Joint Fractional Segmentation and Multi-Tensor Estimation in Diffusion MRI

Xiang Hao and P. Thomas Fletcher

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

Abstract. In this paper we present a novel Bayesian approach for fractional segmentation of white matter tracts and simultaneous estimation of a multi-tensor diffusion model. Our model consists of several white matter tracts, each with a corresponding weight and tensor compartment in each voxel. By incorporating a prior that assumes the tensor fields inside each tract are spatially correlated, we are able to reliably estimate multiple tensor compartments in fiber crossing regions, even with low angular diffusion-weighted imaging (DWI). Our model distinguishes the diffusion compartment associated with each tract, which reduces the effects of partial voluming and achieves more reliable statistics of diffusion measurements. We test our method on synthetic data with known ground truth and show that we can recover the correct volume fractions and tensor compartments. We also demonstrate that the proposed method results in improved segmentation and diffusion measurement statistics on real data in the presence of crossing tracts and partial voluming.

1 Introduction

Diffusion-weighted imaging (DWI) is a magnetic resonance imaging (MRI) modality that can measure the directional diffusion of water in tissue. In diffusion tensor imaging (DTI) a second-order model of diffusion within a voxel is estimated from the DWI data. While the diffusion tensor is an elegant description of anisotropic diffusion in white matter, it is limited to representing only one tract in each imaging voxel. At several regions of the brain, two or more white matter tracts are passing through each other, such as the intersection of the corona radiata and the corpus callosum. In addition, some white matter tracts are mixed with gray matter or cerebrospinal fluid (CSF) at the boundary, such as the corpus callosum and lateral ventricles. In these cases, a single diffusion tensor is incapable of distinguishing between multiple diffusion compartments.

It has been shown that partial volume effects and underestimation of diffusion measurements occur in crossing areas [1,9,8]. For example, Alexander et al. [1] show that the trace of the diffusion tensor will tend to be underestimated in presence of partial voluming, and Metzler-Baddeley et al. [8] show that CSF-based partial volume artifacts have a larger impact for tensors with smaller fractional anisotropy (FA). What's more, underestimation of diffusion tensor measurements could bias statistics of diffusion measurements, providing misleading results in clinical studies. Although the partial volume effects might

2 Xiang Hao and P. Thomas Fletcher

be detected during the DTI analysis [13], it cannot resolve the partial volume effects completely, and other methods that can correct the partial volume effects are needed. High-angular resolution diffusion imaging (HARDI) (see [2] for a review) has been introduced as a means of distinguishing multiple diffusion compartments. Several multi-compartment models of diffusion have been introduced, such as the multi-tensor [12] and ball-and-stick [3] models. However, one drawback to HARDI is the increased imaging time required, which has been a barrier to its introduction in clinical studies.

Several works have proposed correcting the partial volume effects using lowangular resolution images (approximately 30 gradient directions). Pasternak et al [10] used a smooth regularization to reduce partial volume effects using multitensor fitting and showed that the regularization can help solve ill-posedness, but they only test their methods on two fiber crossings. Landman et al [7] used a multi-tensor model and compressed sensing to solve the fiber crossing using low angular DWI, but they use a set of basis directions to map the multiple tensor model to lower dimensionality, which could bias the estimated tensors.

None of these methods combine multiple compartment estimation with white matter tract segmentation, which is often a primary goal of DWI analysis. In this paper, we propose a method that jointly solves the tract segmentation and multi-tensor model estimation problem. In each voxel our model assigns a fractional weight to each fiber tract indicating the proportion of that voxel that belongs to the given tract. These weights also serve as the volume fractions in the multi-tensor model for that voxel, linking the segmentation to the diffusion estimation problem. We impose a Markov random field (MRF) spatial prior to take advantage of spatial redundancy during the estimation and to regularize the multi-tensor field. To the best of our knowledge, ours is the first model that combines fractional segmentation of white matter with multi-tensor estimation. We compare our method with state-of-the-art binary segmentations of white matter tracts, which sometimes under- or over-segment tracts. We demonstrate that the fractional weights can improve statistical analysis of derived measurements, such as fractional anisotropy (FA), by appropriately weighting the data associated with a particular tract. Using synthetic data with known ground truth, as well as real DWI brain data, we show that our method improves tract segmentations, distinguishes multiple tissue compartments, and provides better diffusion estimates even when using only 12 gradient directions.

2 Fractional Segmentation and Multi-Tensor Estimation

In this section we propose a Bayesian approach to simultaneously estimate both the fractional segmentation of white matter tracts and the multi-tensor diffusion model. Our goal is to estimate the volume fractions and diffusion tensors in such a way that the DWI signal from the estimated multi-tensor model matches the measured DWI signal and the estimated tensor compartments in each tract are spatially smooth. This leads to a maximum a posteriori (MAP) estimation approach, in which the log posterior is

$$\log p(\theta|S) \propto \log p(S|\theta) + \log p(\theta),$$

where θ represents the parameters of the model, namely, the multi-tensor field and volume fractions, and S is the original DWI signal. The likelihood, $p(S|\theta)$, models the fit of the multi-tensor model in each voxel to the DWI signal, and the prior, $p(\theta)$, is a Markov Random Field (MRF) smoothness prior on the multi-tensor field. We now describe both the likelihood and prior in detail.

2.1 Likelihood - The Data Attachment Term

To model multiple diffusion compartments within a voxel, whether from crossing white matter fiber tracts or mixtures of white matter with CSF or gray matter, we use a multi-tensor model of the given DWI signals. The multi-tensor model uses n tensor compartments, D_i , and each tensor compartment is associated with a nonnegative volume fraction f_i to model the DWI signals S_i as

$$S_j(x) = S_0(x) \sum_{i=1}^n f_i(x) e^{-bg_j^T D_i(x)g_j}, \text{ with } \sum_i^n f_i = 1,$$
(1)

where b is the b-value, S_0 is the baseline image, g_j is the j-th gradient encoding direction. In contrast to the usual multi-tensor model, we also associate each volume fraction f_i to a white matter tract segmentation. That is, we want to segment the DWI image into n white matter tracts, where $f_i(x)$ represents the fraction of the voxel at position x that is occupied by fibers in the *i*th tract.

Now, assuming the DWI signal is corrupted by additive, i.i.d. Gaussian noise, our log-likelihood is the following DWI signal matching term:

$$\log p(S_j | f_i, D_i) \propto -\int_{\Omega} \sum_{j=1}^m \left(S_0(x) \sum_{i=1}^n f_i(x) e^{-bg_j^T D_i(x)g_j} - S_j(x) \right)^2 dx.$$

This noise model could be replaced with a Rician noise likelihood or a non-central chi-distribution for multichannel MRI. Conceptually this would not change the underlying methodology, but rather make the estimation procedure slightly more complicated. The Gaussian likelihood is a simplifying approximation.

2.2 The Markov Random Field Prior

Inside a white matter tract, it is reasonable to think that the tensors field should be spatially correlated, i.e., the tensor field flows smoothly inside a white matter tract. Thus we incorporate a MRF prior on the tensor fields inside each tract, given by

$$p(D_i) = \frac{1}{Z}e^{-U(D_i)}$$
, and $U(D_i) = \lambda \int_{\Omega} \sum_{i}^{n} \phi(D_i(x)) dx$,

4 Xiang Hao and P. Thomas Fletcher

where Z is a normalization constant, λ is a weighting parameter, and $\phi(D_i(x))$ is a function that measures the correlation of tensors in the *i*th tract, around x.

With a uniform prior on the fractions, f_i , and by combining the log-likelihood term with the prior, we end up minimizing the energy

$$E(f_i, D_i) = \int_{\Omega} \bigg(\sum_{j=1}^m \left(S_0(x) \sum_{i=1}^n f_i(x) e^{-bg_j^T D_i(x)g_j} - S_j(x) \right)^2 + \lambda \sum_i^n \phi(D_i(x)) \bigg) dx$$

We use the full six-component tensor model to represent D_i , but we want to emphasize that D_i could be any other diffusion model, such as the ball-and-stick model. In addition, we use $\phi(D_i(x)) = \|\nabla D_i(x)\|^2 = \sum_{p,q} \sum_{y \in N(x)} (D_i^{pq}(y) - D_i^{pq}(x))^2$, where N(x) is a neighborhood around x, and D_i^{pq} represents component in the *p*th row and the *q*th column of the matrix D_i , but it could be generalized to other functions for measuring the smoothness of D_i .

2.3 Optimizing f_i and D_i

We use a gradient descent algorithm to compute f_i and D_i . The partial derivative of E respect to f_i is

$$\frac{\partial E}{\partial f_i} = 2S_0(x) \sum_{j=1}^m \left(e^{-bg_j^T D_i(x)g_j} \cdot \left(S_0(x) \sum_{i=1}^n f_i(x) e^{-bg_j^T D_i(x)g_j} - S_j(x) \right) \right),$$

and the partial derivative of E respect to D_i^{pq} is

$$\begin{aligned} \frac{\partial E}{\partial D_i^{pq}} &= \bigg(\sum_{j=1}^m 2\Big((S_0(x) \sum_{i=1}^n f_i(x) e^{-bg_j^T D_i(x)g_j} - S_j(x)) f_i e^{-bg_j^T D_i(x)g_j} S_0(x) \\ &\cdot (-bg_j^p g_j^q) \Big) + 4 \cdot \lambda \sum_{y \in N(x)} (D_i^{pq}(x) - D_i^{pq}(y)) \bigg) \cdot \psi(p,q), \end{aligned}$$

where g_j^p represents the *p*th element of the vector g_j , and $\psi(p,q) = 1$ if p = qand $\psi(p,q) = 2$ if $p \neq q$. Special care is required in the update of f_i . In every iteration of our gradient descent algorithm, we project the computed gradient of *E* with respect to f_i onto the constraint hyperplane defined by $\sum_{i=1}^n f_i = 1$. To ensure that the f_i remain positive, we also need to project the gradient on the simplex boundary defined by the constraints $f_i \geq 0$ when necessary. Then we do a line search to compute the optimal step size for updating f_i .

To initialize our gradient descent, we first compute binary segmentations for each of n-1 tracts of interest. The *n*th label is reserved as a "background" label that does not belong to any of the tracts. The binary segmentations just need to roughly capture the tracts, which could be done by several methods. We use a front-propagation geodesic segmentation [5, 6], which uses a Riemannian metric derived from the diffusion tensor field and constructs white matter tracts as geodesics connecting two regions-of-interest (ROI) on the resulting manifold. These geodesics have the desirable property that they tend to follow the main eigenvectors of the tensors, yet still have the flexibility to deviate from these directions when it results in lower costs. This makes such methods more robust to noise and also allows them to pass through crossing regions. After segmenting the n-1 tracts of interest, for each voxel inside the predefined image mask, if no segmentations include the voxel, we initialize f_n as 1, and D_n as the weighted least squares estimate of a single-tensor model. Otherwise, we initialize the fractions f_i equally amongst all segmentations that intersect at the voxel, and we initialize all corresponding tensors D_i in that voxel with the same weighted least-squares estimate of a single-tensor model at that voxel.

2.4 Path Regression Along the Segmented Tracts

Our estimated fraction f_i can also be used to improve statistical summaries of the diffusion data along segmented tracts. Here we demonstrate an example of the use of these fractions in nonparametric regression along a tract. Let $\{x_i\}$ be the collection of voxel locations within a segmented tract. Each voxel has an associated parameter, $s_i = s(x_i) \in [0, 1]$, which denotes the arc-length parameter along the pathway at the spatial location x_i [5]. Denote by d_i a data value at the location x_i . This data may be a full diffusion tensor, or a derived measure, such as FA or MD. We compute a continuous description of the data as a function of s using a nonparametric kernel regression,

$$y(s) = \frac{\sum_{i}^{N} d_{i} f_{i} G(s - s_{i}, \sigma)}{\sum_{i}^{N} f_{i} G(s - s_{i}, \sigma)},$$

where $G(\mu, \sigma)$ denotes a Gaussian kernel with mean μ and standard deviation σ . In the kernel regression, in addition to the Gaussian kernel, each data value is weighted by the fraction f_i , which makes our regression more robust to partial voluming since our fractional segmentation will appropriately assign small f_i to data values influenced by partial voluming. The function y defined above gives a continuous average of the data along the pathway, which can be used to quantify the diffusion measurement along a pathway as shown in Figure 2.

3 Results

In this section we demonstrate the advantages of the fractional segmentations and improvements in multi-tensor estimations computed by our method on both synthetic and real data. On synthetic data with ground truth, we test our method on both low and high angular DWI and compare it with the multi-tensor model estimation implemented in Camino [4] using HARDI. Our measure of quality is the root mean square error (RMSE) between the estimated volume fractions/tensors and true volume fractions/tensors. In addition, we also show that our method improves the diffusion measurement statistics. On real data, we demonstrate that the proposed method results in improved segmentation and 6



Fig. 1. Left: synthetic white matter tract with partial volume effects at the boundary (shown in white). Middle: our estimated tensors. Right: our fractional segmentation. We subsample the tensor field by a factor of two both horizontally and vertically in order to visualize it.

diffusion measurement statistics in crossing tracts and in the presence of partial voluming. To visualize the fractional segmentation, we convert our fractional segmentation to a color-coded RGB image, where each color component is set to a corresponding volume fraction. For visualization purposes, tensors with weight lower than 0.2 are not shown.

3.1 Correcting Partial Volume Effect

To test the ability of our method to correct partial volume effects, we generated one white matter tract whose boundary is mixed with isotropic tensors as shown in Figure 1. The ground truth is a straight white matter tract with width 8 voxels and length 38 voxels, mixed with isotropic diffusion tensors at the boundary (shown in white in Figure 1) with equal volume fractions in the multi-tensor model (1). The sphere tensors used to simulate CSF have eigenvalue $(3,3,3) \times 10^{-3}$ mm²/s. The tensors in the white matter tract have eigenvalue $(1.6,0.4,0.4) \times 10^{-3}$ mm²/s. Using the mixed tensor field, we generate DWIs with 12 gradient directions and slightly blur the DWIs to simulate the point spread function (PSF) arising from imaging. We corrupt the DWIs with Rician noise to get a SNR of 20 (left picture of Figure 1).

Our estimated white matter tensor field is in the middle picture of Figure 1. As we can see, the estimated tensor field can recover the true white matter tensor from the mixed tensors at the boundary. In addition, our estimation algorithm denoises tensors in the interior. Our color-coded fractional segmentation is shown in the right panel of Figure 1. In this example, the segmented white matter tract is the blue channel and the exterior is the green channel. The boundary is correctly assigned a 50% mixture of both compartments. The RMSE of our fractional segmentation of the white matter tract is 7.74×10^{-2} , and the RMSE of our estimated tensors is 5.22×10^{-5} .

In the left two pictures of Figure 2, we do a regression analysis of the FA in the segmented tracts with $\sigma = 0.1$ as mentioned in Section 2.4, where the *x*-axis is the arc-length position along the tract. We first plot the FA values as points. The color is coded by the volume fractions: solid blue denotes 0, and solid red denotes 1, so the binary segmentation values are always shown in red. The regression along the tract is shown in black and the ground truth is shown as a



Fig. 2. Left two pictures: FA regression of simulated straight tract using binary and fractional segmentations, respectively. Ground truth is shown in green and the regression is shown in black. Right two pictures: FA regression of arcuate fasciculus from binary and fractional segmentations, respectively.



Fig. 3. Generated fiber crossing data. Left: 45° crossing. Middle: a curved tract crossing with a straight tract. Right: three orthogonal tracts crossing. We subsample the tensor field by a factor of three both horizontally and vertically in order to visualize it.

dashed green line. It is clear that the partial volume effects bias the regression in the binary segmentations. In our method, the FA values are more tightly distributed around the ground truth, due to the reduction of the partial volume effects and the denoising from our spatial prior.

3.2 Fiber Crossing White Matter Tracts

To test the performance of the our method in the presence of multiple white matter tracts crossing, we generate three fiber crossing tensor fields that have similar properties to many white matter tracts in the brain. They are 1) two bars crossing at the center of the image at an angle of 45° , 2) a curved tract crossing with a bar, and 3) three orthogonal bars crossing at the center of the image. We show a center slice of each tensor field in Figure 3.

The tensors in each white matter tract have eigenvalues $(1.6, 0.4, 0.4) \times 10^{-3} \text{mm}^2/\text{s}$. We generate two sets of DWI datasets, one with 12 gradient directions and the other with 64 gradient directions. Each voxel of the generated crossing DWI at the crossing area was computed based on the multi-tensor model (1), and we use $b = 1000 \text{s/mm}^2$ for both datasets. The true fractions in the crossing regions for the three data sets are (0.4, 0.6), (0.4, 0.6), and (0.3, 0.3, 0.4), respectively. In addition, the DWI was corrupted by Rician noise to simulate SNR of 20.



Fig. 4. Estimated tensors of the crossing region for the three crossing datasets: Top row is for the 45° crossing, Middle row is for the curved tract crossing, and the Bottom row is for the three orthogonal tracts crossing. 1st column: our estimated tensors from noiseless 12-direction DWI. 2nd column: our estimated tensors from noisy 64-direction DWI. 3rd column: our estimated tensors from noisy 12-direction DWI. Last column: estimated tensors from Camino using noisy 64-direction DWI.

In Figure 4, we show the estimated tensors from our method and Camino. We test Camino only on noisy 64-direction DWI (shown in the last column of Figure 4) because Camino cannot estimate a multi-tensor model from only 12 gradient directions (it estimates each voxel independently). However, we can test our method on both noisy 64-direction DWI and 12-direction DWI because our spatial prior utilizes information from multiple neighboring voxels. On the first column of Figure 4, we test our method on noiseless 12-direction DWI to show that we can recover multiple tensors, even when the solution is underdetermined at single voxels because of our spatial prior. On the second and third column, we test our method on noisy 64-direction DWI and noisy 12-direction DWI. The results are much improved over those from Camino, and our method works slightly better on 64-direction DWI since it has more information. The results from Camino are shown in the last column of Figure 4. We gave Camino the advantage of a map of the correct number of tensor compartments in each voxel, but did not provide this extra information to our algorithm.

We also show our fractional segmentation in Figure 5. Since the weights computed from Camino do not correspond to segmentations, we solve a correspondence problem by minimizing the angles between estimated tensors from Camino and the true tensors. After the optimal order is found, we use the corresponding weights as the fractional segmentation. Outside the crossing region, as we mentioned earlier, we tell Camino the number of tensors, so Camino always has the true weights outside. The finding is consistent with the case for tensor estimation. On the noiseless data, the error between our solution and the true



Fig. 5. Estimated fractional segmentations. The order is the same as Figure 4.

	45° Crossing		Curved Crossing		Three Crossing	
	Weight	Tensor	Weight	Tensor	Weight	Tensor
Clean 12-dir	1.16E - 04	5.61E - 07	1.44E - 02	2.97E - 05	1.33E - 05	1.12E - 07
Noisy 64-dir	7.38E - 02	1.39E - 04	4.89E - 02	1.38E - 04	5.16E - 02	2.09E - 04
Noisy 12-dir	8.37E - 02	1.54E - 04	6.18E - 02	1.63E - 04	5.71E - 02	3.07E - 04
Camino	2.6E - 01	4.40E - 03	2.34E - 01	2.89E - 03	$1.45E\!-\!01$	3.14E - 03

Table 1. RMSE of the estimated volume fractions and tensor compartments of our method (first three rows) and Camino using noisy DWI with 64 directions (last row).

solution is close to zero as shown is Table 1. On the noisy data, the solution of our method is worse when we use the 12-direction DWI than the one using the 64-direction DWI, but generally our results are fairly good for all three cases, and they look better than Camino inside the crossing regions.

In Table 1, we do a quantitative comparison of our method and Camino. Since we give Camino the true weight outside the crossing regions, to do a better comparison, we compute the RMSE of the estimated volume fractions only inside the crossing region, but we compute the RMSE of tensors in the whole white matter tracts. As shown in the Table, the RMSE of our method is much smaller than the RMSE of Camino, and our method works best on the noiseless data, but when there is noise, our method has smaller RMSE when using DWI with more gradient directions.

3.3 Real Data

We now show the results of our method with $\lambda = 10^9$ applied on a DWI of a volunteer. DWI data were acquired on a Siemens Trio 3.0 Tesla Scanner with



Fig. 6. The binary segmentation is shown on the 1st and 3rd columns and our fractional segmentation is shown on the 2nd and 4th columns. The overlay tensor fields are the DTI tensor field except for the last one, where our estimated tensors are displayed.

an 8-channel, receive-only head coil. DWI was performed using a single-shot, spin-echo, EPI pulse sequence and SENSE parallel imaging (undersampling factor of 2). Diffusion-weighted images were acquired in 12 non-collinear diffusion encoding directions with diffusion weighting factor b=1000 s/mm² in addition to a single reference image (b=0). Data acquisition parameters included the following: contiguous (no-gap) fifty 2.5mm thick axial slices with an acquisition matrix of 128 x 128 over a FOV of 256 mm (2 x 2 mm² in-plane resolution), 4 averages, repetition time (TR) = 7000 ms, and echo time (TE) = 84 ms. Eddy current distortion and head motion were corrected using an automatic image registration program [11]. Distortion-corrected DW images were interpolated to 1 x 1 x 1 mm³ voxels, and six tensor elements were calculated using weighted least squares. The tensor upsampling is done only for the purposes of numerical computations on the voxel grid; a finer grid results in higher numerical accuracy.

In Figure 6, we test our method on the arcuate fasciculus of the brain. In the 1st and 3rd pictures, we show the initial binary segmentation and the DTI tensor field. As we can see, the binary segmentation over-segments the tract since it includes many voxels (red arrow) below the tract which should be inside the inferior longitudinal fasciculus. In addition, it also under-segments the tract, missing the blue area on the lower right of image (black arrow), which should be included in the segmentation. On the 2nd column, we overlay the DTI tensor field on our fractional segmentation to demonstrate that our segmentation corrects the errors from the binary segmentation. In the last picture, we overly our estimated tensors on top of our segmentation. We can see that our segmentation reduces partial voluming effect at the boundary of tract and has lower weights at the border of the arcuate fasciculus and inferior longitudinal fasciculus, where the tensors from the two different tracts are mixed together. In addition, our segmentation also includes the blue area, which the binary segmentation missed.

In the right two pictures of Figure 2, we do a similar regression analysis as mentioned in Section 3.1. We can see the FA values are more tightly distributed on the right picture than the left one, which is a sign of possible reduction of partial volume effects as shown in Section 3.1. What's more, it may improve statistical power in clinical studies by reducing the within-subject variance.

We also test our method on a brain region with complex white matter organization as shown in Figure 7. In this brain region, we segment three white



Fig. 7. 1st row: the DTI field, our estimated multi-tensor field and their closer views of the crossing region (inside the red box). 2nd row: our fractional segmentation.

matter tracts as our binary input of our method. These three tracts are one branch from the corpus callosum, the corticospinal tract, and the superior longitudinal fasciculus. In the top row of Figure 7, we compare the DTI model with the proposed method. In the 1st and 3rd pictures, we overlay DTI tensor field on the FA background and its closer view of the crossing region. In the 2nd and 4th pictures, we show our estimated tensors and a closer view of the same crossing region. As we can see, the DTI model can not tell the different fibers inside a voxel. However, it is clear that our method can estimate multiple tensor compartments in the crossing region and also reduce the isotropic partial voluming of the estimated tensors in the corpus callosum (blue arrow). Our fractional segmentation is shown in the 2nd row of Figure 7, where we overlay our color-coded fractional segmentation on the top of FA background on the left two pictures and we overlay the estimated tensor compartments on the corresponding fractional segmentation on the right three pictures. We can see that the tensors in each white matter tract are spatially consistent due to our spatial prior, and the fractional segmentation is biologically reasonable.

4 Conclusion and Future Work

We present a Bayesian approach for joint fractional segmentation of white matter tracts and multi-tensor estimation in DWI. Our method can reliably estimate multiple tensor compartments in fiber crossing regions even with low angular DWI. There are three areas we have identified as potential future work. First, our noise model could be replaced with a Rician noise likelihood. Second, we can use an anisotropic spatial prior [14] instead of the isotropic spatial prior to prevent blurring across the whiter matter tracts. Finally, we plan to investigate the sensitivity of our current gradient-based optimization to the initial binary segmentation and explore stochastic optimization schemes.

Acknowledgements. We would like to thank Dr. Janet Lainhart for providing the image data, funded by NIH Grant R01 MH080826. This work was supported by NIH Grant R01 MH084795.

References

- A. L. Alexander, K. M. Hasan, M. Lazar, J. S. Tsuruda, and D. L. Parker. Analysis of partial volume effects in diffusion-tensor MRI. MRM, 45(5):770–780, 2001.
- H.-E. Assemlal, D. Tschumperl, L. Brun, and K. Siddiqi. Recent advances in diffusion MRI modeling: Angular and radial reconstruction. *MedIA*, 15(4):369 – 396, 2011.
- T. E. J. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *MRM*, 50:1077– 1088, 2003.
- P. A. Cook, Y. Bai, Nedjati S. Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *ISMRM*, page 2759, May 2006.
- P. T. Fletcher, R. Tao, W.-K. Jeong, and R. T. Whitaker. A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. In *IPMI*, pages 346–358, 2007.
- 6. X. Hao, R. T. Whitaker, and P. T. Fletcher. Adaptive riemannian metrics for improved geodesic tracking of white matter. In *IPMI*, pages 13–24. 2011.
- B. A. Landman, H. Wan, J. A. Bogovic, P.-L. Bazin, and J. L. Prince. Resolution of crossing fibers with constrained compressed sensing using traditional diffusion tensor MRI. *NeuroImage*, 59:2175–86, 2012.
- C. Metzler-Baddeley, M. J. O'Sullivan, S. Bells, O. Pasternak, and D. K. Jones. How and how not to correct for CSF-contamination in diffusion MRI. *NeuroImage*, 59(2):1394 – 1403, 2012.
- H Oouchi, K Yamada, K Sakai, O Kizu, T Kubota, H Ito, and T Nishimura. Diffusion anisotropy measurement of brain white matter is affected by voxel size: underestimation occurs in areas with crossing fibers. *AJNR*, 28(6):1102–6, 2007.
- O. Pasternak, Y. Assaf, N. Intrator, and N. Sochen. Variational multiple-tensor fitting of fiber-ambiguous diffusion-weighted magnetic resonance imaging voxels. *Magnetic Resonance Imaging*, 26(8):1133 – 1144, 2008.
- G. Rohde, A. Barnett, P. Basser, S. Marenco, and C. Pierpaoli C. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *MRM*, 51:103–114, 2004.
- D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *MRM*, 48(4):577–582, 2002.
- S. B. Vos, D. K. Jones, M. A. Viergever, and A. Leemans. Partial volume effect as a hidden covariate in DTI analyses. *NeuroImage*, 55(4):1566 – 1576, 2011.
- Zhizhou Wang, B.C. Vemuri, Yunmei Chen, and T.H. Mareci. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex dwi. *IEEE Transactions on Medical Imaging*, 23(8):930–939, 2004.

¹² Xiang Hao and P. Thomas Fletcher