



HHS Public Access

Author manuscript

Med Image Comput Comput Assist Interv. Author manuscript; available in PMC 2016 February 03.

Published in final edited form as:

Med Image Comput Comput Assist Interv. 2010 ; 13(0 1): 674–681.

Multi-Diffusion-Tensor Fitting via Spherical Deconvolution: A Unifying Framework

Thomas Schultz¹, Carl-Fredrik Westin², and Gordon Kindlmann¹

¹Computer Science Department and Computation Institute, University of Chicago, Chicago IL, USA

²Laboratory of Mathematics in Imaging, Brigham and Women's Hospital, Harvard Medical School, Boston MA, USA*

Abstract

In analyzing diffusion magnetic resonance imaging, multi-tensor models address the limitations of the single diffusion tensor in situations of partial voluming and fiber crossings. However, selection of a suitable number of fibers and numerical difficulties in model fitting have limited their practical use. This paper addresses both problems by making spherical deconvolution part of the fitting process: We demonstrate that with an appropriate kernel, the deconvolution provides a reliable approximative fit that is efficiently refined by a subsequent descent-type optimization. Moreover, deciding on the number of fibers based on the orientation distribution function produces favorable results when compared to the traditional F-Test. Our work demonstrates the benefits of unifying previously divergent lines of work in diffusion image analysis.

1 Introduction

The diffusion tensor model [1] is widely used for analyzing data from diffusion weighted magnetic resonance imaging (DW-MRI), but is inadequate in situations of partial voluming and fiber crossings. Multi-compartment models provide a natural extension by combining multiple diffusion tensors. They have been used to study the effects of partial voluming [2], and to analyze the diffusion weighted signal in voxels with multiple fiber contributions [3].

Fitting multi-tensor models requires nonlinear optimization, for which previous work has used descent-type algorithms [3–6]. However, these methods only find the global optimum when provided with an initial guess that is sufficiently close to the final solution. At the state of the art, numerous randomized re-starts are tried to reach a suitable optimum with high probability [3, 7, 6], which incurs an excessive computational cost. Alternatively, regularization over spatial neighborhoods [8] requires numerical solution of a partial differential equation.

We present a novel way of applying spherical deconvolution [9] to “kick-start” model fitting, as summarized in Figure 1. We derive a deconvolution kernel to approximate the ball-and-stick model [10, 11], a common variant of the multi-tensor model. The deconvolution result is analyzed by a recent method for fitting discrete models to continuous

orientation distribution functions [12]. This decides the number of fibers and provides a starting point to make subsequent optimization [13, 4–6] more reliable and efficient.

Diffusion image analysis approaches based on spherical deconvolution have generally been distinct from those involving explicit fitting of multi-fiber models. Our main contribution is showing how these two approaches can be adapted and combined to create a unified algorithmic solution that offers advantages of both.

2 Related Work

2.1 Multi-Tensor Models

Multi-tensor models assume k diffusion compartments with little or no exchange during measurement time [2], each parametrized by a symmetric 3×3 diffusion tensor \mathbf{D}_i per compartment, with eigenvalues $\lambda_1 \ \lambda_2 \ \lambda_3 \geq 0$. The signal fractions $f_i \in [0, 1]$ sum to unity. With this, the signal $S(\mathbf{g})$ is predicted as

$$S(\mathbf{g}) = S_0 = \sum_{i=1}^k f_i e^{-b\mathbf{g}^T \mathbf{D}_i \mathbf{g}}, \quad (1)$$

where k is the number of compartments, S_0 is the non diffusion-weighted signal, b is the diffusion weighting, and \mathbf{g} is the diffusion-sensitizing gradient.

A full k -tensor model has $7k-1$ degrees of freedom, but additional constraints are imposed in practice. When a single non-zero b -value is used, a mathematical indeterminacy prevents simultaneous estimation of the isotropic part of \mathbf{D}_i and its volume fraction f_i [7]. Most authors constrain the isotropic part, for example by assuming equal eigenvalues on all \mathbf{D}_i [3, 14, 5], but some prefer to fix the f_i instead [7, 15]. A very common assumption is axial symmetry ($\lambda_2 = \lambda_3$) [3, 7, 8, 10, 14–16]. In many cases, k is limited to $k \leq 2$ [3, 7, 16, 5].

In this work, we focus on the ball-and-stick model [10], which assumes that all \mathbf{D}_i have equal λ_1 ; a single “ball” compartment is completely isotropic ($\lambda_1 = \lambda_2 = \lambda_3$), the remaining “stick” compartments are perfectly linear ($\lambda_2 = \lambda_3 = 0$). For n fiber terms, this leads to $k = n + 1$ compartments and $3n + 1$ degrees of freedom. We consider the model up to $n = 3$.

2.2 Spherical Deconvolution

Rather than assuming a fixed number of compartments, spherical deconvolution reconstructs an orientation distribution function (ODF) $F(\theta, \varphi)$, which specifies a continuous density of volume fractions on the unit sphere. The predicted signal $S(\theta, \varphi)$ is then defined as the convolution of the ODF $F(\theta, \varphi)$ with an axially symmetric single-fiber response function $R(\gamma)$,

$$S(\theta, \varphi) = \int_0^{2\pi} \int_0^\pi F(\theta', \varphi') R(\gamma') \sin(\theta') d\theta' d\varphi' \quad (2)$$

where γ' is the angle between directions given by (θ, φ) and (θ', φ') . Typically, $S(\theta, \varphi)$ and $R(\gamma)$ are estimated from the data and modeled in spherical harmonics and rotational

harmonics, respectively. This reduces spherical deconvolution to simple scalar division, and yields $F(\theta, \varphi)$ [9].

It has been pointed out [10] that Equation (2) describes a continuous version of the ball-and-stick model when substituting

$$R(\gamma) = S_0 e^{-bd \cos^2 \gamma} \quad (3)$$

and adding $S_0 f_{\text{iso}} \exp(-bd)$ outside the integral, where f_{iso} is the volume fraction of the ball compartment. We use this relation to compute an ODF $F(\theta, \varphi)$ that corresponds to the continuous ball-and-stick model. FORECAST [14] has followed a similar goal, though corresponding to a multi-tensor model with nonzero perpendicular diffusivity ($\lambda_2 = \lambda_3 > 0$) in the individual compartments.

2.3 Discrete Approximations of Continuous ODFs

From a continuous ball-and-stick model, the discrete case is recovered by replacing the continuous function $F(\theta, \varphi)$ with a discrete ODF $\tilde{F}(\theta, \varphi)$ that is a finite sum of weighted delta peaks. Even though it is common to recover discrete directions by locating maxima in $F(\theta, \varphi)$ [9, 17], this is not accurate. It ignores the fact that delta peaks represented by finite-order spherical harmonics no longer have negligible width, so the individual peaks interfere with each other.

Therefore, we employ nonlinear optimization based on higher-order tensor representations to find a discrete approximation $\tilde{F}(\theta, \varphi)$, as described in [12]. Selecting the maximum spherical harmonics order involves a tradeoff between increasing peak sharpness and reducing the influence of noise. Unlike maximum extraction, the optimization in [12] explicitly accounts for the blurring of ODF peaks at low orders. We found that because of this, a good tradeoff is already achieved at maximum order four.

3 Using Spherical Deconvolution for Model Fitting

3.1 Fitting the Ball-and-Stick Model

The previous section described how fitting the ball-and-stick model can theoretically be formulated as a deconvolution problem with a discrete ODF $\tilde{F}(\theta, \varphi)$. We now describe an efficient and reliable algorithm based on this insight.

Initial per-voxel estimates of the diffusivity d and the isotropic volume fraction f_{iso} of the ball-and-stick model are obtained from the maximum apparent diffusion coefficient d_{max} and the average diffusion-weighted intensity \bar{S}

$$d_{\text{max}} = -\frac{1}{b} \min_i \log \frac{S_i}{S_0} \quad \bar{S} = \frac{1}{n} \sum_{i=1}^n S_i, \quad (4)$$

where S_i are the diffusion-weighted values. Integrating Equation (1) over the unit sphere shows that \bar{S} varies linearly with f_{iso} between $\bar{S} = S_0 \sqrt{\pi/(4bd)} \text{erf}(\sqrt{bd})$ (at $f_{\text{iso}} = 0$) and $\bar{S} =$

$S_0 \exp(-bd)$ (at $f_{\text{iso}} = 1$), where erf is the Gauss error function. Assuming that $d \approx d_{\text{max}}$, this allows us to compute f_{iso} from \bar{S} . These initial estimates are refined further as part of the final descent-based fitting.

Now, the predicted isotropic part $S_0 f_{\text{iso}} \exp(-bd)$ is subtracted from S_i , and a spherical harmonics representation of $S(\theta, \phi)$ is fit to the remainder. Like [12], we deconvolve to a non-ringing cosine power lobe ($\cos^4 \gamma$) instead of a truncated delta peak. Using computer algebra software, the order-four coefficients of the deconvolution kernel that corresponds to $R(\gamma)$ in Equation (3) are found as:

$$\begin{aligned} R_0 &= \frac{5\sqrt{\pi}S_0}{2\sqrt{bd}} \operatorname{erf}(\sqrt{bd}); & R_2 &= \frac{-35S_0}{32\sqrt{(bd)^3}} \left(6\sqrt{bd}e^{-bd} + \sqrt{\pi}(2bd - 3)\operatorname{erf}(\sqrt{bd}) \right) \\ R_4 &= \frac{105S_0}{512\sqrt{(bd)^5}} \left(-30\sqrt{bd}(2bd+21)e^{-bd} + 9\sqrt{\pi}(4bd(bd-5)+35)\operatorname{erf}(\sqrt{bd}) \right) \end{aligned} \quad (5)$$

Dividing the order- n spherical harmonics coefficients of $S(\theta, \phi)$ by R_n gives $F(\theta, \phi)$ [9], which is then approximated by a discrete ODF $\tilde{F}(\theta, \phi)$ with the algorithm in [12]. This involves nonlinear optimization, but it is much simpler than fitting Equation (1) directly: Instead of fitting to 50–100 DWI values, it considers only 15 coefficients of $F(\theta, \phi)$ (at maximum order four). Moreover, both the objective function and its derivatives involve only additions and multiplications.

The peaks in $\tilde{F}(\theta, \phi)$ approximate fiber directions, their weights provide relative volume fractions. In a final step, these estimates are refined by fitting Equation (1) to the original DWI data, using Levenberg-Marquardt (LM)¹ [13].

3.2 Selecting the Number of Sticks

For best quality, the correct number of compartments needs to be selected in the multi-tensor model: If it is chosen too low, the model might indicate fiber directions that do not align with any true tract. On the other hand, overestimating the true number reduces the accuracy of the result due to overfitting.

In spherical deconvolution, it is common to estimate the number of fiber compartments from the maxima in the ODF $F(\theta, \phi)$ [17]. Since our framework uses deconvolution, we explore a similar strategy based on the discrete ODF $\tilde{F}(\theta, \phi)$. We use a three-stage test with thresholds $t_{\{0,1,2\}}$: If $f_{\text{iso}} > t_0$, no fiber is detected. Otherwise, $F(\theta, \phi)$ is normalized to integrate to unity and a discrete ODF with two delta peaks is extracted, whose weights are $w_1 > w_2$. If $w_2 < t_1$, a single fiber is used. Otherwise, a three-peak approximation is found with $w_1 > w_2 > w_3$. If $w_3 < t_2$, we assume two fibers, else three compartments are used.

¹<http://www.ics.forth.gr/~lourakis/levmar/> [Accessed on 22 Feb 2010]

4 Results

4.1 Synthetic Data

Synthetic diffusion-weighted data with 60 directions, $b = 3000 \text{ s/mm}^2$, and Rician noise at $\text{SNR}_0 = 30$ was created from a multi-cylinder model with non-zero $\lambda_2 = \lambda_3$. Eigenvalues were sampled from Gaussian distributions with parameters according to estimates from real DW-MRI data. Principal eigenvectors and volume fractions were sampled uniformly at random. Cases in which any volume fraction was below 0.2 or any pair of vectors was closer than 30° were rejected.

We created 5000 one-, two-, and three-fiber configurations each and used Levenberg-Marquardt with 100 random restarts, sampled in analogy to data generation, to find the global optimum for the predetermined, correct number of compartments. We then counted the number of times a single randomized run and a run that was kick-started by spherical deconvolution found the correct optimum. Table 1 shows that in the two- and three-fiber cases, the prediction raised the chances of finding the optimum by around 10%. On average, the cost of deconvolution, finding the discrete ODF, and final refinement by LM, is about half the computational cost of finding the optimum by LM alone.

Table 2 lists the mean and median angular deviation of the individual stick compartments, sorted by volume fraction, from the ground truth. It confirms that spherical deconvolution comes close enough to the final result to be a useful seed, but subsequent optimization still improves upon its accuracy, especially in the three-fiber case. To validate our choice of [12] over ODF peak finding, we also list the most accurate result we were able to achieve using that more traditional technique (order 6, truncated delta peak, Laplace-Beltrami regularization [17]).

Finally, we used the synthetic data to compare our criterion for selecting the fiber number to two statistical tests previously applied to diffusion model selection, the Bayesian Information Criterion (BIC) [18], and the F-Test [19, 6, 18]. Automatic relevance determination [11] is not included in our comparison, since it aims at Bayesian model averaging rather than at making a hard decision. As shown in Table 3, the parameter-free BIC exhibited a strong bias towards selecting two or three fibers. The thresholds $t_{\{0,1,2\}}$ of the other two tests were set to balance sensitivity and specificity at each stage. The deconvolution-based test achieved best results for all ground truth configurations.

We also repeated all experiments with $b = 1000 \text{ s/mm}^2$. Without providing the quantitative results, we qualitatively state that the probability of finding the correct optimum increased, but average accuracy with respect to ground truth decreased. Our test for model selection became less reliable, while results of the F-Test improved. Apparently, the F-Test benefits from a low effective noise level, while the deconvolution-based test requires the better separation of the individual compartments afforded by higher b values.

4.2 Real Data

In order to identify the voxels in a real dataset (60 directions, $b = 1000 \text{ s/mm}^2$, isotropic voxel size 1.72 mm) in which the individual tests are most likely to use two- and three-fiber

models, we calibrated the F-Test and our deconvolution-based test to generate 25% no-fiber, 40% one-fiber, 25% two-fiber, and 10% three-fiber voxels within a brain mask.

A detail of the result is shown in Figure 2, the intersection of corpus callosum and corticospinal tract on a coronal slice. For visualization, we map the stick compartments to tensors whose perpendicular diffusivity is scaled by the volume fraction of the ball compartment ($\lambda_2 = \lambda_3 = f_{iso}d$). The fiber fractions are renormalized to unity and color-coded. In (a), standard single diffusion tensors are shown for reference. We found that the F-Test tends to fit multiple stick compartments to voxels that likely just contain a single bending or spreading bundle, e.g., in the body of the corpus callosum (CC). The deconvolution-based test is more robust to such cases, as long as they lead to a clear single ODF peak. We expect that this will allow us to track through crossings like the one presented in Figure 2, while avoiding spurious tracts. However, design of a tractography method that makes use of our novel framework is left as future work.

With the deconvolution-based test, Levenberg-Marquardt is only run once per voxel, which provides an additional speedup. Our complete pipeline, including testing and fitting the final result, processed 1050 voxels per second with the F-Test, 2950 voxels per second with the deconvolution-based test, on a single CPU core of a 2.7 GHz workstation. Finally, 20 randomized restarts of Levenberg-Marquardt improved upon the optimum found by our method in less than 0.2% of all voxels. Therefore, we conclude that our pipeline offers a reliable and efficient solution for fitting ball-and-stick models.

5 Conclusion

Traditionally, multi-fiber models and spherical deconvolution are used as competing methods, each with its own set of advantages and disadvantages: Linear spherical deconvolution is extremely fast and does not require pre-specification of an expected number of fibers. On the other hand, multi-tensor models offer higher accuracy for applications like multi-fiber streamline tractography [4, 7, 17, 15], where it is the primary goal to estimate the most likely fiber directions.

We have presented a framework that combines the best of both worlds: Based on spherical deconvolution, a plausible number of fiber compartments is found automatically. Initializing the fitting with the deconvolution result doubles the speed of the computation, while at the same time increasing the probability of finding the global optimum to more than 95% in the two-fiber case and to more than 90% for three fibers. As a result, we achieve a fully integrated, reliable and efficient algorithmic solution for multi-tensor fitting.

References

1. Basser P, Mattiello J, Bihan DL. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J Magn Reson.* 1994; B(103):247–254.
2. Alexander A, Hasan K, Lazar M, Tsuruda J, Parker D. Analysis of partial volume effects in diffusion-tensor MRI. *Magn Reson Med.* 2001; 45:770–780. [PubMed: 11323803]
3. Tuch D, Reese T, Wiegell M, Makris N, Belliveau J, Wedeen V. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn Reson Med.* 2002; 48:577–582. [PubMed: 12353272]

4. Parker, G.; Alexander, D. *Inf Proc Med Imaging*. Volume 2732 of LNCS. Springer; 2003. Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information; p. 684-695.
5. Peled S, Friman O, Jolesz F, Westin CF. Geometrically constrained two-tensor model for crossing tracts in DWI. *Magn Reson Imaging*. 2006; 24(9):1263–1270. [PubMed: 17071347]
6. Nedjati-Gilani S, Parker G, Alexander DC. Mapping the number of fibre orientations per voxel in diffusion MRI. *Proc Intl Soc Mag Reson Med*. 2006:3169.
7. Kreher B, Schneider J, Mader I, Martin E, Hennig J, Il'yasov K. Multi-tensor approach for analysis and tracking of complex fiber configurations. *Magn Reson Med*. 2005; 54:1216–1225. [PubMed: 16200554]
8. Pasternak O, A Y, Intrator N, Sochen N. Variational multiple-tensor fitting of fiber-ambiguous diffusion-weighted magnetic resonance imaging voxels. *Magn Reson Imaging*. 2008; 26:1133–1144. [PubMed: 18524529]
9. Tournier JD, Calamante F, Gadian D, Connelly A. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*. 2004; 23:1176–1185. [PubMed: 15528117]
10. Behrens T, Woolrich M, Jenkinson M, Johansen-Berg H, Nunes R, Clare S, Matthews P, Brady J, Smith S. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med*. 2003; 50:1077–1088. [PubMed: 14587019]
11. Behrens T, Johansen-Berg H, Jbabdi S, Rushworth M, Woolrich M. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*. 2007; 34:144–155. [PubMed: 17070705]
12. Schultz T, Seidel HP. Estimating crossing fibers: A tensor decomposition approach. *IEEE Trans Vis Comp Graphics*. 2008; 14(6):1635–1642.
13. Marquardt D. An algorithm for least-squares estimation of nonlinear parameters. *SIAM J Applied Math*. 1963; 11(2):431–441.
14. Anderson A. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magn Reson Med*. 2005; 54(5):1194–1206. [PubMed: 16161109]
15. Malcolm J, Michailovich O, Bouix S, Westin CF, Shenton M, Rathi Y. A filtered approach to neural tractography using the Watson directional function. *Medical Image Analysis*. 2010; 14:58–69. [PubMed: 19914856]
16. Hosey T, Williams G, Ansorge R. Inference of multiple fiber orientations in high angular resolution diffusion imaging. *Magn Reson Med*. 2005; 54:1480–1489. [PubMed: 16265642]
17. Descoteaux M, Deriche R, Knösche T, Anwander A. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Trans Med Imaging*. 2009; 28(2):269–286. [PubMed: 19188114]
18. Freidlin R, Özarlan E, Komlosh M, Chang LC, Koay C, Jones D, Basser P. Parsimonious model selection for tissue segmentation and classification applications: A study using simulated and experimental DTI data. *IEEE Trans Med Imaging*. 2007; 26(11):1576–1584. [PubMed: 18041272]
19. Alexander D, Barker G, Arridge S. Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Magn Reson Med*. 2002; 48:331–340. [PubMed: 12210942]




Fig. 1. In our framework, deconvolution and ODF approximation make Levenberg-Marquardt fitting of ball-and-stick models faster and more reliable.

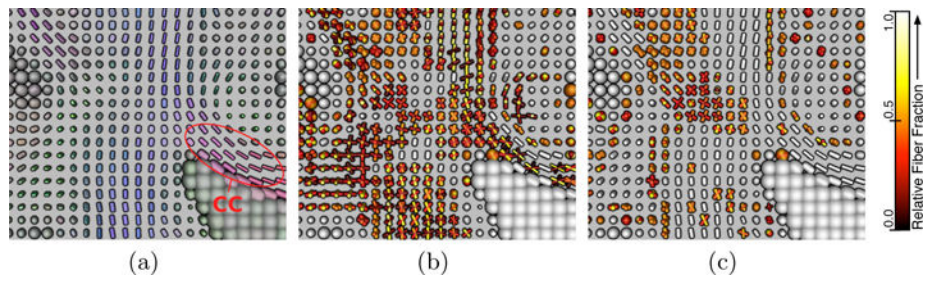


Fig. 2.

In real data (a), our deconvolution-based test (c) produced more plausible results than the F-Test (b), which frequently fits multiple sticks in regions that are generally thought to contain a single bending or spreading compartment.

Percentage of cases in which the global optimum was found by a single randomized run of LM, versus our kick-start method, with “speed” in voxels per second.

Table 1

	1 Fiber		2 Fibers		3 Fibers	
	% Found	Speed	% Found	Speed	% Found	Speed
Randomized Run	99.9	3430	90.3	1496	82.6	805
Kick-Started Run	100	5635	98.6	3264	93.1	1527

Mean and median (in parentheses) of angular deviation from ground truth confirms that spherical deconvolution seeds the optimization close to the final result.

Table 2

	1 Fiber	2 Fibers		3 Fibers		
	Angle	Angle 1	Angle 2	Angle 1	Angle 2	Angle 3
Kick-Started Run	1.2 (1.0)	2.3 (1.9)	5.9 (3.9)	3.9 (3.0)	6.8 (4.9)	8.8 (6.3)
Prediction alone	1.3 (1.1)	3.0 (2.3)	7.5 (5.0)	5.8 (4.4)	9.5 (7.3)	12.0 (9.6)
ODF Peaks	1.5 (1.3)	3.5 (2.6)	12.4 (6.3)	5.7 (4.2)	12.4 (8.1)	16.2 (11.1)

In synthetic data, our deconvolution-based test for the number of fiber compartments produces better results than standard statistical tests

Table 3

Truth	BIC				Prediction F-Test				Our Test			
	0	1	2	3	0	1	2	3	0	1	2	3
No Fiber	98%	2%	0%	0%	98%	2%	0%	0%	100%	0%	0%	0%
1 Fiber	0%	29%	43%	28%	0%	82%	6%	12%	0%	94%	5%	1%
2 Fibers	0%	5%	74%	21%	0%	15%	53%	32%	0%	5%	83%	12%
3 Fibers	0%	5%	23%	73%	2%	31%	13%	54%	0%	0%	16%	84%