

Adaptive Riemannian Metrics for Improved Geodesic Tracking of White Matter

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

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

Abstract. We present a new geodesic approach for studying white matter connectivity from diffusion tensor imaging (DTI). Previous approaches have used the inverse diffusion tensor field as a Riemannian metric and constructed white matter tracts as geodesics on the resulting manifold. These geodesics 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, it also has the serious drawback that geodesics tend to deviate from the major eigenvectors in high-curvature areas in order to achieve the shortest path. In this paper we formulate a modification of the Riemannian metric that results in geodesics adapted to follow the principal eigendirection of the tensor even in high-curvature regions. We show that this correction can be formulated as a simple scalar field modulation of the metric and that the appropriate variational problem results in a Poisson's equation on the Riemannian manifold. We demonstrate that the proposed method results in improved geodesics using both synthetic and real DTI data.

1 Introduction

Front-propagation approaches [11, 10, 4, 8, 13, 3] in diffusion tensor imaging (DTI) infer the pathways of white matter by evolving a level set representing the time-of-arrival of paths emanating from some starting region. 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. [10], 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, shortest paths will tend to prefer to remain tangent to major principal eigenvector of the diffusion tensor. While this is a powerful framework for computing white matter pathways, 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. In this paper we develop a new Riemannian metric, that relies on diffusion tensor

data but resolves this problem by adapting to high-curvature tracts, resulting in geodesic paths that more faithfully follow the principal eigenvectors.

1.1 Background

Front-propagation methods offer several advantages over conventional tractography [9, 1], in which streamlines (sometimes called tracts) are computed by forward integration of the principal eigenvector of the tensor. 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. The front-propagation algorithms are more robust to noise than tractography because 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 it decreases the overall cost.

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. These factors include accumulated errors in the streamlines throwing off the final destination and stopping criteria being triggered, either by low anisotropy tensors, due to noise or partial voluming, or sudden direction changes caused by noise. As shown by Fletcher et al. [3], front propagation methods can be used to segment white matter tracts by solving the geodesic flow from both regions and combining the resulting cost functions. This has the advantage that the solution will not get stuck in regions of noisy data or low anisotropy. This type of analysis is only appropriate if the endpoint regions are well known to be connected by a white matter tract because a white matter path will always be found. Although, if a “false positive” connection is found, this should be detectable as an unusually high cost function incurred by that pathway.

An alternative approach that deals with the problems arising from image noise is probabilistic tractography [6, 2, 12, 7], in which 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 to many hours), whereas efficient implementations of front-propagation solvers are much faster (typically requiring several seconds). The graphics processing unit (GPU) implementation by Jeong et al. [5] even runs at near real-time speeds. Also, probabilistic 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.

1.2 Properties of Front-Propagation

Despite the advantages that front-propagation methods have over tractography, there is one serious drawback. Figure 1 shows a diagram illustrating the prob-




Fig. 1. 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).

lem. 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.

This issue has been addressed previously [3] by “sharpening” the tensor, i.e., increasing the anisotropy by taking the eigenvalues to some power and renormalizing them. This increases the cost of moving in directions other than the principal eigenvector. In fact, the first front-propagation algorithm proposed by Parker et al. [11] 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. However, the amount of sharpening is an ad hoc parameter, and sharpening is applied equally across the image, rather than taking the curvature of the tract into account. Sharpening that increases with the curvature of the tract could be more effective. 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. 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.

Our proposed solution to this problem is to 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, which

requires solving a Poisson equation on the Riemannian manifold—however it does not require any arbitrary choice of parameters. We show that this solution 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 on both synthetic and real DTI data.

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 $e^{\alpha(x)}$, which results in the new geodesic energy functional

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

We call the function e^α the *metric modulating function* because it scales the Riemannian metric at each point. The exponentiation of α is to both 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, there are two reasons why we choose to modify the metric in this fashion. 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 3.

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'll 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 coefficient 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 variational 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., $W(0) = W(1) = 0$. To simplify notation, we will suppress the parameter t in most of the following. Then the variational of the energy functional is

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

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

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

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

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

This can only hold 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 similarly to the classical Poisson equation, with the exception that the div and grad operators used are the Riemannian versions. The divergence of X on M is defined in coordinates as

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

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

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

where $\Delta\alpha = \operatorname{div}(\operatorname{grad} \alpha)$ is the Laplace-Beltrami operator on M . The appropriate boundary conditions for this problem are the Neumann conditions,

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

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 $\nabla \cdot (A \nabla \alpha)$, where A is a symmetric positive-definite matrix and $\nabla \cdot$ and ∇ are the usual Euclidean divergence and gradient operators in the image domain. We solve this equation using a finite-difference scheme with a Jacobi iteration. There are more efficient solvers, such as conjugate gradient or multigrid methods, but the application of these methods to the proposed anisotropic operator with irregular boundary conditions remains an area of future development.

3 Results

In this section we demonstrate the improvement of geodesic flows generated by our metric modulating method compared to those computed with the inverse-tensor metric using both synthetic and real DTI data (Figure 2). Our measure of quality is how well the geodesics from the two methods follow the principal eigenvectors of the tensors. However, 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 the point x . The characteristic vectors of $u(x)$ give the tangent vectors along the geodesic. In the case of the inverse-tensor metric, the characteristic vectors are given by $T(x) = D(x)^{-1} \nabla u(x)$. In the case of our modulated metric, the characteristic vectors are given by $T(x) = e^{\alpha(x)} D(x)^{-1} \nabla u(x)$. Here $\nabla u(x)$ indicates the Euclidean gradient, which we approximate with finite differences, as described in [4].

We compute $u(x)$ by solving a Hamilton-Jacobi equation using the Fast Iterative Method, as described in [3]. For visualization purposes, we compute the geodesics from both methods by integrating their characteristic vectors. Because




Fig. 2. A slice of the synthetic curved tensor field (left). We subsample the tensor field by a factor of 4 both horizontally and vertically in order to visualize it. A slice of the color-coded principal eigenvector image (right). The area inside the red box is the region of interest used in the experiments.

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.

3.1 Synthetic Data

To test our method, we generate a synthetic curved tensor field which 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 smaller radius of 16 voxels and larger radius of 48 voxels. Each tensor in the tensor field has the same eigenvalues (3,1,1). A middle slice of the tensor field is shown in Figure 2. The source region for the geodesic front-propagation is shown in white.

In Figure 3, we compare the characteristic vector field (shown in blue) with the principal eigenvector field (shown in red). On the left, the characteristic vector field is computed using the inverse-tensor metric without modulation. On the right, the characteristic vector field is computed using the metric scaled by our computed metric modulating function. Comparing the two pictures, we can clearly see the characteristic vectors T follow the principal eigenvectors V much better in the right picture. Around the boundary, some characteristic vectors are pointing outwards in both cases. This is caused by aliasing: geodesics at the boundary must take sharp turns around voxel corners.

In this synthetic example, we can compute the analytic solution of $\alpha(x)$, which is $\alpha(x) = -2 \ln r(x) + C$, 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




Fig. 3. Tangent vectors of the geodesics (blue) under the inverse-tensor metric without modulation (left) and with modulation (right). The red vectors are the principal eigenvectors of the diffusion tensors. We subsample the vector field by a factor of 4 both horizontally and vertically in order to visualize it.




Fig. 4. The geodesics emanating from the targets points (right side of the torus) to the source region (white). We subsample the tensor field by a factor of 4 both horizontally and vertically in order to visualize it.

numerical solution and the analytic $\alpha(x)$, and the result is within numerical error. We also compute the root mean square error (RMSE) of the angle between the geodesic tangent vectors and principal eigenvectors. The RMSE with modulation is 10.6 degrees compared to 54.0 degrees without modulation. Most of the RMSE with modulation is caused by the aliasing artifacts at the boundary.

In Figure 4 we visualize the integrated geodesics between some target points (on the right side of the torus) and the source region (shown in white). The left picture shows the geodesics under the metric as the inverse diffusion tensor field. The right picture shows the geodesics under our modulated metric. Under the modulated 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.




Fig. 5. Tangent vectors of the geodesics (blue) under the inverse-tensor metric without modulation (left) and with modulation (right) for a part of the corpus callosum. The red vectors are the principal eigenvectors of the diffusion tensors. The fractional anisotropy (FA) image is shown in the background.

3.2 Real Data

We now show the results of our method applied to a corpus callosum tract from a DTI of a healthy volunteer. DTI data were acquired on a Siemens Trio 3.0 Tesla Scanner with an 8-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 twelve non-collinear diffusion encoding directions with diffusion weighting factor $b=1000$ s/mm^2 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 ($2 \times 2 \text{ mm}^2$ in-plane resolution), 4 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 [14]. Distortion-corrected DW images were interpolated to $1 \times 1 \times 1.25 \text{ mm}^3$ voxels, and six tensor elements were calculated using 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.

In Figure 5, we compare the characteristic vector field T (shown in blue) with the principal eigenvector field V (shown in red) of the corpus callosum. In the left picture, the characteristic vector field is computed using the inverse-tensor metric. In the right picture, the characteristic vector field is computed using our modulated metric. After the modulation, the characteristic vectors tend to follow the main eigendirections. As in the synthetic example, some characteristic vectors are pointing outwards near the boundary, which is caused by aliasing. Again, we compute the root mean square error (RMSE) of the angle between the geodesic tangent vectors and principal eigenvectors. The RMSE with modulation is 23.0 degrees compared to 37.3 degrees without modulation. Much of the RMSE with modulation is caused by the aliasing artifacts at the boundary.

In Figure 6, as in the synthetic example, we track backward from the target points, which are in the middle of corpus callosum in this case to the source region




Fig. 6. The geodesic flow in the corpus callosum from the target points (genu of the corpus callosum) to the source region (in the right frontal forcep). The background images are the FA image and the diffusion tensor field.

on the upper-left of the pictures. 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 our modulated metric more faithfully follow the tensor directions. In the latter case, geodesics are drawn together slightly because the tensor field is thinner around the corner of the corpus callosum.

4 Conclusion and Future Work

We presented a new geodesic front-propagation method for computing white matter pathways in DTI and showed that it results in geodesics that more faithfully follow the principal eigenvectors of the diffusion tensor field, especially in tracts with high curvature. There are two areas we have identified as potential future work. First, the aliasing artifacts along the white matter boundary described in Section 3 have, to the best of our knowledge, not been addressed in the literature. One possible solution to this problem would be to use a fuzzy boundary where the cost function increases to infinity along the normal direction. Currently the cost function changes instantaneously to infinity at the boundary (i.e., moving outside the boundary is infinite cost). Another issue is that the geodesics in front-propagation techniques can, in some cases, cross an edge between two adjacent tracts. We can envision a modification to our metric modulating function, e^α , that would increase the penalty for passing across such edges. This could be achieved by scaling the metric by a larger amount at edges, i.e., increasing the distances in these directions.

5 Acknowledgement

We would like to thank Dr. Janet Lainhart for providing the image data, funded by NIH Grant R01 MH080826, and Fangxiang Jiao for help with visualization.

This work was supported by NIH Grant R01 MH084795 and the National Alliance for Medical Image Computing (NAMIC): NIH Grant U54 EB005149.

References

1. Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A.: In-vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* 44, 625–632 (2000)
2. 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* 50, 1077–1088 (2003)
3. Fletcher, P.T., Tao, R., Jeong, W.K., Whitaker, R.T.: A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. In: Karssemeijer, N., Lelieveldt, B. (eds.) *Information Processing in Medical Imaging (IPMI)*. LNCS, vol. 4584, pp. 346–358. Springer, Heidelberg (2007)
4. Jackowski, M., Kao, C.Y., Qiu, M., Constable, R.T., Staib, L.H.: Estimation of anatomical connectivity by anisotropic front propagation and diffusion tensor imaging. In: Barillot, C., Haynor, D., Hellier, P. (eds.) *MICCAI*. LNCS, vol. 3217, pp. 663–670. Springer, Heidelberg (2004)
5. Jeong, W.K., Fletcher, P.T., Tao, R., Whitaker, R.T.: 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* 13(6), 1480–1487 (2007)
6. Koch, M.A., Norris, D.G., M, H.G.: An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage* 16, 241–250 (2002)
7. Lazar, M., Alexander, A.L.: Bootstrap white matter tractography (BOOT-TRAC). *NeuroImage* 24, 524–532 (2005)
8. 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. (eds.) *MICCAI*. LNCS, vol. 4791, pp. 36–43. Springer, Heidelberg (2007)
9. Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C.M.: Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45, 265–269 (1999)
10. O’Donnell, L., Haker, S., Westin, C.F.: 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. (eds.) *MICCAI*. LNCS, vol. 2488, pp. 459–466. Springer, Heidelberg (2002)
11. Parker, G., Wheeler-Kingshott, C., Barker, G.: Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging* 21, 505–512 (2002)
12. Parker, G.J.M., Haroon, H.A., Wheeler-Kingshott, C.A.M.: A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *J Magn. Reson. Im.* 18, 242–254 (2003)
13. Pichon, E., Westin, C.F., Tannenbaum, A.: A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography. In: Duncan, J., Gerig, G. (eds.) *MICCAI*. LNCS, vol. 3749, pp. 180–187. Springer, Heidelberg (2005)
14. Rohde, G., Barnett, A., Basser, P., Marengo, S., C, C.P.: Comprehensive approach for correction of motion and distortion in diffusion-weighted mri. *Magnetic Resonance in Medicine* 51, 103–114 (2004)