Published in final edited form as:

Med Image Anal. 2009 August ; 13(4): 580–597. doi:10.1016/j.media.2009.05.001.

Using Perturbation Theory to Reduce Noise in Diffusion Tensor Fields

Ravi Bansal^{1,2}, Lawrence H. Staib⁴, Dongrong Xu^{1,2}, Andrew F. Laine³, Jun Liu^{1,2}, and Bradley S. Peterson^{1,2}

¹Department of Psychiatry, Columbia University, New York, NY 10032

²New York State Psychiatric Institute, New York, NY 10032

³Department of Biomedical Engineering, Columbia University, New York, NY 10032

⁴Departments of Electrical Engineering and Diagnostic Radiology, Yale University, New Haven, CT 06512

Abstract

We propose the use of Perturbation theory to reduce noise in Diffusion Tensor (DT) fields. Diffusion Tensor Imaging (DTI) encodes the diffusion of water molecules along different spatial directions in a positive-definite, 3×3 symmetric tensor. Eigenvectors and eigenvalues of DTs allow the *in vivo* visualization and quantitative analysis of white matter fiber bundles across the brain. The validity and reliability of these analyses are limited, however, by the low spatial resolution and low Signal-to-Noise Ratio (SNR) in DTI datasets. Our procedures can be applied to improve the validity and reliability of these quantitative analyses by reducing noise in the tensor fields. We model a tensor field as a three-dimensional Markov Random Field and then compute the *likelihood* and the *prior* terms of this model using Perturbation theory. The *prior* term constrains the tensor field to be smooth, whereas the *likelihood* term constrains the smoothed tensor field to be similar to the original field. Thus, the proposed method generates a smoothed field that is close in structure to the original tensor field. We evaluate the performance of our method both visually and quantitatively using synthetic and real-world datasets. We quantitatively assess the performance of our method by computing the SNR for eigenvalues and the coherence measures for eigenvectors of DTs across tensor fields. In addition, we quantitatively compare the performance of our procedures with the performance of one method that uses a Riemannian distance to compute the similarity between two tensors, and with another method that reduces noise in tensor fields by anisotropically filtering the diffusion weighted images that are used to estimate diffusion tensors. These experiments demonstrate that our method significantly increases the coherence of the eigenvectors and the SNR of the eigenvalues, while simultaneously preserving the fine structure and boundaries between homogeneous regions, in the smoothed tensor field.

© 2009 Elsevier B.V. All rights reserved.

Corresponding Author: Dr. Ravi Bansal Contact Address: Room #2410, Unit 74 New York State Psychiatric Institute 1051 Riverside Dr. New York, NY 10032.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Keywords

Perturbation Theory; Markov Random Field; Diffusion Tensor; Simulated Annealing; Riemannian Distance

1 Introduction

Diffusion Tensor Imaging (DTI) measures quantitatively the diffusional properties of free water molecules along axons in white matter of the brain (Basser et al., 1994b; Pierpaoli et al., 1996; Melhem et al., 2002). Quantitative measurements of diffusion are estimated from a series of diffusion-weighted (DW) images that are acquired through the application of diffusion-sensitizing gradients along varying spatial directions. Diffusion Tensors (DTs) are then reconstructed from DW images under the assumption that water diffuses as an ellipsoid (Neeman et al., 1990; Mattiello et al., 1994). A tensor in this model is represented mathematically by a symmetric 3×3 matrix H , which for the biological diffusion of water must be positive definite (Basser et al., 1994a; Crank, 1975). The eigenvalues and eigenvectors of DTs encode the magnitude and direction, respectively, of the preferential diffusion of water molecules in each voxel of the imaging space. The eigenvector associated with the largest eigenvalue is called the Principal Direction (PD) of a tensor, which is assumed to be oriented along the longitudinal axis of the bundles of nerve fibers in white matter (Peled et al., 1998; Conturo et al., 1999; Mori et al., 1999; Basser and Pajevic, 2000). A quantitative analysis of the PD and other eigenvalues and eigenvectors thus allows investigators to track in vivo the trajectory of fiber pathways in the brain (Xu et al., 2002; Fillard and Gerig, 2003) and to compare those trajectories within and across groups of study participants (Ellis et al., 1999; Lim et al., 1999; Jones et al., 1999a,b; Skare et al., 2000; Pfefferbaum et al., 2000; Klingberg et al., 2000; Eriksson et al., 2001; Rugg-Gunn et al., 2001; Lim and Helpen, 2002; Plessen et al., 2006). In addition, analyses of the Fractional Anisotropy (FA) values computed from eigenvalues allows a quantitative analysis of the integrity of nerve fibers at particular voxels of the brain (Basser, 1995; Plessen et al., 2006). The validity and reliability of these quantitative analyses are limited, however, by the noise that is typically present in the eigenvalues and eigenvectors of the DTs reconstructed from DW images, which inherently contain varying amounts of noise (Westin and Knutsson, 2003; Martin-Fernandez et al., 2004; McGraw et al., 2004; Wang et al., 2004; Poupon et al., 1998). Noise in DW images, for example, introduces significant bias into the estimated eigenvalues, so that isotropic media may appear anisotropic, and vice-versa (Pierpaoli et al., 1996). Furthermore, radio-frequency (RF) noise can bias the mean and variance of a tensor's eigenvectors (Basser and Pajevic, 1999; Martin et al., 1999). Finally, eigenvalues can be calculated as negative in the presence of noise, which from a biological perspective is impossible. Thus, reducing the amount of noise in DTI datasets can increase the validity and reliability of DTI-based analyses of fiber pathways.

A number of advanced methods have been proposed to reduce noise in tensor fields. These methods can be grouped into two general categories: (1) those that smooth DW images and then reconstruct the DT field using those smoothed DW images (McGraw et al., 2004; Chen and Hsu, 2005; Ding et al., 2005; Parker et al., 2000; Tabelow et al., 2008); (2) those that smooth various elements of the tensor fields themselves, which are reconstructed from the original, unsmoothed DW images (Poupon et al., 1998, 2000; Tench et al., 2002; Westin et al., 2002; Martin-Fernandez et al., 2003, 2004; Tschumperle and Deriche, 2001, 2003; Coulon et al., 2004; Taylor and Noble, 2003; Wang et al., 2004; Chen and Hsu, 2005). Although all of these methods can be effective in reducing noise in a reconstructed DT field, they have several limitations. First, methods that smooth DW images using kernels matched to the structure of imaging data (Ding et al., 2005; Tabelow et al., 2008) generate optimally smoothed DW images, but these smoothed DW images do not ensure an optimally smoothed tensor field

because the DW images are nonlinearly related to the corresponding tensor fields. Second, most of the methods that directly smooth tensor fields typically use either a small number of properties of the tensors (most commonly, the PD and its associated eigenvalue) to smooth the field, or they smooth eigenvalues and eigenvectors independently of one another. Because both methods neglect important information in each tensor to smooth the tensor field, the eigenvalues and eigenvectors of the tensors in the smoothed field may not in fact vary smoothly across the field. Smoothing all eigenvalues and eigenvectors simultaneously across a tensor field is an important aim, because diffusion along the directions perpendicular to the PD may carry important information about the structure of white matter (Bansal et al., 2007; Song et al., 2002; Mori and Zhang, 2006). Third, the methods that use anatomical information (Ding et al., 2005; Tabelow et al., 2008) are efficient in smoothing DW images, but these methods may smooth boundaries across homogeneous regions of the brain by unknown amounts, thereby introducing errors in the fiber tracts that are reconstructed from the smoothed tensor fields. Finally, evaluation of the performance of these methods in real-world datasets has generally been limited by the application of simple, qualitative visual assessments of the degree to which they increase alignment of tensors in neighboring voxels of a single image.

We have developed a method that smoothes tensor fields by analyzing simultaneously all the eigenvalues and eigenvectors of tensors in neighboring voxels. Our method models a tensor field as a three-dimensional Markov Random Field (MRF) (Besag, 1986; Geman and Geman, 1984) and then estimates the locally optimal smoothing of the field using the method of simulated annealing (Metropolis et al., 1953). We assume that DTs in homogeneous regions of tissue are perturbed by small amounts of noise. We therefore use Perturbation theory (Sakurai and Fu Tuan, 1994; Bansal et al., 2008) to calculate the *likelihood term* and the *prior term* (also called the *data term* and the *smoothness term*, respectively) of the MRF model. Computation of these two terms involves calculation of the similarity of two DTs. Employing Perturbation theory, we calculate this similarity as the conditional probability $Pr(H_1|H_0)$ of observing the perturbed tensor H_1 , given the unperturbed tensor, H_0 (Bansal et al., 2008). Our method is expected to preserve boundaries across homogeneous regions of the brain, because we use line processes within the MRF model for a field to model these boundaries explicitly. The similarity between tensors across these boundaries is set equal to zero, thereby preserving boundaries between regions of differing tissue types. We quantitatively assess the improvement in smoothness of the eigenvalues and eigenvectors by calculating the Signal-to-Noise Ratio (SNR) for eigenvalues and the coherence of the associated eigenvectors across the DT fields. We generate Coherence Percent-Percent (CPP) plots from these measures of coherence for eigenvectors across the smoothed and unsmoothed DT fields and then quantify the performance of our method by computing the area under the curve in the CPP plots. Finally, we evaluate the performance of our method using both synthetic and real-world datasets.

These evaluations demonstrate that our method significantly improves the alignment of tensors and therefore increases the coherence of both the PD and SD (the secondary direction, or the direction of the eigenvector associated with the second-largest eigenvalue), and it increases the SNR of eigenvalues of tensors across the tensor field. Any of the numerous commercially available statistical tools (Martin et al., 1999; Fletcher and Joshi, 2004; Schwartzman et al., 2005; Arsigny et al., 2005; Schwartzman, 2006; Batchelor et al., 2005; Bansal et al., 2007; Song et al., 2002; Mori and Zhang, 2006) can be used to correlate various clinical and behavioral measures with the quantitative measures of diffusion derived from our smoothing procedures. Furthermore, the amount of smoothing that the procedure produces can be estimated directly from the smoothed field, which is needed when using the theory of Gaussian Random Fields to compute probability values that are corrected for multiple statistical comparisons (Friston et al., 1995).

2 Method

We first briefly describe the probabilistic method that we have developed previously for computing the similarity of two DTs. We then describe our approach for modeling a tensor field as a three-dimensional (3D) Markov Random Field (MRF) and for computing the various terms in this model using the probabilistic method. We then detail a method based on simulated annealing for estimating the smoothed DT field, and we develop procedures for evaluating quantitatively the performance of smoothing on synthetic and real-world datasets. Finally, we compare the performance of our method with one method that uses a Riemannian distance to compute the similarity between two tensors, and with another method that reduces noise in tensor fields by anisotropically filtering the DW images.

2.1 The Probabilistic Similarity of Two Tensors

The calculated smoothness of a tensor field depends critically on which method is used to compute the similarity of tensors in the field. A number of advanced methods have been proposed to compute the similarity of two tensors (Malvern, 1969; Pierpaoli et al., 1996; Alexander et al., 1999; Poupon et al., 1998; Pennec et al., 2006; Fletcher and Joshi, 2004; Moakher, 2005, 2006; Arsigny et al., 2005; Lu, 1998), including our own probabilistic method that is based on Perturbation theory (Sakurai and Fu Tuan, 1994; Anderson, 2001; Bansal et al., 2008). We use this probabilistically based computation of tensor similarity to estimate the smoothness of tensor fields. We model variations in tensors in a homogeneous region as perturbations of their neighbors, and apply the perturbation theory to compute the similarity of two tensors using the conditional probabilities of the eigenvalues and eigenvectors of a perturbed tensor, H , given the eigenspace components of the unperturbed tensor $H^{(0)}$. The probabilistic measure varies symmetrically and smoothly for smooth and symmetric changes in the tensors. Furthermore, the computed similarity relates directly to the diffusion properties of the water because it uses entire tensor to compute similarity values. For the observed tensor H , our conditional probability evaluates the likelihood of $H^{(0)}$ being the tensor in the absence of noise. This probabilistically based computation of similarity accounts inherently for the effects of noise on the measure across the tensor field.

The similarity between tensors could be computed using any one of several advanced methods that have been proposed previously. We have previously demonstrated, however, that our method for the similarity of tensors that is computed using Perturbation Theory varies smoothly and symmetrically when the eigenvalues and eigenvectors of tensors vary smoothly and symmetrically (Bansal et al., 2008). Variations in the log-Euclidean and Riemannian distance (Arsigny et al., 2005; Lu, 1998; Pennec et al., 2006; Fletcher and Joshi, 2004; Moakher, 2005, 2006), in contrast, are not symmetric for symmetric variations in eigenvalues and eigenvectors. Furthermore, we demonstrated that our method performs better than do the log-Euclidean and Riemannian metrics when applied to synthetic and real-subject data (Bansal et al., 2008). We therefore have incorporated this previously described method for computing tensor similarity in our newly proposed procedures for smoothing tensor fields.

We use this probabilistic method to compute the *smoothness term* and the *data term* in our 3D MRF model of a tensor field. Because the three eigenvectors of a DT are orthonormal, every DT is defined entirely by its three eigenvalues and only two eigenvectors. Therefore, the probability $Pr(H^{(0)} | H)$ of $H^{(0)}$ (unperturbed) given H (perturbed) is computed as follows (Bansal et al., 2008):

$$\begin{aligned} Pr(H^{(0)} | H) &= Pr(1^{(0)}, 2^{(0)}, E_1^{(0)}, E_2^{(0)}, E_3^{(0)} | 1, 2, E_1, E_2, E_3) \\ &= Pr(V_1^{(0)} | V_1) \cdot Pr(V_2^{(0)} | V_2) \cdot Pr(E_1^{(0)}, E_2^{(0)}, E_3^{(0)} | E_1, E_2, E_3); \end{aligned} \quad (1)$$

where the eigenvalues (E_1, E_2, E_3) and the eigenvectors (V_1, V_2, V_3) are assumed to be statistically independent, and where the conditional probabilities $Pr(V_n^{(0)}|V_n)$ of the eigenvectors are computed using Perturbation Theory (Bansal et al., 2008). Under the assumptions of Rician noise and heteroskedasticity conditions in the DW images, tensor elements are approximately Gaussian distributed (Bansal et al., 2008). The variance $\sigma_{\Delta_n^{(1)}}$ in the first order perturbation $\sigma_{\Delta_n^{(1)}}$ of an eigenvalue is then computed from the total amount of perturbation of the tensor and its corresponding eigenvector (Bansal et al., 2008). We then use the Gaussian distribution to model the distribution of the unperturbed eigenvalue associated with a given perturbed eigenvalue (Bansal et al., 2008). Eigenvalues are not Gaussian distributed, however, because they are nonlinearly related to the tensor elements. Nevertheless, our experiments using synthetic data and real-subject data demonstrated that our method performs well under the assumption that the eigenvalues are indeed Gaussian distributed (Bansal et al., 2008). In addition, we have previously demonstrated that under this assumption the conditional probability computes the similarity of two tensors that varies smoothly and symmetrically for smooth and symmetric variations in the tensors (Bansal et al., 2008).

The conditional probability $Pr(E_1^{(0)}, E_2^{(0)}, E_3^{(0)}|E_1, E_2, E_3)$ of the unperturbed eigenvalues, given the perturbed eigenvalues, is computed as:

$$\begin{aligned} Pr(E_1^{(0)}, E_2^{(0)}, E_3^{(0)}|E_1, E_2, E_3) &= Pr(E_1^{(0)}|E_1) \cdot Pr(E_2^{(0)}|E_2) \cdot Pr(E_3^{(0)}|E_3) \\ &= \prod_{n=1}^3 \frac{1}{\sigma_{\Delta_n^{(1)}} \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{\Delta_n^{(1)}}{\sigma_{\Delta_n^{(1)}}} \right)^2 \right], \end{aligned} \quad (2)$$

where the eigenvalues are assumed to be statistically independent and $\Delta_n^{(1)}$ is the difference of eigenvalues, which also is assumed to be a Gaussian-distributed random variable having a mean of zero and variance $\sigma_{\Delta_n^{(1)}}$ (Bansal et al., 2008).

2.2 Computing the Similarity of Two Tensors Using Riemannian Distance

We also implement a method that computes the similarity of two tensors H_0 and H_1 as their Riemannian distance $Riemann(H_0, H_1)$ (Pennec et al., 2006; Fletcher and Joshi, 2004; Moakher, 2005, 2006); i.e.

$$Riemann(H_0, H_1) = \| \log(H_0^{-1} H_1) \| = \sqrt{\sum_{i=1}^3 \ln^2 \lambda_i}$$

where λ_i are the eigenvalues of the matrix $H_0^{-1} H_1$. We use this distance within our framework of Markov Random Field (MRF) to generate the smoothed tensor fields, and quantitatively compare the performance of our method for generating the smoothed fields using either the conditional probability or the Riemannian distance to compute the similarity of two tensors. Because the Riemannian distance of two distinct tensors is always non-zero, we expect that using this metric compared with the use of ours based on Perturbation Theory will generate smoothed tensor fields with greater coherence for the principal direction (PDs) and the second principal direction (SD) across the field but will also produce more undesirable smoothing across boundaries of homogeneous regions within the smoothed field.

2.3 Modeling a Diffusion Tensor Field as a 3D MRF

We model a tensor field as a 3D MRF. Let Y denote the original tensor field, X denote the smoothed tensor field, and *line processes* L denote a set of random variables that model the horizontal and vertical edges between DTs in neighboring voxels of the field X , corresponding to discontinuities in orientations across homogeneous regions (Geman and Geman, 1984). Using line processes therefore allows us to generate an optimally smoothed field that preserves discontinuities between different tissue types. The appropriate smoothness that is expected *a priori* in a tensor field is described by *Cliques* C , or sets of locations such that each location in a set is a neighbor to all other locations of the set. In our MRF model, we use a *first-order* neighborhood system in which a tensor at a given location (i, j, k) is a neighbor to tensors at the locations $(i+1, j, k)$, $(i-1, j, k)$, $(i, j+1, k)$, $(i, j-1, k)$, $(i, j, k+1)$, and $(i, j, k-1)$. Using this neighborhood system, calculate the energy of a clique, e.g. $c = \{(i, j, k), (i+1, j, k)\}$, using DTs and line processes at those locations. Summing the energy over all cliques within the brain gives us the smoothness term $E_S(X, L)$, which quantifies the smoothness of the field X based on our *a priori* expectation of the smoothness in the DT field. Following the Hammersley-Clifford theorem (Geman and Geman, 1984), the prior probability $Pr(X, L)$ of observing the smoothed image X is therefore computed using a Gibbs distribution:

$$Pr(X, L) = \frac{1}{Z_S(T)} \exp \left[-\frac{1}{T} E_S(X, L) \right], q$$

where,

$$E_S(X, L) = \sum_{c \in \mathcal{C}} \sum_{u, v \in c, u \neq v} [Pr(X_u | X_v) (1 - l_{uv}) + \gamma l_{uv}],$$

γ is the significance of edges, $Z_S(T)$ is a normalization constant, and $Pr(X_u | X_v)$ is the probability of observing the perturbed tensor X_u given the tensor X_v computed from Perturbation theory (Bansal et al., 2008).

We quantify the *closeness* of the smoothed field X to the original field Y using the conditional probability $P(Y|X)$ of Y given X . This conditional probability is computed using a Gibbs distribution

$$Pr(Y|X) = \frac{1}{Z_D(T)} \exp \left[-\frac{1}{T} E_D(Y|X) \right],$$

where

$$E_D(Y|X) = \sum_u Pr(Y_u | X_u)$$

is called the *data term*, $Z_D(T)$ is a normalization constant, summation is carried over all voxels in the brain, and the conditional probability $Pr(Y_u | X_u)$ is computed using Perturbation Theory.

Finally, we combine the prior probability $Pr(X, L)$ and the conditional probability $Pr(Y|X)$ using Baye's theorem to compute the posterior probability $Pr(X, L|Y)$:

$$Pr(X, L|X) = \frac{pr(X, L) Pr(Y|X, L)}{Pr(Y)}. \quad (3)$$

The *maximum a-posteriori* (MAP) estimates $\hat{X}$ and $\hat{L}$ of the smoothed DT field and line processes, respectively by maximizing the natural logarithm of $Pr(X, L|Y)$:

$$\begin{aligned} (\hat{X}, \hat{L}) &= \arg\max_{(X, L)} \ln Pr(X, L|Y) \\ &= \arg\max_{(X, L)} [\ln Pr(Y|X) + \ln Pr(X, L)] \\ &= \arg\min_{(X, L)} [E_D(Y|X) + E_S(X, L)], \end{aligned}$$

where the second equation follows from the fact that $Pr(Y)$ is a constant. We scale the data term $E_D(Y|X)$ by a parameter η , which is varied between 0.0 and 1.0 in our experiments below, to study the effect of the data term on the smoothed image $\hat{X}$. Thus, we estimate the optimal DT field by minimizing the cost function:

$$\begin{aligned} (\hat{X}, \hat{L}) &= \arg\min_{(X, L)} [\eta E_D(Y|X) + E_S(X, L)] \\ &= \arg\min_{(X, L)} E(X, L, Y). \end{aligned} \quad (4)$$

2.4 Estimating the Smoothed Field Using Simulated Annealing

We estimate the optimally smoothed field $\hat{X}$ and line processes $\hat{L}$ using the method of simulated annealing (Hastings, 1970; Metropolis et al., 1953). Although simulated annealing is proven to converge to the global optima (Hastings, 1970; Metropolis et al., 1953), within limited computational time it may converge to a field that is only locally optimal. Simulated annealing is an iterative procedure that begins with an initial estimate for the smoothed field $x^0 = (X^0, L^0, Y)$ and an initial temperature T_0 , and then iterates a modification of the estimate to generate the final optimal estimate for the smoothed field. Briefly, the first step in modifying the current estimate $x^m = (X^m, L^m, Y)$ in the m th iteration is to generate a proposed new estimate $x^{(m+1)} = (X^{(m+1)}, L^{(m+1)}, Y)$ from that current estimate. The second step sets the current estimate as either x^m or $x^{(m+1)}$ (based on criteria detailed below), and it lowers the temperature T , thus generating a sequence $\{x^m, m = 0, 1, 2, \dots\}$ of MRFs that are sampled from the posterior distribution in Eqn. (3). Denoting $E(x^m) = \eta E_D(Y|X) + E_S(X, L)$, we write the posterior distribution $Pr(x^m)$ as

$$Pr(x^m) = \frac{1}{Z(T)} \exp\left(-\frac{E(x^m)}{T}\right).$$

A more complete description of these two iterated steps in simulated annealing is as follows:

1. First, a candidate state $x^{(m+1)}$ is generated from a proposal distribution $q(x^{(m+1)}|x^m)$ that depends upon the current state x^m . The proposal distribution is assumed in the Metropolis algorithm (Metropolis et al., 1953) to be symmetric, i.e., $q(x^{(m+1)}|x^m) = q(x^m|x^{(m+1)})$. In our implementation, a proposed $x^{(m+1)}$ is generated by perturbing the current state x^m by a small amount ϵ , i.e., $X^{(m+1)} = X^m + \epsilon$, as described below.
2. Second, the generated state $x^{(m+1)}$ is set as the new current state if its energy is less than that of the current state x^m . Otherwise, $x^{(m+1)}$ is set as the new current state with

probability $\alpha(x^m, x^{(m+1)})$, which is an exponential in the difference of their energies, declining asymptotically to zero for the decreasing temperature $T^{(m+1)}$:

$$\alpha(x^m, x^{(m+1)}) = \exp\left(\frac{E(x^{(m+1)}) - E(x^m)}{T^{(m+1)}}\right),$$

where $T^{(m+1)} = \delta T^m$ and δ is the rate at which the temperature declines in each iteration.

2.5 Details of Implementation

The smoothed tensor field is iteratively estimated by slowly decreasing the temperature T in each iteration. The rate δ at which the temperature T is lowered critically affects the rate at which the optimization procedure converges to an estimate $\hat{X}$. Furthermore, the optimal estimate $\hat{X}$ depends upon the perturbation scheme that generates a proposed estimate $x^{(m+1)}$ from the current estimate x^m . In our experiments, we set the initial temperature T_0 to the large value of 5000 and then reduced it gradually at a rate of $\delta = 0.975$ in each iteration, a rate that is sufficiently slow to avoid local minima in the initial iterations and yet sufficiently fast to converge to a specific field in later iterations. Our method iterates the two steps 500 times, at which point the initial temperature is reduced to the very small value of 0.0159. The temperature T is large in the initial iterations, and therefore any proposed tensor field $X^{(m+1)}$ is accepted as the current estimate. In contrast, the temperature is small in later iterations, and therefore only the tensor field $X^{(m+1)}$ that reduces the total energy is accepted as the current estimate in the next iteration. The current estimate X^m , therefore, converges to an estimate that is locally optimal. Our method generates the proposed field by perturbing tensors at all voxels, and it accepts the proposed field with the probability determined by the annealing schedule. By perturbing tensors at all voxels, our method converges to a local optimum in fewer iterations than do methods that perturb a tensor at only a single voxel to generate the proposed tensor field. Our experimental results suggest that this method of optimization generates a smoothed field with the desired amount of smoothness and while simultaneously preserving the boundaries between homogeneous tissues of the brain.

In the m th iteration of the optimization procedure, a proposed estimate $X^{(m+1)}$ is generated from the current estimate X^m by perturbing tensors at each voxel in a checkerboard pattern. For a tensor H_0 at location (i, j, k) , we first define a neighborhood of size β within slice k . We select at random a tensor H_r within this neighborhood and compute the similarity of the two tensors, $Pr(H_r|H_0)$. If $Pr(H_r|H_0) > 0$, we generate a perturbed tensor H_p at location (i, j) by adding a small fraction of H_r to H_0 , i.e., $H_p = 0.993 * H_0 + 0.007 * H_r$, where each element of the tensor H_r is multiplied by 0.007 and then is added to the corresponding element of the tensor H_0 to generate the tensor H_p . The perturbed tensor H_p is positive definite by construction. If H_p locally increases the smoothness in X_m , we replace H_0 by H_p . Otherwise, we generate a new H_p by selecting at random another tensor H_r in the neighborhood β . However, if H_p fails to increase the smoothness locally in 5 independent attempts, then the tensor H_0 is not perturbed, and we move to the next tensor in the checkerboard pattern. After we have perturbed all tensors to generate the proposed tensor field $X^{(m+1)}$, line processes $L^{(m+1)}$ are estimated from $X^{(m+1)}$ to reduce the computational time, as follows:

$$l_{uv}^{(m+1)} = \begin{cases} 1, & \text{if } pr(X_u|X_v) > 0; \\ 0, & \text{otherwise.} \end{cases}$$

We set the significance of the edges $\gamma = 0.5$, which generates fields with the desired degree of smoothing while preserving the boundaries across smoothed regions in the brain. The proposed

tensor field $X^{(m+1)}$ is set as the current field according to Step 2 of the procedure for simulated annealing. Furthermore, because voxels at the interface of gray matter and cerebro-spinal fluid (CFS) have fewer neighboring voxels that contain a valid DT, tensors at these voxels get perturbed less than do voxels within the brain, and they therefore will converge more slowly. However, these voxels typically have low fractional anisotropy (FA) values and therefore are of little interest in the applications of diffusion tensor imaging.

2.6 Assessing the Performance of Our Smoothing Algorithm

We evaluate the performance of our method visually and quantitatively using both a synthetically created tensor field and a tensor image acquired from a single lining individual. A tensor field is a multidimensional dataset that is characterized by six biologically important numbers at each voxel of the imaging space: three representing the three preferential directions of diffusion of water molecules, and three eigenvalues representing the magnitude of diffusion in each of these directions. We expect eigenvalues and eigenvectors within a homogeneous region of tissue to vary smoothly in the absence of noise. We therefore assess these local variations in the eigenvalues and eigenvectors across both the original and smoothed tensor fields to quantify the performance of our method both visually and quantitatively. In addition, we compare the performance of our method quantitatively with one method that computes the similarity between two tensors using the Riemannian distances within our MRF framework and another one that first uses an anisotropic filter to smooth the Diffusion Weighted (DW) images spatially and that then reconstructs a smooth tensor field from the smoothed DW images.

2.6.1 The Synthetic Tensor Field—We generated a synthetic tensor field consisting of 5 identical slices, each with a 32×32 matrix of DTs. In the first 16 rows of the matrix, the PDs of all DTs were aligned along the X-axis, and in the next 16 rows the PDs were aligned along the Y-axis. The PDs of tensors in the columns 8, 9, and 10 across all rows were aligned along the Z-axis (perpendicular to each 2D slice). All DTs were anisotropic with eigenvalues 0.5, 0.25, and 0.05, and fractional anisotropy (FA) equal to 0.7, a typical FA value for DTs within white matter of the brain in human adults (Madden et al., 2004). We set $\lambda \neq \lambda$ for two reasons. First, we specified tensors in our synthetic data with typical FA values of 0.7 in white matter of the brain. Setting $\lambda_2 \neq \lambda_3$ allows tensors to have the specified FA value and yet have λ_1 close to λ_2 , thereby ensuring that noise added to the data could change the PD and coherence of the tensors significantly. In contrast, setting $\lambda_2 = \lambda_3 = 0.25$ would require $\lambda_1 = 2.0$ for $FA = 0.7$, and therefore small perturbations in tensors would not significantly change the PD of the tensors, thereby limiting the utility and interest of the synthetic example. Second, we wanted to assess the performance of our method in improving the coherence of the SD, but the SD is undefined for tensors in which $\lambda_2 = \lambda_3$.

We perturbed each DT independently to generate synthetic DT fields with varying amounts of noise. We use the Cholesky method for decomposition (Press et al., 1992) to perturb each of the DTs, while assuming that the elements of each tensor are distributed according to the multivariate Gaussian distribution with a covariance matrix having an upper block-diagonal form (Basser and Pajevic, 2003; Batchelor et al., 2003)

$$\begin{pmatrix} 0.75\sigma^2 & 0.25\sigma^2 & 0.25\sigma^2 & 0 & 0 & 0 \\ 0.25\sigma^2 & 0.75\sigma^2 & 0.25\sigma^2 & 0 & 0 & 0 \\ 0.25\sigma^2 & 0.25\sigma^2 & 0.75\sigma^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma^2 \end{pmatrix}.$$

To ensure that the DTs are perturbed by only small amounts, we set $\sigma = 0.05$ and $\sigma = 0.65$. We then estimate the smoothed tensor fields using our algorithm with a neighborhood size $\beta = 5$ and a data weight $\eta = 0.5$. We visually and quantitatively assess the improvement in the coherence of the PDs across the smoothed tensor field. We expect that our algorithm will significantly increase the coherence of the DTs in the smoothed field while preserving the boundaries between homogeneous regions. We also expect that our method will smooth and preserve within the smoothed tensor field the DTs in columns 8, 9, and 10, where the PDs are aligned along the Z-axis.

2.6.2 A Real-World Dataset

Pulse Sequence: The DTI data for a single human subject was acquired on a GE Signa 3 Tesla whole body scanner (Milwaukee, WI) equipped with the pulse sequences and a 16-channel receiver for parallel imaging. All images were acquired with an 8-channel, phased array head coil. The parameters for the DTI pulse sequence were: Data matrix = 132×128 zero-padded to 256×256 , Repetition Time (TR) = 15700 ms, Echo Time (TE) = 74 ms, Field of View (FOV) = 24 cm, Flip = 90° , Slices = 60, Slice thickness = 2.5 mm, Spacing = 0 mm; Number of Averages = 2, Percent FOV = 1.0. Furthermore, 3 baseline images having $b=0$ s/mm² and 15 directions of diffusion gradients from a uniformly arranged direction scheme at $b=1000$ s/mm². Phase correction and the array spatial sensitivity encoding technique (ASSET) were applied to reduce susceptibility artifacts and spatial distortion, which also slightly reduced the image acquisition time to 9 minutes 26 seconds.

Imposing a Positivity Constraint on Tensor Estimations: To ensure that all reconstructed DTs were positive-definite, we masked out the background in the images and reconstructed only DTs within the brain. For the few reconstructed DT within the brain that were not positive-definite we inverted the sign of the negative eigenvalue and recomputed the tensor to satisfy the positivity constraint. Such an approach reduces computation time without compromising the imaging data or our findings. More advanced methods, including constrained nonlinear least-squares estimation (Koay et al., 2006; Mangin et al., 2002), could be used to ensure that the reconstructed DTs are positive-definite; however, these approaches are computationally expensive, and therefore we adopt the simpler approach to ensure the positive-definite property of tensors.

2.6.3 Assessing the Smoothness of the Eigenvectors—We quantitatively assess the smoothness of the eigenvectors in each tensor by computing the coherence of the eigenvector along the PD and the eigenvector associated with the SD. Because our image slices are 2.5 mm thick, we compute coherence measures $\mu^{PD}(i, j)$ and $\mu^{SD}(i, j)$ by comparing the eigenvectors of DTs at location (i, j) and its 24 first- and second-order neighbors within a slice (i.e. in a $4.5\text{mm} \times 4.5\text{mm}$ neighborhood within the slice):

$$\mu^{PD}(i, j) = \frac{1}{24} \sum_{(k,l)} PD_{(i,j)}, i-2 \leq k \leq i+2, \quad j-2 \leq l \leq j+2,$$

where $PD_{(i,j)}$ and $PD_{(k,l)}$ are the unit eigenvectors along the PDs at locations (i, j) and (k, l) , respectively, and summed over 24 DTs that neighbor the tensor at location (i, j) . A coherence measure varies from 0 to 1.0, and within homogeneous regions of tissue, coherence values are close to 1.0, the larger coherence measures signifying that tensors are better aligned across the tensor field. The coherence measure at a voxel can also be rewritten as the inner product of $PD_{(i,j)}$ at the voxel (i, j, k) and the average of PDs at its 24 neighbors. If the PDs of the neighbors are uniformly distributed along all directions, then the average PD equals zero and therefore the coherence equals zero. On the other hand, if the average PD equals $PD_{(i,j)}$, then the coherence equals 1. The coherence measure therefore reflects the amount of noise in a

homogeneous region of a given tissue type in the brain. Therefore, to quantify the performance of our algorithm in improving the coherence of eigenvectors, we generate a histogram of these coherence values. Because smoothing will increase the coherence of the eigenvectors, we expect to see a shift towards larger values in the histogram for the smoothed field. We also compute the percent increase in the number of DTs having coherence values greater than either 0.95 or 0.90 across the smoothed and original fields. We expect a significant increase in the number of DTs having coherences greater than 0.95 in the smoothed image.

Using the percent increase in DTs having coherences greater than specific values has two limitations, however. First, this quantitative measure of performance will depend upon the choice of the particular threshold value for the coherence measure. In our experiments, for example, the number of DTs with a PD coherence > 0.95 increased by 7.01% after smoothing, whereas the number of DTs having a coherence > 0.90 increased by only 0.47% (using a neighborhood size $\beta = 5$ and data weight $\eta = 0.7$, Fig. 14). The second limitation, related to the first, is that the number of tensors having a coherence greater than some value may actually decrease, leading to the erroneous conclusion that the method reduces coherence of the DTs across the smoothed field. The number of tensors having coherences greater than some predetermined value may decline because the computed coherence of DTs along the boundaries of the regions may decline as the DTs within regions become better aligned by smoothing. For PDs in our experiments, for example, the number of DTs having coherences > 0.90 did in fact decline in the smoothed field (neighborhood size $\beta = 5$ and data weight $\eta = 0.2$, Fig. 14). To avoid these erroneous conclusions, we propose the use of Coherence Percent-Percent (CPP) plots to quantify the overall improvement in coherence of DTs across the entire smoothed field.

Coherence Percent-Percent Plots: To generate a CPP plot, we compute the percentage of DTs having a coherence value greater than or equal to a specified value in both the smoothed and the original fields. The percent of DTs having coherence values greater than 1.0 equals zero, whereas 100% of the DTs have a coherence value greater than zero. We calculate a set of percent values for each tensor field by decreasing the coherence measure from 1.0 to 0.0, and then plot values in the smoothed field as fractions against values in the original field. The area under the curve in the graph is a unique number that quantifies unambiguously the improvement in coherence associated with smoothing compared with the coherence in the original field. This area equals 0.5 if the coherence across the smoothed field equals that across the original field. The area will be greater than 0.5 if the coherence is greater across the smoothed field, and it will equal 1.0 if the DTs are perfectly aligned in the smoothed field. A value less than 0.5 will signify that coherence is lower in the smoothed field than in the original. We expect our method will increase the coherence of eigenvectors substantially, and therefore we expect the area to be greater than 0.5 in the CPP plots for the PD and SD.

In addition to this quantitative analysis, we color-encode the PD at each voxel of the smoothed and original DT fields. These color maps allow visual assessment of the improvement in coherence that smoothing produces and of the preservation of boundaries between homogeneous regions across the tensor fields. Thus, both visual and quantitative assessments are helpful in assessing the performance of methods for noise reduction.

2.6.4 Assessing the Smoothness of Eigenvalues—To assess visually the effect of smoothing on eigenvalues, we generate a map of FA values at each voxel of the smoothed and original fields, as well as a difference map of FA values in the two fields. We expect that these maps will demonstrate that our algorithm spatially smoothes FA values while preserving the sharp boundaries between tissue types, and that the FA map for the smoothed field is close to that of the original field.

To assess quantitatively the increased in the spatial smoothness of the eigenvalues, we compute their SNRs, thereby quantifying the spatial variation in eigenvalues for the largest and second-largest eigenvalues at each voxel in a small neighborhood across each of the two fields. Larger amounts of noise will perturb eigenvalues by larger amounts in this neighborhood, thereby reducing the computed SNR. Increasing the spatial smoothness of eigenvalues will reduce variations in eigenvalues, and therefore will increase SNR. To compute the SNR for an eigenvalue at a voxel (i, j, k) , we first calculate the average and variance of the eigenvalues in a $7 \times 7 \times 7$ neighborhood of that voxel. The SNR at that voxel is then estimated as the ratio of the average value to the variance. Then for both the largest and the second-largest eigenvalues, we generate a histogram of the SNR at each voxel by plotting SNR values along the X-axis and the *Number of Voxels* along the Y-axis. In assessing the smoothing that our algorithm produces, we expect a shift toward higher SNR values in the histogram for the smoothed tensor field compared with the histogram for the original field, thereby documenting and quantifying the increase in SNR in eigenvalues across the smoothed field. In addition, we compute the percent increase in the number of voxels having an SNR greater than or equal to 8.0 across the two fields, as well as the mean and median SNR in the smoothed and original fields, all of which we expect to increase significantly in the smoothed field compared to the original field.

2.6.5 Optimizing Parameters—The performance of our method depends on the size β of the neighborhood used to generate a proposed estimate $X^{(m+1)}$ and on the weighing term η used to scale the data term $E_D(Y/X)$ in Eqn. (4). We expect that increasing values of β and decreasing values of η will produce increasing degrees of smoothing. To study the effects of these two terms, we increase β from 3 to 11 (while holding constant $\eta = 0.5$) in increments of 2 for one set of experiments, and we increase η from 0.0 to 1.0 (while holding constant $\beta = 5$) in increments of 0.1 for the second set of experiments. These analyses help us to determine optimal values for these two parameters, and we used those values in quantitative analyses of our algorithm's performance.

2.6.6 Anisotropic Filtering of Diffusion Weighted Images—We compare the performance of our method with that of a method that reconstructs a DT field after smoothing the diffusion weighted (DW) images with an anisotropic filter. Our implementation of anisotropic filtering is similar to several excellent methods for anisotropic filtering proposed in the literature (Gerig et al., 1992; Ding et al., 2005; Lee et al., 2005; Tabelow et al., 2008). The amount of smoothing varies with the number of iterations of the filter. Using the default settings in the ANALYZE software platform (Mayo foundation, Mayo clinic, Rochester, MN), we iterate the filter either 1 or 10 times, and then we reconstruct DT fields from the smoothed DW images. We apply these filters using the default settings in the Analyze software using the exponential diffusion function to model the discontinuities and the parameters $\kappa = 0.8$ and time interval = 0.25, which were determined empirically to work well for a wide range of images. To assess the performance of this method visually, we generate color maps that encode the PD of DTs across the smoothed fields. We also generate maps that use grey-scale values to encode FA measures. We assess the performance of this method quantitatively by computing the area under the CPP curve for the PD and SD and by plotting the histogram for the calculated SNRs of the largest and second-largest eigenvalues across the smoothed DT field. We expect that increasing the number of iterations will increase the coherence of the PD and SD, as well as increase the SNR of the eigenvalues, across the smoothed DT field. We also expect, however, that increasing the number of iterations will reduce discrimination of the boundaries between homogeneous tissues.

3 Results

We first visually and quantitatively assess the increase in coherence of eigenvectors and the increase in smoothness of eigenvalues using both synthetic and real-world datasets. We set the

data weight $\eta = 0.5$ and the neighborhood size $\beta = 5$, which we determined to be optimal in our experiments described below (Section 3.4) to be optimal. We quantitatively assess the effect of the data weight η and the neighborhood size β on the smoothed tensor field. Finally, we quantitatively assess the performance of a method that anisotropically filters DW images to smooth tensor fields, and we compare its performance with that of our own method.

3.1 Smoothing Synthetic Data

Visually, our algorithm for noise reduction seems to preserve well the boundaries between homogeneous regions, which have not shifted at all from their true locations in the smoothed DT field (Fig. 1). The algorithm even preserves alignment of the three-voxel-wide region in which PDs are directed along the Z-axis (through plane) (Fig. 1). Quantitatively, smoothing and noise reduction increases significantly the coherence of both the PD and the SD of tensors across the smoothed synthetic field (Fig. 2). Thus, our method significantly reduces noise without smoothing boundaries between regions, and it preserves and smoothes even small homogeneous regions in the synthetic tensor fields that are perturbed by varying amounts of noise.

3.2 The Effects of Smoothing on Eigenvectors

Visually, the coherence maps of the PD and SD show that the tensors across the smoothed field are better aligned compared with the tensors in the original field (Figs. 3 & 5). In addition, the color map encoding the PD shows that DTs are aligned better within homogeneous regions, and the sharp boundaries between those regions are well preserved, when compared with those in the original image. In addition, the PDs in the smoothed DT field are visually similar to the PDs in the original DT field (Fig. 4). Quantitative analysis of the coherence values shows that the peak in the histogram of the coherence measures shifted toward 1, indicating that more DTs in the smoothed field have higher coherence values than do those in the original image (Fig. 6). The number of DTs having coherence values for the PD greater than 0.95 is 7.78%, greater in the smoothed field than in the original field and for the SD is 40.9% greater in the smoothed field than in the original field (Fig. 6). Furthermore, the areas under the graph in the CPP plots for both the PD and the SD are greater than 0.5, indicating that the DTs are aligned better in both directions in the smoothed than in the original field (Fig. 7). Thus, our method significantly improves the coherence of DTs, it preserves the sharp boundaries between homogeneous tissues in the smoothed field, and it keeps the orientations of DTs in the smoothed field close to those in the original.

3.3 The Effect of Smoothing on Eigenvalues

The maps of Fractional Anisotropy (FA) values show that smoothing significantly reduces the spatial variation in the FA values (Figs. 8 & 9). Moreover, our algorithm preserves the boundaries between homogeneous tissues, without displacement. The maps and histograms of SNR for the largest and the second largest eigenvalues show that our algorithm significantly increases the SNR of eigenvalues in the smoothed field (Figs. 10 & 11). The number of DTs with SNR greater than 8 in the smoothed field increased by 108.09% for the largest eigenvalue and by 342.29% for the second largest eigenvalue. Our method also preserves the fine, hair-like structures in the posterior regions of the brain in the smoothed DT field (Fig. 8).

3.4 Parameter Effects on Optimization

Increasing the neighborhood size β and the data weight η has little influence on the visual appearance of the PD and FA maps. In addition, even very small values of η preserve the sharp boundaries between homogeneous regions in those maps, and neither map for the smoothed dataset is displaced from the corresponding map for the original field (Fig. 12), indicating that our algorithm is robust to the choice of these two parameters. Nevertheless, the graph of the

area under the Coherence Percent-Percent (CPP) plot suggests that a neighborhood size $\beta = 5$ produces an optimal alignment of DTs (Fig. 13).

The coherence of the PDs and SDs across the smoothed field increases monotonically with a decreasing data weight η (Fig. 13). Moreover, as expected, the number of DTs having a coherence value greater than 0.95 or 0.90 for the SD also increases as the data weight decreases (Fig. 14). Although the number of DTs having a coherence value for the PD greater than 0.95 also increases monotonically as the data weight decreases, the number of DTs having coherences greater than 0.90 actually decreased as the data weight decreases (Fig. 14), primarily because coherence values for the PD in the original field are high (Fig. 6). Therefore, even though additional smoothing had little influence on increasing the coherence values of DTs within homogeneous tissues, the coherence of DTs along the boundaries of these regions declines, thereby decreasing the number of DTs having coherence values ≥ 0.90 . Thus, because the number of DTs having coherence values ≥ 0.95 or 0.90 increases for a data weight $\eta = 0.5$, and because the area under the graph in the CPP is also large; therefore, we select $\eta = 0.5$ as the optimal estimate of this parameter in our experiments.

3.5 Anisotropic Filtering of Diffusion Weighted Images

Increasing the number of applications of the anisotropic filter to the DW images visually increases the coherence of the PDs and SDs (Fig. 15), and it increases the smoothness of their eigenvalues (Fig. 15). However, increasing the number of iterations also progressively increases the smoothing of the boundaries between homogeneous tissues (Fig. 15). The area under the CPP plot for the PDs increased from 0.545 to 0.755 when the number of iterations increased from 1 to 10 (Fig. 16). Similarly, the number of DTs having SNRs > 8 for the largest eigenvalue increased from 10,497 to 31,348 when the iterations increased from 1 to 10 (Fig. 17). Thus after one iteration, anisotropic filtering produced a smaller reduction in noise than did our method. After 10 iterations, however, anisotropic filtering produced a greater reduction in noise than did our method, but at the substantial cost of smoothing the boundaries between homogeneous tissues in the smoothed DT field. Moreover, large amounts of smoothing can change the orientations of fiber pathways unpredictably. Therefore, even though anisotropic filtering for some number of iterations between 1 and 10 will perform as well as our method in terms of reducing noise, it also invariably smoothes the boundaries between homogeneous regions. Anisotropic filtering also does not constrain displacement of the smoothed field. Our algorithm, in contrast, reduces noise in the DT field while preserving boundaries between tissues and, through the *data term*, it ensures that the smoothed tensor field is not displaced from the original field. In contrast, iterative application of anisotropic filtering increases the smoothness of the field and in limit will generate a field without boundaries.

3.6 Smoothing When Using Riemannian Distance

As expected, in the smoothed fields generated using Riemannian distance, the PDs and the SDs are more coherent than are those in the smoothed field generated using conditional probability (Figs. 18 & 7). Therefore, use of Riemannian distance generates fields with a greater degree of smoothness, and a larger reduction of noise in the tensor field, than does our algorithm. Similar to anisotropic filtering, however, Riemannian distance also produces a greater blurring of the boundaries between homogeneous tissues in the brain (Fig. 20).

To assess the performance of our method for preserving the boundaries between homogeneous regions, we plot the edges map across the original and the smoothed tensor fields (Figs. 1 & 19). These maps demonstrate that our algorithm significantly reduces edges within homogeneous regions while preserving boundaries between these regions. Our method also preserves the narrow regions in the brain and the edges along narrow fiber tracts, especially in the posterior regions of the brain (arrows in Fig. 19). Thus, edge processes within our

framework of MRF using conditional probability models discontinuities within tensor fields with excellent accuracy and generates smoothed tensor fields with an appropriate trade-off of smoothing the tensor field while preserving the boundaries between homogeneous tissues.

4 Discussion

We have developed an algorithm that uses Perturbation theory to smooth a tensor field. We model that field as a 3D Markov Random Field (MRF) using a *smoothness term*, which describes quantitatively the expected smoothness of the tensor field and a *data term*, which describes the fidelity of the location and orientation of tensors in the smoothed field relative to those in the original field. A weighted sum of the two terms defines a cost function that is optimized, through simulated annealing, to produce a smoothed DT field having the desired degree of smoothing. By varying the importance of the data term compared with the smoothness term in the cost function, we can vary the degree of smoothing in the estimated tensor field while ensuring the fidelity of the location and orientation of tensors. Our quantitative analyses of eigenvalues and eigenvectors demonstrate that homogeneous tissues in the smoothed tensor field are indeed smooth, while the distinctness and locations of the boundaries between those tissues are well preserved.

We have proposed a method within the framework of MRF that computes the similarity of two tensors as a conditional probability. We quantitatively and qualitatively evaluated the performance of our method using synthetic and real-world datasets. Because in general any method that reduces noise in a tensor field will generate a smoother tensor field, we have in addition compared the performance of our method with the performance of a method that applies anisotropic filtering to the diffusion weighted images (**Section 2.6.6**), and with the performance of another method that uses Riemannian distance within our framework of MRF to compute the similarity of two tensors (**Section 2.2**). This comparison demonstrated that these latter two methods may produce smoothed tensor fields with greater smoothness across a tensor field (Figs. 16 & 18)), but both methods also produced an unacceptable degree of smoothing across boundaries between homogeneous tissues of the brain, tending to blur these boundaries and to impair the discrimination of differing types of tissue (Figs. 19, 20, & 15). Our algorithm, in contrast, generated smoothed fields that had a more acceptable trade-off between the degree of smoothing that it produces and the preservation of boundaries between homogeneous tissues, while also remaining close to the structure of the original fields. Moreover, our experiments demonstrated that our algorithm is robust to variation in its various operating parameters over a large range of values.

Our algorithm also significantly increased the coherence of the second principal direction (SD) of the tensors in the smoothed field. The SD is not defined when the tensor is prolate in shape and its significance is unclear in regions devoid of fiber crossings. In the homogeneous tissues where the tensor is neither prolate nor oblate, however we expect high coherence for SDs across tensors in a local neighborhood. Furthermore, analyses of diffusion along the SD may yield important information about the myelination of axons or the density of fiber tracts (Bansal et al., 2007; Song et al., 2002; Mori and Zhang, 2006). For a specified amount of noise in the tensor field, the SNRs for the SD and its associated eigenvalue are lower than are those for the PD and its associated eigenvalue. Therefore, methods for noise reduction ideally should smooth the SDs and their associated eigenvalues by larger amounts than the PDs and their eigenvalues, and our algorithm does this. Thus, in addition to increasing the coherences of the PDs and the SNRs of their associated eigenvalues, our algorithm significantly increases the coherences of the SDs (Figs. 3 & 6) and improves the SNRs of their associated eigenvalues (Fig. 11). The proportionately larger increase in the SNRs for both the SDs and their associated eigenvalues, compared with the SNRs for the PDs, reflects the lower SNR of the SDs in the original field.

We use CPP plots to quantify the increase in coherence of DTs that smoothing produces relative to the original field. A diffusion tensor is a multidimensional datum that comprises 6 numbers (3 eigenvalues and 3 eigenvectors) at each voxel of the image. To quantify the increased smoothness of the eigenvectors that an algorithm for noise reduction produces, we compute the coherences of the PDs and SDs across the smoothed and original tensor fields. We then plot the percent of DTs having a coherence greater than a specified value in the smoothed image against the corresponding percentage in the original image. The area under the graph in this plot provides a quantitative index of the improvement in coherence of DTs that smoothing and noise reduction produce. We also note the percent increase in DTs having a coherence value greater than 95% in the smoothed field compared with the original field. Our noise reduction algorithm significantly increases the coherence of both PDs and SDs.

The coherence measure for tensors along tissue boundaries is small, and indeed optimal smoothing may actually reduce the coherence of tensors along the boundaries of homogeneous tissues. This potential reduction in coherence along tissue boundaries means that a single value of the coherence measure cannot be used to quantify the performance of our method or any method that is used to smooth tensor fields. This difficulty motivated our use of the area under the CPP plot to quantify the smoothing of tensor fields. Although this area would equal 1.0 for a completely smoothed and uniform field, our MRF-based model for smoothing explicitly models boundaries between regions of homogeneous tissues using line processes, and therefore it will generate a smoothed field without smoothing its boundaries excessively. Its area under the CPP plot will in practice almost never approach 1.0. In addition to use of quantitative measures to assess smoothing, simple visual inspection of the smoothed field is also vitally important in assessing the quality of smoothing. We therefore applied quantitative and visual procedures to assess the performance of the various algorithms.

Our algorithm for smoothing a tensor field uses Perturbation theory to compute the *smoothness* and *data terms* that generate a proposed state $X^{(m+1)}$ from the current state X^m in a model of the tensor field that is based on a 3D Markov Random Field. We have previously developed a method based on Perturbation theory to compute the similarity of two DTs as the conditional probability $Pr(H_1|H_0)$ of observing a perturbed tensor H_1 given an unperturbed tensor H_0 (Bansal et al., 2008). This probabilistic similarity $Pr(H_1|H_0)$ varies between 0 and 1 for DTs within homogeneous tissues, equaling 1 for identical tensors and 0 for tensors that differ markedly, such as those from different tissues in the brain. Having a similarity that equals 0 is the key property that preserves boundaries between homogeneous tissues and that preserves fine, hair-like structures through the white matter of the brain (Fig. 8). Because the power spectrum of white noise is nonzero for all frequencies, and because physiological signals are typically located in the lower range of the frequency spectrum, reducing noise amounts essentially to applying a low-pass filter to the dataset. Typically, low-pass filtering is implemented as an averaging of data across neighboring voxels. The indiscriminate averaging of values across neighboring voxels using a rotationally symmetric Gaussian filter, however, would smooth the boundaries between homogeneous tissues. Smoothing with an anisotropic filter could in principle alleviate this problem somewhat, but it cannot do so completely (Fig. 15). Our method, in contrast, preserves the distinctness of boundaries between tissues because we generate the new estimate $X^{(m+1)}$ of the optimally smoothed image by perturbing a tensor H_0 in the X^m using a tensor H_r in its neighborhood, such that the similarity of H_r and H_0 is greater than 0. Therefore, the morphology of the perturbed tensor H_p is similar to that of the tensor H_0 . Because DTs in the proposed field $X^{(m+1)}$ are close at each step to the current estimated field X^m , the boundaries between homogeneous tissues and the fine structures within white matter are preserved in the smoothed field (Fig. 8). Thus, the careful application of a probabilistic measure of similarity, which reduces to zero for two markedly different tensors, ensures the generation of a smoothed tensor field that has the desirable features of being similar

to the original field along tissue boundaries and yet smooth within tissues that are relatively homogeneous.

Diffusion tensors within the brain are always positive definite, because eigenvalues encode along various directions the magnitude of the diffusion of water molecules and these are always positive values. Some of the methods that have been proposed previously for smoothing tensor fields may destroy this important property of a tensor. To ensure the positivity of DTs, these methods therefore often explicitly impose a positivity constraint in their formulations. This constraint will of course ensure that the DTs are positive definite, but the eigenvalues in the smoothed field can still be arbitrarily close to zero, which is biologically improbable. Our method ensures that DTs are positive definite because the generation of a new estimate for the tensor field in the simulated annealing procedure adds small perturbations to the current best estimate. Furthermore, the *data term* of the MRF model ensures that the smoothed tensor field is similar to the original field, thereby ensuring that the eigenvalues will not be arbitrarily close to zero. Thus, our method generates smoothed fields of tensors that are biologically relevant, and it ensures that the tensors are positive definite across the smoothed field.

Acknowledgments

This work was supported in part by NIMH grants MH068318 and K02-74677, NIDA grant DA017820, and the Suzanne Crosby Murphy Endowment at Columbia University.

References

- Alexander, D.; Gee, J.; Bajcsy, R. Similarity measures for matching diffusion tensor images; British Machine Vision Conference (BMVC'99); 1999.
- Anderson AW. Theoretical analysis of the effects of noise on diffusion tensor imaging. *Magnetic Resonance in Medicine* 2001;46:1174–1188. [PubMed: 11746585]
- Arsigny, V.; Fillard, P.; Pennec, X.; Ayache, N.; Duncan, J.; Gerig, G. Fast and simple calculus on tensors in the log-Euclidean framework; Proceedings of MICCAI 2005, Eighth International Conference on Medical Image Computing and Computer-Assisted Intervention (LNCS 3749); Palm Springs, CA, USA: 2005. p. 115-122.
- Bansal R, Staib L, Plessen K, Xu D, Royal J, Peterson B. Voxel-wise comparisons in the morphology of diffusion tensors across groups of experimental subjects. *Psychiatric Research, Neuroimaging*. 2007In Press
- Bansal R, Staib L, Xu D, Laine A, Royal J, Peterson B. Using perturbation theory to compute the morphological similarity of diffusion tensors. *IEEE Transactions on Medical Imaging* 2008;27(5): 589–607. [PubMed: 18450533]
- Basser P. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *NMR in Biomedicine* 1995;8(7):333–344. [PubMed: 8739270]
- Basser P, Mattiello J, LeBihan D. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J. Magn. Reson. B* 1994a;103:247–254. [PubMed: 8019776]
- Basser P, Mattiello J, LeBihan D. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* 1994b;66:259–267. [PubMed: 8130344]
- Basser, P.; Pajevic, S. 7th Annual ISMRM. Philadelphia: 1999. Method to reduce eigenvalue sorting bias in DT-MRI; p. 1788
- Basser P, Pajevic S. Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise. *Magnetic Resonance in Medicine* 2000;44(1):41–50. [PubMed: 10893520]
- Basser P, Pajevic S. A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI. *IEEE Transactions on Medical Imaging* 2003;22(7):785–794. [PubMed: 12906233]
- Batchelor P, Atkinson D, Hill D, Calamante F, Connelly A. Anisotropic noise propagation in diffusion tensor MRI sampling schemes. *Magnetic Resonance in Medicine* 2003;49:1143–1151. [PubMed: 12768593]

- Batchelor P, Moakher M, Atkinson D, Calamante F, Connelly A. A rigorous framework for diffusion tensor calculus. *Magnetic Resonance in Medicine* 2005;53(1):221–225. [PubMed: 15690523]
- Besag J. On the statistical analysis of dirty pictures. *J. of the Royal Statistical Society, B* 1986;48(3):259–302.
- Chen B, Hsu E. Noise removal in magnetic resonance diffusion tensor imaging. *Magnetic Resonance in Medicine* 2005;54:393–407. [PubMed: 16032670]
- Conturo T, Lori N, Cull T, Akbudak E, Snyder A, Shimnoy J, McKinstry R, Burton H, Raichle M. Tracking neuronal fiber pathways in the living human brain. *Proc. Natl. Acad. Sci. USA* 1999;96:10422–10427. [PubMed: 10468624]
- Coulon O, Alexander D, Arridge S. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis* 2004;8(1):47–67. [PubMed: 14644146]
- Crank, J. *The Mathematics of Diffusion*. Oxford University Press; 1975.
- Ding Z, Gore J, Anderson A. Reduction of noise in diffusion tensor images using anisotropic smoothing. *Magnetic Resonance in Medicine* 2005;53:485–490. [PubMed: 15678537]
- Ellis C, Simmons A, Jones DK, Bland J, Dawson JM, Horsfield MA, Williams S, Leigh P. Diffusion tensor MRI assesses corticospinal tract damage in ALS. *Neurology* 1999;53:1051–1058. [PubMed: 10496265]
- Eriksson SH, Rugg-Gunn FJ, Symms MR, Baker GJ, Duncan JS. Diffusion tensor imaging in patients with epilepsy and malformations of cortical development. *Brain* 2001;124:617–626. [PubMed: 11222460]
- Fillard P, Gerig G. Analysis tool for diffusion tensor MRI. *MICCAI'03*. 2003
- Fletcher P, Joshi S. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. *Workshop on Computer Vision Approaches to Medical Image Analysis (CVAMIA) 2004*;3117:87–98. LNCS
- Friston K, Holmes A, Worsley K, Poline J, Frith C, Frackowiak R. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping* 1995;2:189–210.
- Geman S, Geman D. Stochastic relaxation, gibbs distributions and the bayesian restoration of images. *IEEE Trans. PAMI* 1984;6(6):721–741.
- Gerig G, Kubler O, Kikinis R, Jolesz F. Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging* 1992;11(2):221–232. [PubMed: 18218376]
- Hastings WK. Monte carlo sampling methods using markov chains and their applications. *Biometrika* 1970;57(1):97–109.
- Jones D, Horsfield M, Simmons A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn. Reson. Med* 1999a;42(3):515–525. [PubMed: 10467296]
- Jones D, Simmons A, Williams S, Horsfield M. Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magn. Reson. Med* 1999b;42(1):37–41. [PubMed: 10398948]
- Klingberg T, Hedehus M, Temple E, Salz T, Gabrieli J, Moseley M, Poldrack R. Microstructure of temporo-parietal white matter as a basis for reading ability: Evidence from diffusion tensor magnetic resonance imaging. *Neuron* 2000;25(2):493–500. [PubMed: 10719902]
- Koay C, Carew J, Alexander A, Basser P, Meyerand M. Investigation of anomalous estimates of tensor-derived quantities in diffusion tensor imaging. *Magnetic Resonance in Medicine* 2006;55:930–936. [PubMed: 16526013]
- Lee J, Chung M, Oakes T, Alexander A. Anisotropic kernel smoothing of DTI data. *International Society for Magnetic Resonance in Medicine (ISMRM)*. 2005
- Lim K, Hedehus M, Moseley M, de Crespigny A, Sullivan E, Pfefferbaum A. Compromised white matter tract integrity in schizophrenia inferred from diffusion tensor imaging. *Arch. Gen. Psychiatry* 1999;56:367–374. [PubMed: 10197834]
- Lim K, Helpert J. Neuropsychiatric applications of DTI - a review. *NMR in biomedicine* 2002;15(2):587–593. [PubMed: 12489105]
- Lu Y. Computing the logarithm of a symmetric positive definite matrix. *Applied numerical mathematics* 1998;26(4):483–496.

- Madden D, Whiting W, Huettel S, White L, MacFall J, Provenzale J. Diffusion tensor imaging of adult age differences in cerebral white matter: relation to response time. *NeuroImage* 2004;21:1174–1181. [PubMed: 15006684]
- Malvern, LE. Introduction to the mechanics of a continuous medium. Prentice–Hall, Inc.; Englewood Cliffs, N.J.: 1969.
- Mangin J, Poupon C, Clark C, Bihan DL, Bloch I. Distortion correction and robust tensor estimation for MR diffusion imaging. *Medical Image Analysis* 2002;6:191–198. [PubMed: 12270226]
- Martin K, Papadakis N, Huang C, Hall LD, Carpenter T. The reduction of the sorting bias in the eigenvalues of the diffusion tensor. *Magnetic Resonance in Imaging* 1999;17(6):893–901.
- Martin-Fernandez M, Alberola-Lopez C, Ruiz-Alzola J, Westin C-F. Regularization of diffusion tensor maps using a non-Gaussian Markov random field approach. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 2003;2879:92–100.LNCS
- Martin-Fernandez M, Westin C, Alberola-Lopez C. 3D Bayesian regularization of diffusion tensorMRI using multivariate gaussian markov random fields. *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* 2004;3216:351–359.LNCS
- Mattiello J, Basser P, LeBihan D. Analytical expression for the b matrix in NMR diffusion imaging and spectroscopy. *J. Magn. Reson. A* 1994;108:131–141.
- McGraw T, Vemuri B, Chen Y, Rao M, Mareci T. DT-MRI denoising and neuronal fiber tracking. *Medical Image Analysis* 2004;8:95–111. [PubMed: 15063860]
- Melhem ER, Mori S, Mukundan G, Kraut MA, Pomper MG, van Zijl PCM. Diffusion tensor MR imaging of the brain and white matter tractography. *American Journal Of Radiology (AJR)* 2002;178:3–16.
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equation of state calculations by fast computing machines. *J. Chem. Phys* 1953;21:1087–1092.
- Moakher M. A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications* 2005;26(3):735–747.
- Moakher M. On the averaging of symmetric positive-definite tensors. *Journal of Elasticity* 2006;82:273–296.
- Mori S, Crain B, Chacko V, Van Zijl P. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 1999;45(2):265–269. [PubMed: 9989633]
- Mori S, Zhang J. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron* 2006;51:527–539. [PubMed: 16950152]
- Neeman M, Freyer J, Sillerud L. Pulsed–gradient spin–echo studies in NMR imaging. effects of the imaging gradients on the determination of diffusion coefficients. *J. Magn. Reson* 1990;90:303–213.
- Parker P, Schnabel J, Symms M, D.J. W, Barker G. Nonlinear smoothing for reduction of systematic and random errors in diffusion tensor imaging. *J. of Magnetic Resonance Imaging* 2000;11:702–710.
- Peled S, Gudbjartsson H, Westin C, Kikinis R, Jolesz F. Magnetic resonance imaging shows orientation and asymmetry of white matter fiber tracts. *Brain Res* 1998;780:27–33. [PubMed: 9473573]
- Pennec X, Fillard P, Ayache N. A Riemannian framework for tensor computing. *International Journal of Computer Vision* 2006;66(1):41–66.
- Pfefferbaum A, Sullivan E, Hedehus M, Lim K, Adalsteinsson E, Moseley M. Age-related decline in brain white matter anisotropy measured with spatially corrected echo-planar diffusion tensor imaging. *Magnetic resonance in medicine* 2000;44(2):259–268. [PubMed: 10918325]
- Pierpaoli C, Jezzard P, Basser P, Barnett A, Chiro G. Diffusion tensor MR imaging of human brain. *Radiology* 1996;201(3):637–648. [PubMed: 8939209]
- Plessen KJ, Gruner R, Lundervold A, Hirsch J, Xu D, Bansal R, Hammer T, Wentzel-Larsen T, Atle LS, Peterson B, Hugdahl K. Reduced white matter connectivity in the corpus callosum of children with Tourette syndrome. *Journal of Child Psychology and Psychiatry* 2006;47(12):1013–1022. [PubMed: 17073980]
- Poupon C, Clark A, Frouin V, Regis J, Bihan DL, Bloch I, Mangin JF. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage* 2000;12:184–195. [PubMed: 10913324]

- Poupon C, Mangin J-F, Frouin V, Regis J, Poupon F, Pachot-Clouard M, Bihan DL, Bloch I. Regularization of MR diffusion tensor maps for tracking brain white matter bundles. *MICCAI'98* 1998;1496:489–498.LNCS
- Press, W.; Teukolsky, S.; Vetterling, W.; Flannery, B. *Numerical Recipes in C. The Art of Scientific Computing*. Cambridge University Press; 1992.
- Rugg-Gunn FJ, Eriksson SH, Symms MR, Barker GJ, Duncan JS. Diffusion tensor imaging of cryptogenic and acquired partial epilepsies. *Brain* 2001;124:627–636. [PubMed: 11222461]
- Sakurai, JJ.; Fu Tuan, S. *Modern Quantum Mechanics*. Addison–Wesley; 1994. ISBN: 0–201–53929–2
- Schwartzman, A. *Random Ellipsoids and False Discovery Rates: Statistics for Diffusion Tensor Imaging Data*. Stanford University; 2006. PhD thesis
- Schwartzman A, Dougherty R, Taylor J. Cross-subject comparison of principal diffusion direction maps. *Magnetic Resonance in Medicine* 2005;53:1423–1431. [PubMed: 15906307]
- Skare S, Hedehus M, Moseley M, Li T. Condition number as a measure of noise performance of diffusion tensor data acquisition schemes with MRI. *J. Magn. Reson* 2000;147(2):340–352. [PubMed: 11097823]
- Song S, Sun S, Ramsbottom M, Chang C, Russell J, Cross A. Dysmyelination revealed through MRI as an increased radial (but unchanged axial) diffusion of water. *Neuroimage* 2002;17:1429–1436. [PubMed: 12414282]
- Tabelow K, Polzehl J, Spokoiny V, Voss H. Diffusion tensor imaging: Structural adaptive smoothing, *NeuroImage*. 2008
- Taylor, C.; Noble, J. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from DWI, volume *Information Processing in Medical Imaging*. Springer-Verlag; Heidelberg: 2003. p. 600
- Tench C, Morgan P, Blumhardt L, Constantinescu C. Improved white matter fiber tracking using stochastic labeling. *Magnetic Resonance in Medicine* 2002;48:677–683. [PubMed: 12353285]
- Tschumperle, D.; Deriche, R. Constrained and unconstrained PDEs for vector image restoration; *Proceeding of the 12th Scandinavian Conference on Image Analysis*; 2001. p. 92-105.
- Tschumperle D, Deriche R. DT-MRI images: Estimation, regularization and application. *MICCAI'03* 2003;2809:530–541.LNCS
- Wang Z, Vemuri B, Chen Y, Mareci T. A constrained variational principal for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Transactions on Medical Imaging* 2004;23(8):930–939. [PubMed: 15338727]
- Westin, C.; Knutsson, H.; Diaz, RM.; Arencibia, AQ. *Computer Aided Systems Theory (EUROCAST'03)*, LNCS 2809. Springer Verlag; Las Palmas de Gran Canaria, Spain: 2003. Tensor field regularization using normalized convolution; p. 564-572.
- Westin C, Maier S, Mamata H, Navabi A, Jolesz F, Kikinis R. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis* 2002;6(2):93–108. [PubMed: 12044998]
- Xu D, Mori S, Solaiyappan M, van Zijl P, Davatzikos C. A framework for callosal fiber distribution analysis. *NeuroImage* 2002;17:1131–1143. [PubMed: 12414255]

Figure 1. Visual Assessment of the Smoothed Synthetic Data

The synthetic data consisted of 5 identical slices, each with a 32×32 array of anisotropic Diffusion Tensors (DTs). The eigenvalues of DTs were 0.5, 0.25, and 0.05, and their fractional anisotropy (FA) was 0.7, a value typical for FA value within white matter of the brain. The DTs were independently perturbed by random noise with standard deviations $\sigma = 0.05$ (*left, bottom row*) and $\sigma = 0.065$ (*right, bottom row*). We color encode and map the principal direction (PD) of the DTs (left two columns – *Red*: PD along X-axis; *Green*: PD along Y-axis; *Blue*: PD along Z-axis). The edges detected by our method, both before and after smoothing, are shown in the right two columns. In the smoothed tensor field (top row), the sharp boundaries between homogeneous regions are typically preserved. The smoothed field is similar to the synthetic field without noise. Even the narrow region that consists of 3 columns of tensors oriented parallel to the z-axis is preserved in the smoothed field. The smoothed fields were estimated with data weight $\eta = 0.5$, and neighborhood of size $\beta = 5$.

Figure 2. Quantitative Assessment of Smoothed Synthetic Data

The area under the Coherence Percent-Percent (CPP) plot demonstrates marked improvement in the coherence of the Principal Direction (PD) and Second Direction (SD) of tensors in the synthetic fields containing differing amounts of noise. The eigenvalues in the synthetic fields were 0.5, 0.25, and 0.05, with fractional anisotropy = 0.7, for all tensors before the addition of noise. The percentages of tensors having coherence values greater than a specific coherence value are plotted as fractions along the axes, for coherence decreasing from 1.0 to 0.0. For PDs, the area = 0.795 when the noise standard deviation $\sigma = 0.05$ (*left*), and the area = 0.7928 with $\sigma = 0.065$ (*right*). For SDs, the area = 0.8513 with $\sigma = 0.05$, and the area = 0.834 with $\sigma = 0.065$. The smoothed DT field was estimated with the data weight $\eta = 0.5$ and the neighborhood size $\beta = 5$. These plots demonstrate that our method significantly increased the coherence of DTs across the smoothed tensor field.

Figure 3. Maps of the Coherence Values for the Principal Direction (PD) and the Second Direction (SD)

These are shown for the smoothed and original tensor fields of our DTI dataset from a healthy individual. The smoothed tensor field was estimated for data weight $\eta = 0.5$ and neighborhood size $\beta = 5$. We computed the coherence values only for those tensors with fractional anisotropy (FA) > 0.2 in the original image. The magnified views of the maps from two regions (third row of images) show significant improvement in coherence of both PDs and SDs across the smoothed field. Furthermore, the small islands of high coherence values in the gray matter of original field have been removed in the smoothed field because we plotted only the coherence values for tensors with FA > 0.2 . These islands of tissue had FA value greater than 0.2 in the original field because of the noise. Our method appropriately reduced the noise and therefore the tissue within gray matter had FA < 0.2 , and thus these islands are not displayed in the maps for smoothed field.

Figure 4. Color Maps of the PD Before and After Smoothing

At each voxel in the smoothed and the original tensor fields, we color encoded the PD, as follows: *Red* = PD oriented parallel to the X-axis; *Green* = PD oriented parallel to the Y-axis; *Blue* = PD oriented along the Z-axis, perpendicular to the imaging plane. The smoothed tensor field was estimated for data weight $\eta = 0.5$ and neighborhood size $\beta = 5$. Because the eigenvalues are also adjusted in the smoothed field, the colors were not weighted with the fractional anisotropy values of the tensors. The color maps show visually that the PDs are more coherent in the smoothed tensor field and that the edges of homogeneous regions having differing orientations are well preserved in the smoothed field. Moreover, PDs in the smoothed field are similar to the PDs at corresponding voxels of the original field.

Figure 5. Maps of the Principal Directions (PDs) in the smoothed field (*left*) and the original field (*right*)

Arrows denote the direction of the PD at each voxel in the two tensor fields. The bottom row magnifies a small region in the map for the entire field within a slice. These magnified views show that the PDs are more coherent in the smoothed field and yet are close to those in the original field. Although the PDs in the smoothed field are changed only slightly from the original field, these small changes could produce significant changes in the fiber tracts that pass through these voxels. Therefore, we expect that reducing the amount of noise in the data and making PDs more coherent will generate tracts that are biologically more meaningful when tracking fibers. The smoothed tensor field was generated using our algorithm with neighborhood size $\beta = 5$ and data weight $\eta = 0.5$.

Figure 6. Statistics of the Coherence Values for PDs and SDs

We generated the histogram of the coherence values in each voxel of the smoothed and original tensor fields. The shift to the right in the histogram shows visually that the coherence of tensors improved in the smoothed field. For PDs, the number of DTs having a coherence value greater than 0.95 was 48373 after smoothing and 44881 before smoothing, a 7.78% increase. For SDs, the number of DTs having a coherence value greater than 0.95 were 28327 after smoothing and 20104 before smoothing, a 40.9% increase. Thus, DTs were better aligned in the smoothed field compared with those in the original field.

Figure 7. Coherence Percent-Percent Plots for PDs and SDs

These plots show the percentage of voxels having a coherence greater than a specified value in the smoothed image compared with the percentage in the original image. To generate the graph, percent-percent values are plotted for decreasing coherence values from 1.0 to 0.0. The graph therefore passes through 0.0 and 1.0. The dashed green line is the 45° line that has an area under the line equal to 0.5. *Left:* The CPP plot for PDs. The area under the blue line is 0.575. *Right:* The CPP plot for SDs. The area under the blue line is 0.572. Smoothing therefore increased coherence of both PDs and SDs. Although smoothing dramatically improved coherence of the SDs (Fig. 6), the PDs were better aligned overall in the smoothed field. The smoothed tensor field was estimate with neighborhood size $\beta = 5$ and data weight $\eta = 0.5$.

Figure 8. The Effects of Smoothing on FA Values

Maps of FA values across the smoothed field (top row) and original field (middle row).

Magnified maps for two brain regions show that FA values are significantly smoothed while preserving the distinctness and locations of the boundaries between homogeneous tissues. Fine, hair-like structures in white matter (arrows) are preserved, while FA values in their proximity have been smoothed. The smoothed tensor field was estimated with data weight $\eta = 0.5$ and neighborhood size $\beta = 5$.

Figure 9. Voxel-Wise Differences in FA Across the Smoothed and Original Tensor Fields

The FA map of the smoothed image was subtracted from the FA map of the original tensor field at each voxel. These difference values are shown in regions of the brain where FA values are greater than 0.2. The bottom row of images show magnified views of the difference maps from two regions near the lateral ventricles. Because FA values were substantially smoothed in the smoothed tensor field (Fig. 8), the difference map represents variations in FA values that are caused predominantly by noise in the original field.

Figure 10. The Effects of Smoothing on the Largest Eigenvalue

We computed the signal-to-noise ratio (SNR) for the largest eigenvalue at each voxel in the smoothed and original tensor fields. SNR was computed in regions of the original tensor field where fractional anisotropy was greater than 0.2. The top row of images shows maps of SNR in a single slice, and the bottom row shows the histograms of these values. The number of tensors in this slice having an SNR > 8 was 10226 before smoothing and 21280 after, an increase of 108.09%. The mean SNR in the smoothed field increased significantly (the P-value ≈ 0.0 because of the large number of tensors) as compared to the SNR in the original field (see the table below).

Figure 11. The Effects of Smoothing on the Second-Largest Eigenvalue

We computed and plotted the signal-to-noise ratio (SNR) of the second-largest eigenvalue at each voxel in Slice 27. The top row of images show the maps of these values and the bottom row shows their histograms. The number of tensors with SNR greater than 8 was 2944 in the original field and 13021 in the smoothed field, a 342.29% increase in the number of tensor having a large SNR. Therefore, smoothing considerably increased not only the coherence of the SD (Fig. 6) but also the SNR of the corresponding eigenvalue. The mean SNR in the smoothed field increased significantly ($P\text{-value} \approx 0.0$) compared to the SNR in the original field (see the table below).

Figure 12. The Effects of Neighborhood Size on Smoothing of FAs and PDs

We increased the neighborhood size β in our noise reduction algorithm from 3 to 11. The maps show color-encoded PDs (top row) and gray scale-encoded FA values (bottom row) in a single slice of the smoothed and original tensor fields. *Red* = PDs parallel to the X-axis; *Green* = PDs parallel to the Y-axis; *Blue* = PDs parallel to the Z-axis, perpendicular to the imaging plane. The color coding for the largest component of the PD was not weighted by the FA value, because FA values changed with increasing neighborhood size. Although smoothness increases with increasing neighborhood size, these maps show that our method preserves the distinctness and locations of the boundaries between homogeneous tissues. The smoothed field also appears close to the original tensor field (Fig. 4).

Figure 13. Variation in the Coherence of Eigenvectors with Increasing Neighborhood Size β and Data Weight η

We plot the area under the curve in the Coherence Percent-Percent (CPP) plots for the Principal Direction (PD, blue line) and Second Direction (SD, purple line) for (1) increasing β (left, with $\eta = 0.5$), and (2) increasing η (right, with $\beta = 5$). For increasing β (left), graphs for both PDs and SDs peak at $\beta = 5$ (and $\eta = 0.5$), suggesting that the optimal coherence in PDs and SDs is obtained when $\beta = 5$ and $\eta = 0.5$. For increasing η (right), the areas under the curve decrease monotonically. Thus, the optimal neighborhood size for our algorithm and DTI dataset is $\beta = 5$.

Figure 14. The Effects of Data Weight η on Coherence

For increasing data weight η , we plotted the percent increase in the number of tensors having a coherence value greater than 0.95 (blue) or 0.90 (purple), for both PDs (left) and SDs (right). For SDs, as expected, this percentage decreased with increasing values of η . Therefore, as expected, the SDs of tensors were better aligned with decreasing data weights. For PDs, the number of tensors having a coherence greater than 0.95 increased with decreasing η , as expected (i.e., alignment of tensors improves with decreasing data weights). In contrast, the number of tensors having coherence values greater than 0.90 declined with decreasing values of η . This decline in coherence with greater smoothing was attributable in this instance to the good alignment of PDs across the original image, so that further adjustment of their alignment decreased the coherence values for tensors located along the boundaries of homogeneous tissues.

Figure 15. Visual Assessment of the Improved Coherence of PDs and Effect on FA with Anisotropic Filtering

The anisotropic filter was applied either once (*top row*) or 10 times (*bottom row*) to the DW images. Visually, increasing the number of iterations of the filter increased the smoothness of FA values (*left*) and the coherence of tensors in homogeneous tissues (*right*). The boundaries of these regions were also smoothed and blurred, however, with the larger number of iterations. In contrast, our method preserves boundaries between homogeneous tissues (Figs. 8 & 12) compared to our implementation of anisotropic filtering. Color coding is as defined in Fig. 12.

Figure 16. Quantitative Assessment of the Improved Coherence in Tensors with Anisotropic Filtering

The smoothed tensor field was estimated from DW images that were smoothed with an anisotropic filter iterated either once (*top row*) or 10 times (*bottom row*). For smoothed tensor fields estimated with one iteration of the filter, the area under the curve = 0.5455 in the Coherence Percent-Percent (CPP) plot for Principal Directions (PDs), and the area = 0.5512 for Second Directions (SDs). When the filter was iterated 10 times, the area = 0.732 for PDs, and area = 0.7771 for SDs in the smoothed field. Therefore, increasing the number of iterations for the anisotropic filter significantly increased the coherence of tensors across the smoothed tensor field. Using our algorithm for noise reduction, the area = 0.5789 for PDs, and area = 0.5792 for SDs (Fig. 7). Thus, our method performs better than our implementation of the anisotropic filtering with a single iteration.

Figure 17. The Effects of Anisotropic Filtering on Eigenvalues

The filter was applied either once (*top row*) or 10 times (*bottom row*) to the original DW images. We generated the histograms of Signal-to-Noise Ratio (SNR) for the largest (*left column*) and second-largest (*right column*) eigenvalues computed at each voxel in the smoothed field. The table shows the number of tensors having an SNR > 8 for the largest and the second-largest eigenvalues in the original field as well as in the three smoothed fields estimated using either a single iteration or 10 iterations of the anisotropic filter, and by our algorithm. Thus, our method increased substantially the SNR of eigenvalues compared with a single iteration of the anisotropic filter, although the increase in SNR with our algorithm was less than that obtained with 10 iterations of the anisotropic filter.

Figure 18.

Quantitative Assessment of Smoothing Using Riemannian Distance to compute the similarity of two tensors. The area under the Coherence Percent- Percent (CPP) plot is 0.6474 and 0.629 for the Principal Direction (PD, *left*) and Second Direction (SD, *right*), respectively. Therefore, use of the Riemannian distance within our framework of Markov Random Field (MRF) generated a tensor field having a greater coherence of both the PD and SD of tensors compared to the smoothed field generated using conditional probability based on perturbation theory (Fig. 7). The method using Riemannian distance, however, smoothed across boundaries of homogeneous tissues in the brain (Fig. 19).

Figure 19.

Visually Comparing the Edge Maps in the Smoothed Tensor Fields generated using both the conditional probability (*middle*) and the Riemannian distance (*right*) within our MRF framework. The edge map for the original tensor field is on the left. Because of the large amount of noise in gray matter, our method seemed to place an edge between every voxel in gray matter. Smoothing using both measures generated tensor fields with significant reduction in the edges across white matter. The Riemannian distance, however, generated a tensor field with regions of thinner and broken edges and broken contours (arrows). Thus, using Riemannian distance generated fields with higher coherence across PD and SD of tensors (Fig. 18), but at the expense of smoothing across boundaries of homogeneous tissues.

Figure 20.

Visual Comparison of the Fractional Anisotropy Maps generated by our method using both conditional probability (*left*) and Riemannian distance (*right*) for computing similarity of two tensors within our framework of Markov Random Field. These maps show significantly larger amount of smoothing in the FA values using Riemannian distance compared with the smoothing of FA values in the tensor field generated using our perturbation measure. Thus, the larger coherence of PDs in smoothed field that is generated using Riemannian distance comes at the expense of more smoothing across differing homogeneous tissues in the brain. Our algorithm, in contrast, generates smoothed tensor fields while preserving the boundaries between homogeneous tissues.

SNR	Field	Smoothed	Original
	Median	6.801	6.0445
	Mean	6.9735	5.933

SNR	Field	Smoothed	Original
	Median	6.0445	5.1488
	Mean	5.975	5.09