we then solve for the particular scaling function that causes geodesics to follow the desired directions. In the end, we see that the appropriate function is computed by solving a Poisson equation on the Riemannian manifold.

2.1. The Metric Modulating Function

On a Riemannian manifold, M , the geodesic between two points $p, q \in M$ is defined by the minimization of the energy functional

$$E(\gamma) = \int_0^1 \langle T(t), T(t) \rangle dt,$$

where $\gamma: [0, 1] \rightarrow M$ is a curve with fixed endpoints, $\gamma(0) = p$, $\gamma(1) = q$, $T = d\gamma/dt$, and the inner product is given by the Riemannian metric. In our case the manifold $M \subset \mathbb{R}^3$ is the image domain, and the Riemannian metric can be equated with a smoothly-varying, positive-definite matrix $g(x)$ defined at each point $x \in M$. Letting $T_x M$ denote the tangent space at a point $x \in M$, the inner product between two tangent vectors $u, v \in T_x M$ is given by $\langle u, v \rangle = u^t g(x) v$. As mentioned above, previous front-propagation approaches to DTI have used the inverse of the diffusion tensor field as a metric, i.e., $g(x) = D(x)^{-1}$ (or a sharpened or modified version), and this choice of metric leads to geodesics that bend inwards around curves. To rectify this problem, we will scale the Riemannian metric by a positive function $e^{a(x)}$, which results in the new geodesic energy functional

$$E_a(\gamma) = \int_0^1 e^{a(\gamma(t))} \langle T(t), T(t) \rangle dt. \quad (1)$$

We call the function e^a the *metric modulating function* because it scales the Riemannian metric at each point. The exponentiation of a is to ensure that this scaling factor is positive and to make the solution to the variational problem come out simpler in the end. While it is possible to envision more complicated modifications of the metric tensor, we choose to modify the metric in this fashion for three reasons. First, the shape of the diffusion tensor provides information about the relative preference in diffusion directions, and a scaling operation allows us to keep this information intact. Second, the modification in (1) is sufficient to correct for the effects of curvature. In other words, if the tensors are following a curved path, but not changing shape, the metric modulating function can be chosen in such a way that the resulting geodesics perfectly follow the principal eigenvector. We demonstrate this property empirically using a synthetic example in Section 5. Third, on a Riemannian manifold M , if there exist two Riemannian metrics g_1, g_2 satisfying $g_1 = f g_2$ for some positive function f on M , we call these two metrics conformally equivalent and the function f a conformal factor. So, our modulated Riemannian metric is a conformal transformation of inverse-tensor metric and the computed e^a can be seen as a conformal factor.

2.2. Computing the Geodesic Equation

To minimize the new geodesic energy functional given in (1), we use two tools of Riemannian geometry. The first is the affine connection $\nabla_X Y$, which is the derivative of a vector field Y in the direction of a vector field X . We will write the vector fields X, Y in terms of a coordinate system $(x^1, x^2, \dots, x^n)$; note that superscripts here are indices, not

exponentiation. We write $X = \sum a^i E_i$ and $Y = \sum b^j E_j$, where $E_i = \frac{\partial}{\partial x^i}$ are the coordinate basis vectors, and a^i and b^j are smooth coefficients functions. Then the affine connection is given by

$$\nabla_X Y = \sum_k \left(\sum_i a^i \frac{\partial b^k}{\partial x^i} + \sum_{i,j} \Gamma_{ij}^k a^i b^j \right) E_k.$$

The terms Γ_{ij}^k are the Christoffel symbols, which are defined as

$$\Gamma_{ij}^k = \frac{1}{2} \sum_{l=1}^n g^{kl} \left(\frac{\partial g_{jl}}{\partial x^i} + \frac{\partial g_{il}}{\partial x^j} - \frac{\partial g_{ij}}{\partial x^l} \right),$$

where g_{ij} denotes the entries of the Riemannian metric, g , and g^{ij} denotes the entries of the inverse metric, g^{-1} . Again, the intuition behind this affine connection is that it is like a directional derivative of vector fields. In the special case of $Y = X$, $\nabla_X X$ measures how the vector field X bends along its integral curves.

The second tool that we employ is the Riemannian gradient of a smooth function f , which we denote $\text{grad } f$. The gradient of a function on a Riemannian manifold looks like the standard Euclidean gradient, except with a multiplication by the inverse of the metric, i.e.,

$$\text{grad } f = g^{-1} \left(\frac{\partial f}{\partial x^1}, \frac{\partial f}{\partial x^2}, \dots, \frac{\partial f}{\partial x^n} \right).$$

The gradient is defined in this way so that the inner product with a unit vector u results in the usual directional derivative, $\nabla_u f = \langle \text{grad } f, u \rangle$.

Using the affine connection and Riemannian gradient, we take the variation of the energy (1). Let W be a vector field defined along the curve γ that represents an arbitrary perturbation of γ , keeping the endpoints fixed, i.e., $W(0) = W(1) = 0$. Notice that W and T are partial derivatives of the variation of γ , and therefore they commute, i.e., $\nabla_W T = \nabla_T W$. To simplify the notation, we will suppress the parameter t in most of the following equations. Then the variational of the energy functional is

$$\begin{aligned} \nabla_W E_\alpha(\gamma) &= \nabla_W \int_0^1 e^\alpha \langle T, T \rangle dt \\ &= \int_0^1 \nabla_W e^\alpha \langle T, T \rangle dt \\ &= \int_0^1 \nabla_W e^\alpha \cdot \langle T, T \rangle + e^\alpha \nabla_W \langle T, T \rangle dt \\ &= \int_0^1 \langle W, \text{grad } e^\alpha \rangle \cdot \langle T, T \rangle + 2e^\alpha \langle \nabla_W T, T \rangle dt \\ &= \int_0^1 \langle W, \text{grad } e^\alpha \rangle \cdot \|T\|^2 + 2 \langle \nabla_W T, e^\alpha T \rangle dt \\ &= \int_0^1 \langle W, \text{grad } e^\alpha \rangle \cdot \|T\|^2 + 2 \langle \nabla_T W, e^\alpha T \rangle dt \\ &= \int_0^1 \left\langle W, e^\alpha \|T\|^2 \text{grad } \alpha \right\rangle - 2 \langle W, \nabla_T (e^\alpha T) \rangle dt \\ &= \int_0^1 \left\langle W, e^\alpha \|T\|^2 \text{grad } \alpha - 2e^\alpha d\alpha(T) \cdot T - 2e^\alpha \nabla_T T \right\rangle dt. \end{aligned}$$

Now, setting this last line to zero and dividing through by e^α results in the geodesic equation

$$\text{grad } \alpha \cdot \|T\|^2 = 2 \nabla_T T + 2 d\alpha(T) \cdot T. \quad (2)$$

If we assume, without loss of generality, that geodesics have unit-speed parameterization, i.e., $\|T\| = 1$, then $\nabla_T T$ will be normal to T . Now, assuming this parameterization and taking the inner product with T on both sides of (2), we obtain

$$\langle \text{grad } \alpha, T \rangle = 2 d\alpha(T) = 2 \langle \text{grad } \alpha, T \rangle.$$

This can hold only if the tangential component, $\langle \text{grad } \alpha, T \rangle = 0$. Therefore, the last term in (2) must vanish, and we get the final, simplified geodesic equation

$$\text{grad} \alpha = 2 \nabla_T T. \quad (3)$$

2.3. Computing the Metric Modulating Function

Now that we have the geodesic equation for the modulated Riemannian metric, we introduce the property that we would like to enforce: that the tangent vectors, T , follow the unit principal eigenvector directions, V . Satisfying this property directly would result in the equation $\text{grad} \alpha = 2 \nabla_V V$, which we would need to solve for α . However, given an arbitrary unit vector field V , there may not exist such a function with the desired gradient field.

Instead we minimize the squared error between the two vector fields, i.e., we minimize the functional

$$F(\alpha) = \int_M \|\text{grad} \alpha - 2 \nabla_V V\|^2 dx. \quad (4)$$

As before, the norm here is given by the Riemannian metric. The Euler-Lagrange solution to this problem is derived in the following derivation similarly to the classical Poisson equation,

$$\begin{aligned} \left. \frac{dF(\alpha + \epsilon h)}{d\epsilon} \right|_{\epsilon=0} &= \frac{d}{d\epsilon} \int_M \|\text{grad}(\alpha + \epsilon h) - 2 \nabla_V V\|^2 dx \Big|_{\epsilon=0} \\ &= \frac{d}{d\epsilon} \int_M \langle \text{grad}(\alpha + \epsilon h) - 2 \nabla_V V, \text{grad}(\alpha + \epsilon h) - 2 \nabla_V V \rangle dx \Big|_{\epsilon=0} \\ &= 2 \int_M \left\langle \frac{d}{d\epsilon} \text{grad}(\alpha + \epsilon h), \text{grad}(\alpha + \epsilon h) - 2 \nabla_V V \right\rangle dx \Big|_{\epsilon=0} \\ &= 2 \int_M \langle \text{grad} h, \text{grad} \alpha - 2 \nabla_V V \rangle dx \\ &= -2 \int_M \langle h, \text{div}(\text{grad} \alpha) - 2 \text{div}(\nabla_V V) \rangle dx \\ &\quad + \int_M \text{div}(h \cdot (\text{grad} \alpha - 2 \nabla_V V)) dx \\ &= -2 \int_M \langle h, \Delta \alpha - 2 \text{div}(\nabla_V V) \rangle dx \\ &\quad + \int_{\partial M} h \cdot \langle \text{grad} \alpha - 2 \nabla_V V, \vec{n} \rangle dx. \end{aligned}$$

The div operator used above is the Riemannian divergence, and the divergence of X on M is defined in coordinates as

$$\text{div}(X) = \frac{1}{\sqrt{|g|}} \sum_i \frac{\partial}{\partial x^i} \left(\sqrt{|g|} a^i \right),$$

where $|g|$ is the determinant of the Riemannian metric, which represents the appropriate volume element. Finally, the equation of the metric modulating function that minimizes (4) is given by

$$\Delta \alpha = 2 \text{div}(\nabla_V V), \quad (5)$$

where $\Delta \alpha = \text{div}(\text{grad} \alpha)$ is the Laplace-Beltrami operator on M , defined as

$$\Delta(\alpha) = \frac{1}{\sqrt{|g|}} \sum_i \frac{\partial}{\partial x^i} \left(\sqrt{|g|} \sum_j g^{ij} \frac{\partial \alpha}{\partial x^j} \right)$$

in coordinates. From the boundary integral in the last step of the above derivation, the appropriate boundary conditions for this problem are the Neumann conditions,

$$\frac{\partial \alpha}{\partial \vec{n}} = \langle \text{grad } \alpha, \vec{n} \rangle = \langle 2\nabla_V V, \vec{n} \rangle.$$

Finally, we solve (5) with the above boundary conditions to compute α for a data set, and the detailed implementation is discussed in Section 3. In the bottom image of Figure 1, we show a slice of our α solution for the synthetic torus in Section 5.1. The voxels are color coded from red (low value) to yellow (high value), and we can see that the interior of the α field has a higher value than the exterior. Scaling the inverse tensor metric with e^α makes the geodesics follow the desired directions. This is because it has a higher cost for a pathway to travel along the interior of the torus than the exterior. In addition, the α field is consistent with equation $\text{grad } \alpha = 2\nabla_V V$. As we mentioned earlier, $\nabla_V V$ measures how V is bending along its integral curve. Since V is rotating only in this case, the $\text{grad } \alpha$ points inward to the torus center, which means α should increase as we move from the exterior of the torus to the interior. In another perspective, since the curvature is higher in the interior of the torus than the exterior, the α also penalizes higher curvature. The higher the curvature, the higher the α . This also happens on real data. For example, in the axial slice of the brain as shown in Figure 2, we can see that the α is higher in the interior of the genu because the curvature is higher in the interior.

3. Numerical Implementation

A closer look at (5) reveals that it is nothing but an anisotropic Poisson equation on the image domain. The right-hand side is constant in α , and the Laplace-Beltrami operator on the left-hand side can be expressed as $\nabla \cdot (G \nabla \alpha)$, where G is a symmetric positive-definite matrix, is the usual Euclidean divergence operator, and is the gradient $\nabla \cdot$ operator in domain. ∇ the image Since the Poisson equation is on the image grid, we approximate both sides of the Poisson equation using a finite-difference approach.

It is common to use central, forward, or backward differences to approximate the first derivative. These techniques perform well on noiseless data, but they do not give adequate results on noisy data. So instead of using these differences, we use the noise-robust differentiator (Holoborodko, 2008) as our finite-difference approximation of the first derivative, which suppresses the high frequencies of the noise signal, is precise on low frequencies, and is particularly beneficial for noisy data.

Care must be taken to handle the sign ambiguity of the eigenvectors V when computing $\nabla_V V$. When we compute $\nabla_V V$ at a node, we need to make the eigenvectors around the node consistent. Our approach selects the eigenvector at the current node as a reference eigenvector and then chooses the directions that have smaller angle with the reference eigenvector of the eigenvectors in the neighborhood. Remember that since $\nabla_V V = \nabla_{(-V)}(-V)$, it does not matter which direction the reference eigenvector points to of the two possible directions.

After the discretization of the Poisson equation with the Neumann boundary condition, we will get a sparse linear system $A \cdot \alpha = b$ to solve. For sparse linear systems, there are many efficient iterative solvers, such as conjugate gradient (Hestenes and Stiefel, 1952) (CG). However, CG can apply only to symmetric and positive-definite matrices or at least positive-semidefinite matrices. In our discretization, the Laplace-Beltrami operator A on the left-hand side is asymmetric due to the asymmetric interaction between g^{-1} and the computation of $\nabla \alpha$ in different neighborhoods, so we have to use other solvers that can be

applied to asymmetric linear systems, such as the generalized minimal residual method (Saad and Schultz, 1986) (GMRES) and biconjugate gradient method (Fletcher, 1976) (BICG). For BICG, we need to compute A^T , which is not trivial, so we use GMRES as implemented in PETSc (Balay et al., 2013a,b, 1997) in our implementation.

GMRES solves a linear system by constructing a vector space with an orthonormal basis, and the solution can be written as the linear combination of the basis vectors. However, one drawback of the GMRES is that it also needs to save a certain number of basis vectors. In our real data study, to get a good solution, the GMRES usually runs hundreds of iterations and we need to save several hundreds of basis vectors. Each basis vector usually requires tens of megabytes of memory, which means we need at least several gigabytes of memory. However, the amount of memory usage and computing time depends on the number of iterations, and we can reduce the number of iterations by using a good initialization of α . Our initialization of α is computed by minimizing the L_2 norm of the difference between $\nabla \alpha$ and $2\nabla_v V$, i.e., we will minimize

$$\int \|\nabla \alpha - 2\nabla_v V\|^2 dx$$

by solving the Euclidean Poisson equation

$$\Delta \alpha = 2 \operatorname{div}(\nabla_v V),$$

where $\Delta \alpha = \operatorname{div}(\nabla \alpha)$ is the Laplace operator and div is just the Euclidean divergence. The reason we choose this initialization is that the Euclidean Poisson equation with Neumann conditions can be easily discretized to a symmetric positive definite linear system, which can be solved using CG efficiently in terms of both memory requirements and computing time. With the good initialization, we usually need tens of iterations instead of hundreds of iterations till the GMRES converges, and we have to save only tens of basis vectors. In our real data study, if we use the GMRES with the good initialization, the total cost is about two gigabytes of memory and five minutes instead of more than ten gigabytes of memory and forty minutes while using the GMRES with zero initialization.

4. Segmentation of White Matter Tracts Using Adaptive Geodesics

Front-propagation methods do not explicitly compute the geodesic curves, but instead compute a cost function $u(x)$, which is the time-of-arrival of the geodesic flow at the point x . The characteristic vectors of $u(x)$ give the tangent vectors along the geodesic, and the characteristic vectors are given by $T(x) = g(x)^{-1} \nabla u(x)$, where $g(x)$ is the Riemannian metric and $\nabla u(x)$ indicates the Euclidean gradient.

To segment a white matter tract, Fletcher et al. (2007) propose the idea of first defining two ROIs at each end of the tract and then thresholding the addition $u(x) = u_1(x) + u_2(x)$, where $u_1(x)$ and $u_2(x)$ are the two cost functions from the two ROIs. This threshold, however, is preset and the segmentations of white matter tracts vary hugely with the changes of the threshold value. Thus, users often need to manually adjust the threshold value, which is not only time-consuming but also may reduce the reliability and repeatability of the segmentations. This issue is critical especially for longitudinal studies where we need to analyze the changes of a white matter tract over time and we want the segmentations of the white matter tract to be consistent for different time points.

To overcome this drawback, we propose to segment a white matter tract based on the angles of the two characteristic vector fields from the two ROIs. Intuitively, the characteristic

vectors from the two ROIs would tend to point against each other inside the white matter tract, and tend to point to similar directions outside the tract as shown in the second image of Figure 3. In theory, for each point on the Riemannian manifold, there must exist a geodesic from each ROI to the point, which means we have two characteristic vectors at each point. If the point is on a geodesic from one ROI to the other ROI, the two characteristic vectors must point to opposite directions, though, in practice, we do not require that the two characteristic vectors be exactly opposite but have a large angle. When we compute the angle image of the two characteristic vector fields, since some characteristic vectors are pointing outwards at the boundary of the white matter boundary (Hao et al., 2011), we apply a median filter to the angle image to fix this artifact.

In order to segment the tract, we propose to threshold the angle image of the two characteristic vector fields. In addition, we do not want to manually set this threshold, as this may introduce reliability issues across different subjects or time points. Rather, we want to use an algorithm for automatically finding the optimal threshold for the angle image. Since the boundary of the tract of interest is very clear in the angle image (one example is shown in the second image of Figure 3), we propose to use Otsu's method (Otsu, 1979) to automatically segment the tract from the white matter.

Otsu's method automatically thresholds a gray-level image to a binary image by calculating the optimal threshold that minimizes the within-class variance

$$\sigma_w^2(t) = \omega_1(t) \sigma_1^2(t) + \omega_2(t) \sigma_2^2(t),$$

where $\omega_1(t)$ and $\omega_2(t)$ are the probabilities of the two classes thresholded by t , which are computed from the histogram, and $\sigma_1^2(t)$ and $\sigma_2^2(t)$ are the variance of the two classes. When the image to be thresholded contains two classes of voxels with well-separated intensity distributions, which holds for our angle images, Otsu's method performs well.

With Otsu's thresholding, we can segment the white matter tract of interest pretty well in most cases, but there is still an issue: the two characteristic vector fields sometimes could also meet at some white matter tracts other than the one of interest, and the segmented image would contain parts of some other white matter tracts (one example is shown in the third image of Figure 3). In the case that these extra parts are not connected to the tract of interest, we can exclude these parts by choosing only the component containing the two ROIs using the connected component method. If these extra parts, however, are connected to the tract of interest, the connected component method does not help.

To resolve this issue, we introduce a step before the Otsu's thresholding. We propose to threshold the cost image $u(x)$ to the limit where all the voxels inside the ROIs are included in the thresholded region. We do this because more and more voxels inside the ROIs are included in the thresholded region as the threshold increases, and once all the voxels inside the ROIs are included, voxels that do not belong to the tracts will be included if we continue increasing the threshold. This threshold could be the max cost of all the voxels inside the two ROIs, but in case there are some outlier voxels inside the ROIs, we use the 95% largest cost of the voxels inside the two ROIs instead. Usually this thresholding of the cost image results in a region that is slightly larger than the tract of interest. One example is shown in the fourth image of Figure 3. Though the choice of 95% threshold seems ad hoc and the increase of the threshold over-segments the tract of interest even more, our method is not sensitive to this specific threshold value. The reason is that the following Otsu's thresholding of the angle image from the thresholded region reduces the amount of over-segmentation as

shown in the last image of Figure 3. Moreover, this thresholding strategy can be replaced with some outlier detection algorithms.

To summarize, our segmentation algorithm has two steps:

1. Segment a relatively small region containing the tract of interest by automatically thresholding the cost image $u(x)$ as described above.
2. Based on the angle image, use Otsu's method to automatically segment the tract of interest from the thresholded region from step 1.

5. Results

In this section, we demonstrate the improvement of geodesic flows generated by our metric modulating method compared to flows computed with the inverse-tensor metric and the sharpened-tensor metric (Fletcher et al., 2007) using both synthetic and real DTI data. Our measure of quality is how well the geodesics from the three methods follow the principal eigenvectors of the noise free tensors. However, as mentioned in Section 4, front-propagation methods do not explicitly compute the geodesic curves, but instead compute a function $u(x)$, which is the time-of-arrival of the geodesic flow at point x . The characteristic vectors of $u(x)$ give the tangent vectors along the geodesic. In the case of the inverse-tensor metric, $D(x)^{-1}$, the characteristic vectors are given by $T(x) = D(x)\nabla u(x)$. In the case of the sharpened-tensor metric, the characteristic vectors are given by $T(x) = M(x)\nabla u(x)$, where

$$M(x) = |D(x)|^{\frac{1}{3}} \left(\frac{D(x)}{|D(x)|^{\frac{1}{3}}} \right)^{\beta}, \quad (6)$$

and we use $\beta = 3$ in all our experiments. In the case of our adaptive metric, the characteristic vectors are given by $T(x) = e^{-\alpha(x)}D(x)\nabla u(x)$. Here $\nabla u(x)$ indicates the Euclidean gradient, which we approximate with finite differences, as described in Jackowski et al. (2005).

We compute $u(x)$ by solving a Hamilton-Jacobi equation using the Fast Iterative Method, as described in Fletcher et al. (2007). For visualization purposes, we compute the geodesics from both methods by integrating their characteristic vectors. Because these vectors always point away from the source region, we compute geodesic curves by integrating the characteristic vectors backward from any target point in the tensor field. These integral curves of the negative characteristic vectors are guaranteed to end up in the source region.

In addition, we also demonstrate the improvement of segmentations generated based on our metric modulating method compared to those computed with the inverse-tensor metric, the sharpened-tensor metric, the BF and multi-ROI approach, and stochastic tractography (Friman et al., 2006) using both synthetic and real data. The segmentation algorithms based on geodesic tracking are mentioned in Section 4. We also test the BF and multi-ROI approach because it increases the validity of tractography as shown in Huang et al. (2004). In this approach, a fiber is traced for all pixels in the image, but only those fibers that penetrate the predefined ROIs are selected. Also, for a fair comparison with other segmentations, we truncate the selected fibers if they pass beyond the ROIs. For the deterministic tractography algorithm used in the BF + two-ROI approach, we test continuous tracking (FACT) (Mori et al., 1999a) and tensor line (Weinstein et al., 1999) methods for both synthetic and real data. In addition, since we use a two-tensor model (7) to simulate several fiber crossing data in Section 5.3, we also do two-tensor estimation in the crossing regions and test two-tensor FACT and two-tensor tensor line methods for the generated fiber crossing data. In the tractography algorithms used in our experiments, our step size is 0.1 of the slice thickness, and for the deterministic tractography algorithms, the streamline stops

when the change of angle is larger than 75° during tracking. The stochastic tractography algorithm we used in our experiment is Bayesian tractography with 10,000 tracts initiated from each seed point. To compare with other two ROIs approaches, we select only those fibers starting from one ROI passing through the other ROI, and the connectivity map is computed from the selected fibers. In addition, for quantitative validation purposes, we also truncate the selected fibers as mentioned earlier, and threshold the connectivity maps with threshold value zero to get binary segmentations of the tracts. We use the implementation in CAMINO (Cook et al., 2006) for the two-tensor estimation, deterministic tractography, and stochastic tractography.

In the following sections, we first validate our adaptive Riemannian metric on clean synthetic data in Section 5.1 and on noisy synthetic data in Section 5.2. Then, we compare our segmentation method with other approaches and perform quantitative evaluation of the segmentations both on synthetic data in Section 5.3 and on real data in Section 5.4.

5.1. Clean Synthetic Curved Tensor Data

To test our method, we generate a synthetic curved tensor field that has similar properties to many white matter tracts in the brain. The synthetic data is the top half of a solid torus, where the tensors rotate along the large circle of the torus. The torus has a minor radius of 8 voxels and a major radius of 40 voxels. Each tensor in the tensor field has the same eigenvalues $(16, 4, 4) \times 10^{-4} \text{ mm}^2/\text{s}$. A middle slice of the tensor field is shown in the first column of Figure 4. The source region for the front-propagation method is shown in white.

In the third column of Figure 4, we compare the characteristic vector field (shown in blue) of the generated noise free data with the principal eigenvector field (shown in red). Comparing the three images, we can clearly see the characteristic vectors T follow the principal eigenvectors V much better in the middle and bottom images. We do not show the vectors at the boundary of the white matter, for some characteristic vectors are pointing outwards because of the aliasing artifacts (Hao et al., 2011), and we will use this convention in the following descriptions.

In this synthetic example, we can compute the analytic solution of $a(x)$, which is $a(x) = -2 \ln r(x) + C$, where $r(x)$ is the distance from x to the center of the torus, and C is some constant. We computed the difference between our numerical solution and the analytic $a(x)$, and the result was within numerical error. We also computed the root mean square error (RMSE) of the angles between the geodesic tangent vectors and principal eigenvectors. The RMSE for our adaptive metric is 1.62° compared to 12.21° for the inverse-tensor metric and 0.84° for the sharpened-tensor metric. Again, we did not count the angles at the boundary when we computed RMSE, and we will use this convention in the following descriptions.

In the first column of Figure 4, we visualize the integrated geodesics of the generated noise free data between some target points (on the right side of the torus) and the source region (shown in white). Under the sharpened-tensor metric and adaptive metric, the geodesics follow the principal eigenvectors of the tensor field and arrive at a point inside the source region. In contrast, the geodesics under the inverse-tensor metric without modulation, starting from the same target points, take a shortcut and end up at the closest point inside the source region by closely following the boundary constraints.

As shown in Figure 4, the sharpened-tensor metric and our adaptive metric have better characteristic vector fields and geodesics than the inverse-tensor metric. The sharpened-tensor metric is even better than our adaptive metric in terms of RMSE of the angles as shown in Table 1. This makes sense as we make the tensor sharper by increasing β in (6), the cost of moving in directions other than the principal eigenvector will increase, and the

geodesics will tend to follow the principal eigenvectors more closely. In the limit case, the geodesics will follow the principal eigenvectors perfectly. However, we do not want to let the geodesics follow the principal eigenvectors too closely on noisy data, for this would decrease the robustness of the front-propagation method to noise as we demonstrate in Section 5.2. Besides, the amount of sharpening is an ad hoc parameter and even if there exists an optimal amount of sharpening, it is not clear how to compute this optimal value.

5.2. Noisy Synthetic Curved Tensor Data

To test the robustness of our method, we simulate Rician noise on the clean synthetic data we generated in Section 5.1, with signal-to-noise ratio (SNR) of 10, 15, and 20. To simulate Rician noise, we use the Stejskal-Tanner equation to generate 12 “clean” diffusion weighted images (DWIs) with $b = 1000\text{s/mm}^2$ and one non-diffusion-weighted image. We then simulate Rician noise on each DWI and estimate the noisy synthetic tensor data.

We do the same comparison for the noisy synthetic data as we did for clean synthetic data. In the fourth column of Figure 4, we compare the characteristic vector field (shown in blue) of the generated noisy data at an SNR of 15 with the principal eigenvector field (shown in red) of the *clean synthetic data*. We can clearly see the characteristic vectors T follow the principal eigenvectors V much better for the sharpened-tensor metric and adaptive metric compared to inverse-tensor metric, but we can also see some characteristic vectors T have a relative large angle with the principal eigenvectors V for the sharpened-tensor metric. Again, we compute the RMSE of the angles between the geodesic tangent vectors and clean principal eigenvectors. The RMSE for our adaptive metric is 5.94° compared to 16.35° for the inverse-tensor metric and 6.97° for the sharpened-tensor metric.

In addition, in Table 1, we can see that the RMSE of the angles between the geodesic tangent vectors and *clean principal eigenvectors* increases as the SNR decreases for all three metrics. Though our modulating metric has a somewhat larger RMSE than the sharpened-tensor metric on the clean data, our modulating metric always has the lowest RMSE on the noisy data.

In the second column of Figure 4, we visualize the integrated geodesics of the generated noisy data at an SNR of 15 between a target region (on the right side of the torus) and a source region (shown in white). Under our adaptive metric, the geodesics follow the principal eigenvectors of the tensor field closely and arrive at a point inside the source region. In contrast, the geodesics under the inverse-tensor metric take a shortcut and end up at the closest point inside the source region by closely following the boundary constraints. The geodesics under the sharpened-tensor metric are better than the ones under the inverse-tensor metric, but are slightly worse than the ones under our adaptive metric.

5.3. Fiber Crossing Tensor Data

To test the effectiveness of our segmentation algorithm presented in Section 4 on crossing tracts, we generate three fiber crossing tensor fields with similar properties to many white matter tracts in the brain. One goal of this experiment was to test the influence of oblate tensors, where the principal eigenvectors that our adaptive metric attempts to follow are not well-defined. The synthetic images are two bars crossing at the center of the image at an angle of 60° and 90° , and a curved torus crossing with a cylinder. The bar has width 8 voxels, the cylinder has radius of 8 voxels, and the torus has minor radius of 8 voxels and major radius of 40 voxels. We show a center slice of each tensor field along with the ROIs we use in this experiment in the first column of Figure 5.

The tensors in each tract of the two crossing tracts have eigenvalues $(16,4,4)\times 10^{-4}\text{mm}^2/\text{s}$. Each voxel of the generated DWI in the crossing area was computed based on the two-tensor model,

$$S=S_0 \left(f e^{-bg^t D_1 g} + (1-f) e^{-bg^t D_2 g} \right), \quad (7)$$

where S_0 is the baseline image, g is the gradient direction, and D_1 and D_2 are the crossing tensors at a voxel. In this case, we use 64 gradient directions with $b = 1000\text{s/mm}^2$ and $f = 0.5$. In addition, the DWI was corrupted by Rician noise to simulate SNR of 10, 15, and 20.

In Tables 2 and 3, we compare the sensitivity, specificity, and Dice coefficient of the segmentations computed by eight methods. While the sensitivity and specificity represent only the true positive rate and the true negative rate of the segmentations, respectively, the Dice coefficient measures how similar the segmentations and the ground truth are overall. A segmentation with either a higher sensitivity or a higher specificity does not mean this segmentation is better. For example, on one hand, if a segmentation has every pixel, it will have sensitivity 1 and specificity 0; on the other hand, if a segmentation has no voxel, it will have sensitivity 0 and specificity 1, which is the case for the single-tensor FACT method in Table 3. So the Dice coefficient is of most interest and a high Dice coefficient should have both high sensitivity and high specificity.

From the two tables, we can see that it is very difficult for the single-tensor FACT and tensor line methods to pass the 60° crossing area. The stochastic tractography method works somewhat better, but its Dice coefficients are still low compared to our method. The major reason is that for these three methods, most of the fibers from one tract enter the wrong tract by following the ill-defined principal eigenvectors in the crossing region. The two-tensor methods perform much better than the single-tensor methods. For the 90° crossing data, most of the fibers of the single-tensor FACT method cannot pass through the crossing region because the principal eigenvectors are random at the crossing region, but the tensor line method works well because the tensors in the crossing region looks like pancakes and the fibers tend to go straight. The two two-tensor methods work better compared to the 60° crossing case, for the crossing angle is larger, which is easier for the two-tensor estimation algorithm to separate the two tensors. For the image containing a curved tract crossing with a cylinder, the single-tensor FACT and stochastic tractography methods work better compared to the other two cases because the principal eigenvectors are not random except in the center slice of the crossing region, and most of the fibers do not go to the wrong tract by following the principal eigenvectors.

For the segmentations based on the geodesic tracking with three metrics, we can see that these three methods outperform the other five tractography algorithms and that the three methods have very similar Dice coefficients, which are all above 99%, for both 60° and 90° crossing. For the data where a curved tract crosses with a cylinder, although these three methods still outperform the other five tractography algorithms, the segmentations based on our modulating metric have the best Dice coefficient, and their differences from the Dice coefficients of the segmentations based on the other two metrics are greater.

Besides, the geodesic methods have very stable segmentations as the noise level changes, but the tractography methods are relatively more sensitive to noise. In addition, for different crossing data, the best tractography algorithm changes, so it is unclear how to choose the most appropriate tractography algorithm if we want to use tractography to segment tracts.

In the last three columns of Figure 5, we show a slice of the segmentations using the best tractography methods (third column) in terms of the Dice coefficient, using the sharpened-

tensor metric (fourth column), and using our adaptive metric (fifth column). We can see that the segmentations based on our adaptive metric are favorable in all three cases. In the second column of Figure 5, we show the connectivity maps of the stochastic tractography. We can see that the stochastic results miss numerous voxels that are inside the tract of interest.

5.4. Real Data

We now show the results of our method applied to DTI of 10 healthy volunteers. DTI data were acquired on a Siemens Trio 3.0 Tesla Scanner with an eight-channel, receive-only head coil. DTI was performed using a single-shot, spin-echo, EPI pulse sequence and SENSE parallel imaging (undersampling factor of 2). Diffusion-weighted images were acquired in 12 non-collinear diffusion encoding directions with diffusion weighting factor $b=1000$ s/mm² in addition to a single reference image ($b=0$). Data acquisition parameters included the following: contiguous (no-gap) fifty 2.5mm thick axial slices with an acquisition matrix of 128×128 over a FOV of 256 mm (2×2 mm² in-plane resolution), four averages, repetition time (TR) = 7000 ms, and echo time (TE) = 84 ms. Eddy current distortion and head motion of each data set were corrected using an automatic image registration program (Rohde et al., 2004). Distortion-corrected DW images were interpolated to $1 \times 1 \times 1$ mm³ voxels, and six tensor elements were calculated using weighted least squares. The tensor upsampling is done only for the purposes of numerical computations on the voxel grid; a finer grid results in higher numerical accuracy.

5.4.1. Comparison of Characteristic Fields and Geodesics—In the left column of Figure 6, we compare the characteristic vector field T (shown in blue) with the principal eigenvector field V (shown in red) of the corpus callosum. Both sharpened tensor and adaptive metrics result in characteristic vectors that tend to follow the main eigendirections better. We computed the RMSE of the angles between the geodesic tangent vectors and principal eigenvectors. The RMSE with our modulation is 9.31° compared to 17.74° without modulation and 7.28° with sharpened modulation. Note here that there is no ground truth, so a lower RMSE just gives a sense of how closely the characteristic vectors follow the noisy eigenvectors. As discussed earlier, following the principal eigenvectors too closely results in a decrease in robustness to noise.

In the right column of Figure 6, as in the synthetic example, we track backward from target points (in the upper right side of the image) to a source region (upper left). Again, the geodesics under the inverse-tensor metric take a shortcut and merge into the closest point in the source region. In contrast, the geodesics under the sharpened-tensor metric and our adaptive metric more faithfully follow the tensor directions. Geodesics from our adaptive metric are drawn together slightly because the tensor field is thinner around the corner of the corpus callosum. The main paths of the adaptive metric and sharpened-tensor metric are similar, which is consistent with our Section 5.1 experiments on the synthetic torus without noise.

5.4.2. Comparison of White Matter Tract Segmentations—We selected four white matter tracts—the uncinate fasciculus, genu, arcuate fasciculus, and corticospinal tract—to validate our segmentation algorithm on the real DTI from 10 subjects. We compare the segmentations of different algorithms as we did in Section 5.3, except that we do not test the two-tensor methods, for the diffusion-weighted images were acquired in only 12 diffusion encoding directions. To get a quantitative validation, we need the ground-truth segmentation on these real data, and for this we use the Johns Hopkins University DTI-based white matter atlases (Mori et al., 2005; Wakana et al., 2007; Hua et al., 2008). There are both a white matter labels atlas, created by hand segmentation, and a white matter tractography atlas,

created by averaging the results of deterministic tractography. We first transform both labels to each of the 10 subject's anatomy by registering the atlas to each subject's anatomy using diffeomorphic registration. We then use these registered labels as a starting point for manual segmentations of the 10 images. This allowed for manual correction of segmentations due to anatomical variability and registration errors.

Table 4 gives the Dice coefficients and specificity/sensitivity of all methods, compared to the manual segmentation ground truth. Here we see similar behavior as we found in the segmentation of synthetic data in Section 5.3. First, we can see that the tractography methods have difficulties in connecting two regions, especially in the presence of partial voluming or fiber crossing. For example, the Dice coefficients of the tractography methods are all below 0.6 for the uncinate fasciculus, arcuate fasciculus, and corticospinal tract, where there is either fiber crossing or fiber kissing inside each tract. The tractography methods work better on the genu, in terms of Dice coefficients, because there are no fiber crossings inside this tract. In all four tracts, the front-propagation approaches perform better than the tractography methods in terms of Dice coefficients, with our adaptive metric performing the highest. For the corticospinal tract, compared to the other three highly curved tracts, the three different geodesic methods have very close Dice coefficients, and the one with our modulated metric does not improve the segmentation by much. This is consistent with what we found on the synthetic data, where the geodesic methods have similar Dice coefficients when the tracts are straight, and our adaptive metric is better when the tracts are curved.

In Figures 7, 9, and 10, we show slices of the ground-truth segmentations, the segmentations of the best tractography methods in terms of the Dice coefficient, the connectivity map from stochastic tractography, and the segmentations based on geodesic tracking. We can see that in all cases, the BF + two-ROI methods tend to under-segment the white matter tracts. They often miss numerous voxels that are obvious inside the tracts of interest. This is especially obvious in the segmentations of uncinate fasciculus (Figure 10). These two methods can usually segment the central part of the tracts, but often miss voxels around the boundary of the tracts because the fibers passing these voxels are likely to stop at the boundary of the white matter or enter into other tracts because of partial voluming or fiber crossing. In addition, we can also see that the BF + two-ROI approach sometimes includes some voxels outside the tracts of interest. For example, in the segmentation of the genu from the tensor line BF + two-ROI method in Figure 7, we can see that there is an isolated segment in the segmentation. This is caused by some fibers that first went to other tracts of the brain and then made a U-turn before they arrived at the ROI as shown in Figure 8. This problem might be solved by changing our stop criteria, i.e., decreasing the stop angle threshold for tractography, but this results in worse results by also decreasing the volume of the segmentations.

The stochastic tractography has the same under-segmentation problems, as shown by their resulting probability maps. First of all, they miss many voxels around the boundary of the tracts, and we can also see that the central part of the tract has higher probabilities than the voxels close to the boundary of the white matter tracts. For the connectivity map of the uncinate (Figure 10), only a few voxels have intensity in the shown slice, and even in the maximum intensity projection along the sagittal plan (Figure 11), the connectivity map still looks very thin. It is obvious that it is very difficult to get a good segmentation based on the probability map. This makes sense because if a fiber from deterministic tractography stops at the boundary of the white matter, even if we initiate a lot more fibers, there is still a big chance that the fibers arrive at the boundary and stop there. Moreover, both the BF method and stochastic tractography are computationally-intensive approaches. Though the BF

process is a preprocessing step and it needs to be done only once, we usually need several gigabytes of storage space to save the whole brain tractography result.

On the other hand, for geodesic-based segmentation methods, the segmentations based on the inverse-tensor metric and sharpened-tensor metric are better than the BF + two-ROI approach in terms of missing voxels. For example, they delineate the structure of the uncinate fasciculus better, although they sometimes also miss some voxels as shown in Figure 10. However, they also clearly under-segment the genu of the corpus callosum in Figure 7. This is from missing numerous voxels at the posterior side of the tract, due to the same curvature effects and shortcutting seen in the synthetic torus. Also, they over-segment the tracts as shown in the segmentations of the arcuate fasciculus in Figure 9. Here they both over-segment the tract because they include many voxels below the tract which should be inside the inferior longitudinal fasciculus. However, our adaptive metric results in improved segmentations of these tracts. In the segmentations of the genu in Figure 7, our method corrects the shortcutting and includes most voxels of the genu. Our method's segmentations of the direct arcuate fasciculus are also improved, as seen in Figure 9. While the BF + two-ROI approach under-segments the tract, and the other two geodesic-based methods over-segment the tract, our method finds a good balance between the two (although it also includes a few voxels inside the inferior longitudinal fasciculus).

6. Conclusion and Future Work

We presented a new geodesic-based algorithm for computing white matter pathways and segmenting white matter tracts in DTI. We formulate a simple scalar field modulation of the Riemannian metric, and the appropriate variational problem results in a Poisson's equation on the Riemannian manifold. We also develop a segmentation framework based on the computed geodesics. Since the modification and segmentation algorithm are formulated on a general Riemannian manifold, our method can also be applied to problems for which there exist a Riemannian metric (which is just the Euclidean metric in Euclidean space) and preferred geodesics.

We have demonstrated the advantages of our method over the state-of-the-art methods on both synthetic and real data. We have shown that our adaptive Riemannian metric results in geodesics that more faithfully follow the principal eigenvectors of the diffusion tensor field, resulting in segmentations that better delineate the white matter tracts, especially in tracts with high curvature. In addition, there are almost no parameters in our entire framework, and the only user intervention we need is identification of ROIs. As such, our method can be very useful for large data clinical studies, where fewer parameters and user intervention are preferred. Furthermore, our method is less sensitive to local perturbations, such as noise, partial volume effects, or fiber crossing, compared to tractography algorithms. First, our adaptive metric is computed by minimizing functional (4), which makes our metric less sensitive to the noise as shown in Section 5.2. Second, the front-propagation method computes the pathways by optimizing a global functional using the whole DTI data, and hence is less influenced by noise, partial volume effects, or fiber crossing.

We have identified three areas as potential future work. First, as the proposed method performs well on highly curved tracts, it might be able to recover U-shaped tracts between neighboring gyri, which are poorly recovered by classical tensor based tractography (Conturo et al., 1999; Catani et al., 2012). Second, although our method is based on a DTI model, it is possible to generalize our method to high-angular resolution diffusion imaging (HARDI), such as the multi-tensor model. Moreover, since the proposed method computes binary segmentation of the tracts, which sometimes under- or over-segment tracts, it might help to incorporate some anatomical prior, such as a white matter atlas, into our

segmentation framework. Another way to solve this problem is to compute fractional segmentation of the tracts instead of binary segmentation. Hao and Fletcher (2013) propose a fractional segmentation method to jointly solve the tract segmentation and multi-tensor model estimation in DWI. This method can both reduce the partial volume effect and reliably estimate multiple tensor compartments in fiber crossing regions even with low angular DWI, but this method needs more validation.

Acknowledgments

We would like to thank Dr. Janet Lainhart for providing the image data, funded by NIH Grant R01 MH080826. This work was supported by NIH Grant R01 MH084795 and the National Alliance for Medical Image Computing (NAMIC): NIH Grant U54 EB005149.

References

- Awate, SP.; Zhang, H.; Gee, JC. Proceedings of the 10th international conference on Medical image computing and computer-assisted intervention - Volume Part I. Springer-Verlag; Berlin, Heidelberg; 2007. Fuzzy nonparametric DTI segmentation for robust cingulum-tract extraction; p. 294-301.
- Balay, S.; Brown, J.; Buschelman, K.; Eijkhout, V.; Gropp, WD.; Kaushik, D.; Knepley, MG.; McInnes, LC.; Smith, BF.; Zhang, H. PETSc Users Manual. Argonne National Laboratory; 2013a. Technical Report ANL-95/11 - Revision 3.4
- Balay S, Brown J, Buschelman K, Gropp WD, Kaushik D, Knepley MG, McInnes LC, Smith BF, Zhang H. PETSc. 2013b Web page. <http://www.mcs.anl.gov/petsc>.
- Balay, S.; Gropp, WD.; McInnes, LC.; Smith, BF. Efficient management of parallelism in object oriented numerical software libraries. In: Arge, E.; Bruaset, AM.; Langtangen, HP., editors. Modern Software Tools in Scientific Computing. Birkhäuser Press; 1997. p. 163-202.
- Barbieri S, Bauer MHA, Klein J, Moltz JH, Nimsy C, Hahn HK. DTI segmentation via the combined analysis of connectivity maps and tensor distances. *NeuroImage*. 2012; 60:1025–1035. [PubMed: 22293133]
- Barnea-Goraly N, Menon V, Eckert M, Tamm L, Bammer R, Karchemskiy A, Dant CC, Reiss AL. White matter development during childhood and adolescence: a cross-sectional diffusion tensor imaging study. *Cerebral Cortex*. 2005; 15:1848–1854. [PubMed: 15758200]
- Basser PJ, Pajevic S, Pierpaoli C, Duda J, Aldroubi A. In-vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*. 2000; 44:625–632. [PubMed: 11025519]
- Bazin PL, Ye C, Bogovic JA, Shiee N, Reich DS, Prince JL, Pham DL. Direct segmentation of the major white matter tracts in diffusion tensor images. *Neuroimage*. 2011; 58:458–468. [PubMed: 21718790]
- Behrens T, Woolrich M, Jenkinson M, Johansen-Berg H, Nunes R, Clare S, Matthews P, Brady J, Smith S. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*. 2003; 50:1077–1088. [PubMed: 14587019]
- Catani M, DellAcqua F, Vergani F, Malik F, Hodge H, Roy P, Valabregue R, Thiebaut de Schotten M. Short frontal lobe connections of the human brain. *cortex*. 2012; 48:273–291. [PubMed: 22209688]
- Conturo TE, Lori NF, Raichle ME. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences*. 1999; 96:10422.
- Cook, PA.; Bai, Y.; Gilani, NS.; Seunarine, KK.; Hall, MG.; Parker, GJ.; Alexander, DC. Camino: Open-source diffusion-mri reconstruction and processing; 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine; 2006; p. 2759
- Corouge I, Fletcher PT, Joshi S, Gouttard S, Gerig G. Fiber tract-oriented statistics for quantitative diffusion tensor MRI analysis. *Medical Image Analysis*. 2006; 10:786–798. [PubMed: 16926104]
- Fletcher, PT.; Tao, R.; Jeong, WK.; Whitaker, RT. A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. In: Karssemeijer, N.; Lelieveldt, B.,

- editors. Information Processing in Medical Imaging (IPMI). Springer; Heidelberg: 2007. p. 346-358.
- Fletcher, R. Conjugate gradient methods for indefinite systems. In: Watson, G., editor. Numerical Analysis. Springer Berlin; Heidelberg: 1976. p. 73-89. volume 506 of *Lecture Notes in Mathematics* 10.1007/BFb0080116
- Friman O, Farnebeck G, Westin CF. A bayesian approach for stochastic white matter tractography. *IEEE Trans. Med. Imaging*. 2006;965–978. [PubMed: 16894991]
- Hagmann P, Kuran M, Gigandet X, Thiran P, Wedeen VJ, Meuli R, Thiran JP. Mapping human whole-brain structural networks with diffusion mri. *PloS one*. 2007; 2:e597. [PubMed: 17611629]
- Hao, X.; Fletcher, P. Joint fractional segmentation and multi-tensor estimation in diffusion mri. In: Gee, J.; Joshi, S.; Pohl, K.; Wells, W.; Zillei, L., editors. Information Processing in Medical Imaging. Springer; Berlin Heidelberg: 2013. p. 340-351. volume 7917 of *Lecture Notes in Computer Science*
- Hao, X.; Whitaker, RT.; Fletcher, PT. Adaptive riemannian metrics for improved geodesic tracking of white matter. Proceedings of the 22nd international conference on Information processing in medical imaging; Springer-Verlag, Berlin, Heidelberg. 2011; p. 13-24.
- Hestenes MR, Stiefel E. Methods of Conjugate Gradients for Solving Linear Systems. *Journal of Research of the National Bureau of Standards*. 1952; 49:409–436.
- Holoborodko P. Smooth noise robust differentiators. 2008 <http://www.holoborodko.com/pavel/numerical-methods/numerical-derivative/smooth-low-noise-differentiators/>.
- Hua K, Zhang J, Wakana S, Jiang H, Li X, Reich DS, Calabresi PA, Pekar JJ, van Zijl P, Mori S. Tract probability maps in stereotaxic spaces: analyses of white matter anatomy and tract-specific quantification. *Neuroimage*. 2008; 39:336–347. [PubMed: 17931890]
- Huang H, Zhang J, van Zijl PC, Mori S. Analysis of noise effects on DTI-based tractography using the brute-force and multi-roi approach. *Magnetic Resonance in Medicine*. 2004; 52:559–565. [PubMed: 15334575]
- Jackowski M, Kao CY, Qiu M, Constable RT, Staib LH. White matter tractography by anisotropic wavefront evolution and diffusion tensor imaging. *Medical Image Analysis*. 2005; 9:427–440. ! ce:title; Medical Image Computing and Computer-Assisted Intervention - {MICCAI} 2004! / ce:title; ! xocs:full-name; Medical Image Computing and Computer-Assisted Intervention! / xocs:full-name;. [PubMed: 16040268]
- Jbabdi S, Bellec P, Toro R, Daunizeau J, Pélérini-Issac M, Benali H. Accurate anisotropic fast marching for diffusion-based geodesic tractography. *Journal of Biomedical Imaging*. 2008; 2:1-2–12. 2008.
- Jeong WK, Fletcher PT, Tao R, Whitaker RT. Interactive visualization of volumetric white matter connectivity in diffusion tensor MRI using a parallel-hardware Hamilton-Jacobi solver. *IEEE Transactions on Visualization and Computer Graphics*. 2007; 13:1480–1487. [PubMed: 17968100]
- Jones DK. Tractography gone wild: Probabilistic fibre tracking using the wild bootstrap with diffusion tensor mri. *IEEE Trans. Med. Imaging*. 2008; 27:1268–1274. [PubMed: 18779066]
- Koch MA, Norris DG, M HG. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage*. 2002; 16:241–250. [PubMed: 11969331]
- Lazar M, Alexander AL. Bootstrap white matter tractography (BOOT-TRAC). *NeuroImage*. 2005; 24:524–532. [PubMed: 15627594]
- Lenglet C, Rousson M, Deriche R, Faugeras O, Lehericy S, Ugurbil K. A riemannian approach to diffusion tensor images segmentation. In *Proc. IPMI*. 2005:591602.
- Melonakos, J.; Mohan, V.; Niethammer, M.; Smith, K.; Kubicki, M.; Tannenbaum, A. Finsler tractography for white matter connectivity analysis of the cingulum bundle. In: Ayache, N.; Ourselin, S.; Maeder, A., editors. MICCAI. Springer; Heidelberg: 2007a. p. 36-43.
- Melonakos, J.; Niethammer, M.; Mohan, V.; Kubicki, M.; Miller, J.; Tannenbaum, A. Locally-constrained region-based methods for dw-mri segmentation; *Computer Vision*, 2007. ICCV 2007. IEEE 11th International Conference on; 2007b. p. 1-8.
- Mori S, Crain BJ, Chacko VP, Van Zijl PCM. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*. 1999a; 45:265–269. [PubMed: 9989633]

- Mori S, Crain BJ, Chacko VP, van Zijl PCM. Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*. 1999b; 45:265–269. [PubMed: 9989633]
- Mori S, Wakana S, Van Zijl PC, Nagae-Poetscher L. MRI atlas of human white matter. *Am Soc Neuroradiology*. 2005
- Niethammer M, Zach C, Melonakos J, Tannenbaum A. Near-tubular fiber bundle segmentation for diffusion weighted imaging: Segmentation through frame reorientation. *NeuroImage*. 2009; 45:S123–S132. [PubMed: 19101640]
- O'Donnell, L.; Haker, S.; Westin, CF. New approaches to estimation of white matter connectivity in diffusion tensor MRI: elliptic PDEs and geodesics in a tensor-warped space. In: Dohi, T.; Kikinis, R., editors. *MICCAI*. Springer; Heidelberg: 2002. p. 459-466.
- Otsu N. A threshold selection method from gray-level histograms. *IEEE Transactions On Systems Man And Cybernetics*. 1979; 9:62–66.
- Parker G, Wheeler-Kingshott C, Barker G. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging*. 2002; 21:505–512. [PubMed: 12071621]
- Parker GJM, Haroon HA, Wheeler-Kingshott CAM. A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *J Magn. Reson. Im*. 2003; 18:242–254.
- Pichon, E.; Westin, CF.; Tannenbaum, A.; Duncan, J.; Gerig, G. *MICCAI*. Springer; Heidelberg: 2005. A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography; p. 180-187.
- Rohde G, Barnett A, Basser P, Marengo S, C CP. Comprehensive approach for correction of motion and distortion in diffusion-weighted mri. *Magnetic Resonance in Medicine*. 2004; 51:103–114. [PubMed: 14705050]
- Rousson, M.; Lenglet, C.; Deriche, R. *ECCV Workshops CVAMIA and MMBIA'04*. 2004. Level set and region based surface propagation for diffusion tensor mri segmentation; p. 123-134.
- Saad Y, Schultz MH. Gmres: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput*. 1986; 7:856–869.
- Wakana S, Caprihan A, Panzenboeck MM, Fallon JH, Perry M, Gollub RL, Hua K, Zhang J, Jiang H, Dubey P, et al. Reproducibility of quantitative tractography methods applied to cerebral white matter. *Neuroimage*. 2007; 36:630–644. [PubMed: 17481925]
- Wang Z, Vemuri B. DTI segmentation using an information theoretic tensor dissimilarity measure. *Medical Imaging, IEEE Transactions*. 2005; 24:1267–1277.
- Weinstein, D.; Kindlmann, G.; Lundberg, E. *Proceedings of the conference on Visualization '99: celebrating ten years*. IEEE Computer Society Press; Los Alamitos, CA, USA: 1999. Tensorlines: advection-diffusion based propagation through diffusion tensor fields; p. 249-253.
- Zhukov L, Museth K, Breen D, Whitaker R, Barr AH. Level set modeling and segmentation of dt-mri brain data. *JOURNAL OF ELECTRONIC IMAGING*. 2003; 12:125–133.
- Ziyan, U.; Tuch, D.; Westin, CF. Segmentation of thalamic nuclei from DTI using spectral clustering. In: Larsen, R.; Nielsen, M.; Sporring, J., editors. *Medical Image Computing and Computer-Assisted Intervention MICCAI 2006*. Springer; Berlin/Heidelberg: 2006. p. 807-814.

Highlights

- We develop a new Riemannian metric that resolves high-curvature tracts deviation.
- The adaptive metric is computed by solving a variational problem on the manifold.
- Our method computes geodesics that more faithfully follow the principal eigenvectors.
- We present an automatic segmentation method that does not require user intervention.
- Our method improves segmentations of white matter tracts, including crossing tracts.

Figure 1.

Top: Diagram of various pathways between two points in a curved tensor field: the desired path following the principal eigenvectors (blue), the shortest path under the Euclidean metric (red), and the compromise path taken when using the inverse tensor field as metric (magenta). Bottom: A slice of our $a(x)$ solution for the synthetic data in Section 5.1. (Voxels are color coded from red (low value) to yellow (high value).)

Figure 2.

Top: Slices of the α solution for a read data in different views (Voxels are color coded from red (low value) to yellow (high value).). Bottom: Corresponding color-coded principal eigenvector image.

Figure 3.

From left to right: A slice of the cost image $u(x)$; the angle image between two characteristic vector fields, voxels colors are from white (small angle) to red (large angle); result of Otsu's thresholding directly on the angle image (shown in white); result of 95% thresholding of the cost image $u(x)$ (shown in red); and result of Otsu's thresholding of the angle image from the 95% thresholded region.

Figure 4.

First and second columns: geodesics of the generated noise free data (first column) and noisy data at an SNR of 15 (second column) emanating from the targets points (right side of the torus) to the source region (white). Third and fourth columns: tangent vectors of the geodesics (blue) of the generated noise free data (third column) and noisy data at an SNR of 15 (fourth column) under the inverse-tensor metric without modulation (top row), sharpened tensor metric (middle row), and with our modulation (bottom row). The red vectors are the principal eigenvectors of the diffusion tensors. We subsample the tensor field by a factor of 4 both horizontally and vertically in order to visualize it.

Figure 5.

A slice of the fiber crossing data at an SNR of 20, and the corresponding connectivity map and segmentations (shown white on top of the color-coded principal eigenvector images). First column: the simulated fiber crossing tensor fields with two ROIs (White and Gray). We subsample the tensor field by a factor of 3 both horizontally and vertically in order to visualize it. Second column: the connectivity map (voxels are color coded from red (low intensity) to yellow (high intensity)) from stochastic tractography. The background image is the fractional anisotropy (FA) image. Third column: the best tractography segmentation results in terms of the Dice coefficient, which are two-tensor FACT BF + two-ROI method, single-tensor tensor line BF + two-ROI method, and two-tensor tensor line BF + two-ROI method from top to bottom. The fourth (fifth) column show the segmentation results based on the sharpened-tensor metric (our modulating metric)

Figure 6.

Left column: tangent vectors of the geodesics (blue) under the inverse-tensor metric (top), sharpened-tensor metric (middle) and our adaptive metric (bottom) for a part of the corpus callosum. The red vectors are the principal eigenvectors of the diffusion tensors. The FA image is shown in the background. Right column: the geodesic flow in the corpus callosum from the target points (in the left frontal forcep) to the source region (in the right frontal forcep). The background images are the FA image and the diffusion tensor field.

Figure 7.

A slice of segmentations (shown in white) and connectivity map of the genu of the corpus callosum. Left column (from top to bottom): ground-truth segmentation, segmentation from single-tensor tensor line BF + two-ROI method, connectivity map from stochastic tractography; Right column (from top to bottom): segmentations based on geodesic tracking with inverse-tensor metric, sharpened-tensor metric and our adaptive metric.

Figure 8. Fibers generated from single-tensor tensor line BF + two-ROI method for the genu of the corpus callosum.

Figure 9.

A slice of segmentations (shown in white) and connectivity map of the direct arcuate fasciculus. Left column (from top to bottom): ground-truth segmentation, segmentation from single-tensor tensor line BF + two-ROI method, connectivity map from stochastic tractography; Right column (from top to bottom): segmentations based on geodesic tracking with inverse-tensor metric, sharpened-tensor metric and our adaptive metric

Figure 10.

A slice of segmentations (shown in white) and connectivity map of the uncinate fasciculus. Left column (from top to bottom): ground-truth segmentation, segmentation from single-tensor tensor line BF + two-ROI method, connectivity map from stochastic tractography; Right column (from top to bottom): segmentations based on geodesic tracking with inverse-tensor metric, sharpened-tensor metric and our adaptive metric.

Figure 11. Maximum intensity projection of the connectivity map along the sagittal plane for the uncinate fasciculus.

Table 1

The RMSE of the angles (in degree) between the geodesic tangent vectors and *clean principal eigenvectors* under different metrics and different noise level

	No Noise	SNR20	SNR 15	SNR 10
Inverse-tensor	12.21	15.76	16.35	18.35
Sharpened-tensor	0.84	5.31	6.97	10.70
Our adaptive	1.62	4.85	5.94	8.36

Table 2

Quantitative evaluation (Dice, sensitivity and specificity) of the segmentations on the generated noisy data at an SNR of 10 based on 1: inverse-tensor metric, 2: sharpened-tensor metric, 3: our modulating metric, 4: single-tensor FACT BF + two-ROI, 5: single-tensor tensor line BF + two-ROI, 6: stochastic tractography, 7: two-tensor FACT BF + two-ROI, 8: two-tensor tensor line BF + two-ROI

	60° Crossing			90° Crossing			Curved Crossing		
	Dice	SENS	SPEC	Dice	SENS	SPEC	Dice	SENS	SPEC
1	0.997	1	0.992	0.996	1	0.991	0.977	0.960	0.989
2	0.997	1	0.993	0.996	1	0.991	0.978	0.960	0.993
3	0.997	1	0.993	0.997	1	0.991	0.993	0.990	0.991
4	0.044	0.023	0.994	0.143	0.078	0.991	0.665	0.500	0.989
5	0.189	0.105	0.999	0.991	0.989	0.991	0.958	0.919	0.999
6	0.327	0.195	1	0.269	0.155	1	0.860	0.754	0.999
7	0.810	0.681	0.998	0.901	0.821	0.998	0.950	0.905	0.997
8	0.777	0.637	0.996	0.973	0.954	0.991	0.958	0.921	0.994

Table 3

Quantitative evaluation (Dice, sensitivity and specificity) of the segmentations on the generated noisy data at an SNR of 20 based on 1: inverse-tensor metric, 2: sharpened-tensor metric, 3: our modulating metric, 4: single-tensor FACT BF + two-ROI, 5: single-tensor tensor line BF + two-ROI, 6: stochastic tractography, 7: two-tensor FACT BF + two-ROI, 8: two-tensor tensor line BF + two-ROI

	60° Crossing			90° Crossing			Curved Crossing		
	Dice	SENS	SPEC	Dice	SENS	SPEC	Dice	SENS	SPEC
1	0.998	1	0.996	0.996	1	0.991	0.977	0.958	0.992
2	0.998	1	0.996	0.996	1	0.991	0.977	0.959	0.992
3	0.997	1	0.992	0.996	1	0.991	0.993	0.990	0.992
4	0	0	1	0.030	0.015	1	0.705	0.544	0.999
5	0.053	0.027	1	0.991	0.991	0.989	0.952	0.908	0.998
6	0.204	0.124	1	0.536	0.366	1	0.882	0.790	0.997
7	0.929	0.872	0.993	0.949	0.907	0.995	0.967	0.939	0.992
8	0.845	0.749	0.995	0.982	0.974	0.987	0.968	0.941	0.992

Table 4

Mean and standard deviation of segmentation metrics (Dice, sensitivity and specificity) for the uncinate fasciculus, genu, arcuate fasciculus and corticospinal tract from 10 subjects. Methods shown are 1: Inverse-tensor metric, 2: Sharpened-tensor metric, 3: Our modulating metric, 4: Single-tensor FACT BF + two-ROI, 5: Single-tensor Tensor line BF + two-ROI, 6: Stochastic tractography.

	uncinate			genu			arcuate fasciculus			corticospinal tract		
	Dice	SENS	SPEC	Dice	SENS	SPEC	Dice	SENS	SPEC	Dice	SENS	SPEC
1	.793 ± .054	.720 ± .087	.999 ± 0	.706 ± .132	.625 ± .163	.993 ± .004	.729 ± .052	.621 ± .074	.999 ± .001	.778 ± .050	.752 ± .089	.997 ± .001
2	.829 ± .047	.800 ± .099	.999 ± 0	.726 ± .078	.637 ± .114	.993 ± .003	.733 ± .045	.650 ± .080	.999 ± .001	.758 ± .050	.733 ± .088	.996 ± .001
3	.845 ± .036	.841 ± .044	.999 ± .001	.745 ± .046	.709 ± .111	.988 ± .006	.750 ± .034	.687 ± .054	.998 ± .001	.782 ± .041	.866 ± .049	.994 ± .002
4	.460 ± .121	.401 ± .119	.998 ± .001	.628 ± .091	.511 ± .117	.994 ± .004	.429 ± .107	.322 ± .094	.998 ± .001	.373 ± .120	.260 ± .098	.998 ± .001
5	.545 ± .096	.506 ± .126	.998 ± .001	.698 ± .046	.599 ± .075	.993 ± .003	.530 ± .110	.410 ± .112	.999 ± .001	.367 ± .095	.246 ± .083	.999 ± .001
6	.469 ± .158	.341 ± .136	.999 ± .001	.613 ± .074	.474 ± .103	.996 ± .004	.568 ± .123	.474 ± .122	.997 ± .003	.462 ± .098	.341 ± .095	.997 ± .002