

Regularized, Fast, and Robust Analytical Q-Ball Imaging

Maxime Descoteaux,^{1*} Elaine Angelino,² Shaun Fitzgibbons,² and Rachid Deriche¹

We propose a regularized, fast, and robust analytical solution for the Q-ball imaging (QBI) reconstruction of the orientation distribution function (ODF) together with its detailed validation and a discussion on its benefits over the state-of-the-art. Our analytical solution is achieved by modeling the raw high angular resolution diffusion imaging signal with a spherical harmonic basis that incorporates a regularization term based on the Laplace–Beltrami operator defined on the unit sphere. This leads to an elegant mathematical simplification of the Funk–Radon transform which approximates the ODF. We prove a new corollary of the Funk–Hecke theorem to obtain this simplification. Then, we show that the Laplace–Beltrami regularization is theoretically and practically better than Tikhonov regularization. At the cost of slightly reducing angular resolution, the Laplace–Beltrami regularization reduces ODF estimation errors and improves fiber detection while reducing angular error in the ODF maxima detected. Finally, a careful quantitative validation is performed against ground truth from synthetic data and against real data from a biological phantom and a human brain dataset. We show that our technique is also able to recover known fiber crossings in the human brain and provides the practical advantage of being up to 15 times faster than original numerical QBI method. Magn Reson Med 58:497–510, 2007. © 2007 Wiley-Liss, Inc.

Key words: diffusion tensor imaging (DTI); high angular resolution diffusion imaging (HARDI); Q-ball imaging (QBI); fiber tractography, orientation distribution function (ODF); regularization; Funk Radon transform; spherical harmonic

INTRODUCTION

Diffusion MRI (1,2) is a noninvasive tool widely used in medical imaging applications to obtain information about the neural architecture in vivo and to understand functional coupling between cortical regions of the brain. Because of the well-known limitations of diffusion tensor imaging (DTI) in regions of low anisotropy and multiple fiber crossings, recent works have attempted to generalize the existing DTI model (3) with new higher resolution acquisition techniques such as Q-space imaging (QSI) (4,5), high angular resolution diffusion imaging (HARDI) (6) and Q-Ball Imaging (QBI) (6). There are currently two classes of high order processing methods for these high resolution acquisition techniques. The first is based on apparent diffusion coefficient (ADC) modeling (7–13) and the other is based on the estimation of the probability density function (PDF) of the average spin displacement of water molecules

or variants of this function (14–29). This article focuses on the second class of methods.

Existing techniques in this class approximate the PDF or variants of it arising from various types of high angular resolution signal data. The existing functions in the literature are the persistent angular structure (PAS) (16) of the PDF, the fiber orientation distribution (FOD) (18,19), diffusion orientation transform (DOT) (27), and the diffusion orientation distribution function (ODF) (17). We mainly focus on the diffusion ODF, which captures important angular information and can be viewed as the probability that the water molecule will diffuse into any solid angle. For all these high angular resolution functions, the important property is that their maxima agree with the underlying fiber distribution. However, these methods are based on numerical methods, lack a straightforward regularization process and fail to take into account the useful tools for both estimation and regularization that have been developed for the fitting of the ADC profile (13).

In this paper, the backbone tool used to analyze HARDI data is the spherical harmonic (SH). Frank (8), Alexander et al. (7), Zhan et al. (10), Tournier et al. (18), Ozarslan et al. (27), Anderson (23), Hess et al. (28,29), and others seem to have converged to this methodological tool which has many powerful properties. This is reassuring and SH seem to be a natural way to decompose signals that live on the sphere, just as the Fourier transform is the widespread tool used to decompose images. Hence, it is not surprising that Anderson (23), Hess et al. (28,29), and our group (25,26) have recently developed separately and in parallel an analytical solution for the ODF reconstruction in QBI. Despite the fact that the three analytical solutions are similar, the regularized estimation part, the derivation, the experimental results, and the validation phase are quite different. In the rest of the article, we carefully present our work and put up front the differences in derivations and contributions between our proposed method and the two papers (23,28).

More precisely, our solution is obtained by modeling the signal with high order SH series using a Laplace–Beltrami regularization method developed for the ADC profile estimation in (13). This leads to an elegant mathematical simplification of the Funk–Radon transform which approximates the ODF. We prove a new corollary of the Funk–Hecke theorem to obtain this simplification. We then obtain a fast algorithm for the extraction of a robust regularized model-independent ODF approximation at each voxel of the raw HARDI data, which offers advantages over previous numerical approaches. Another contribution is its quantitative validation on synthetic data and on real data obtained from a biological rat phantom and from a human brain dataset. Overall, the contributions of the paper are threefold: (1) Our ODF estimation is up to 15 times faster than Tuch's numerical method. (2) Our ODF estimation is regularized with the Laplace–Beltrami operator, which is theoretically and practically better than

¹Odyssée Project Team, INRIA/ENPC/ENS, INRIA Sophia Antipolis, France

²Division of Engineering and Applied Sciences, Harvard University, Boston, Massachusetts, USA.

Grant sponsors: CRSNG Canada graduate scholarship; the FQRNT-INRIA; the INRIA (International internships program)

*Correspondence to: Maxime Descoteaux, INRIA, Odyssée Project Team, 2004 route des Lucioles, 06902 Sophia Antipolis, France. E-mail: maxime.descoteaux@sophia.inria.fr

Received 21 April 2006; revised 1 March 2007; accepted 2 April 2007.

DOI 10.1002/mrm.21277

Published online in Wiley InterScience (www.interscience.wiley.com).

Tikhonov regularization (28) as well as more robust to noise. (3) At the cost of slightly reducing angular resolution, our ODF estimation reduces errors and improves fiber detection while reducing angular error in the ODF maxima detected.

THEORY

Q-Ball Imaging

QBI (6,17), introduced by D. Tuch, is a HARDI technique that reconstructs the diffusion ODF of the underlying fiber population of a biological tissue. QBI aids in the inference of fiber bundles with crossing, kissing or diverging configurations, with advantages over DTI in these situations. The ODF is also intuitive and gives a good representation of the underlying fiber distribution. The ODF in a unit direction \mathbf{u} , $\Psi(\mathbf{u})$, is given by the radial projection of the diffusion PDF in q -space, where $q = \mathbf{q}/|\mathbf{q}|$ is given by $\mathbf{q} = \gamma\delta\mathbf{G}/2\pi$, with γ the nuclear gyromagnetic ratio, \mathbf{G} the applied diffusion gradient vector and δ the gradient pulse duration. Tuch (17) showed that this ODF could be estimated directly from the raw HARDI signal S on a single sphere of q -space by the Funk–Radon transform (FRT) \mathcal{G} . This FRT is essentially a smoothed version of the true ODF (17). In practice, the FRT value at a given spherical point \mathbf{u} is the great circle integral of the signal on the sphere defined by the plane through the origin with normal vector \mathbf{u} . This can be written as

$$\mathcal{G}[S](\mathbf{u}) = \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) S(\mathbf{w}) d\mathbf{w}. \quad [1]$$

Figure 1 shows the qualitative effect of the Funk–Radon transform on 1 and 2 fiber examples of synthetic signals.

The QBI reconstruction has several advantages that have made it a popular high angular resolution reconstruction in recent works (17,22–25,28,30–32) for fiber tracking and for characterizing white matter architecture. In theory, samples are only taken on a single sphere in q -space and thus, the imaging time is much smaller than that of the QSI despite significantly higher angular resolution measurements. Furthermore, since a relatively small constant value of q can be chosen, the signal to noise ratio is greatly improved. Finally, QBI also has the advantage of being model-independent which is not the case for multi-fiber Gaussian models (14) and spherical deconvolution approaches (18,19) where a fiber response function needs to be assumed a priori to use a deconvolution kernel. In (19), the response function is a standard Gaussian kernel whereas in (18), the response function is estimated from the signal attenuation profile in 300 voxels with highest fractional anisotropy, that is regions likely to have a single coherent fiber population.

Analytical Q-Ball Imaging

As mentioned in the introduction, we are proposing a similar analytical QBI solution for the ODF reconstruction as Anderson (23) and Hess et al. (28,29). However, despite the similarity of the analytical solution, the derivations are

quite different and use other properties of the SH basis that are worth reviewing.

First, Anderson (23, Appendix B) proposes an analytical solution using SH. It is not the focus of the article and this is the reason why Anderson has not studied the full potential of the solution as Hess et al. did and as we do in this article. Nonetheless, the derivation is simple and different from that of both Hess et al. and ours. The idea is that the signal on the sphere can be represented in terms of the standard polar angles (θ, ϕ) or in terms of coordinates in a rotated frame. Using the Wigner rotation matrix, there is a simple relation between SHs evaluated at (θ, ϕ) and SHs evaluated at the same physical angle but expressed in the rotated frame. This is used to evaluate and simplify the FRT by expressing the integral in a carefully chosen rotated frame.

Second, the solution of Hess et al. (28,29) is based on the analytical solution to great circle integrals over SHs. This is based on a previous work by Backus (33) for an application to geophysical data. Two important properties of the SH are exploited in this proof: (1) any rotated SH can be uniquely expressed as a linear combination of SH of the same degree and (2) SH satisfy the addition theorem, that is for any two directions (θ, ϕ) and (θ', ϕ') separated on the sphere by the angle (Θ, Φ) , $Y_\ell^0(\theta', 0) = (1/Y_\ell^0(0, 0)) \sum_{m=-\ell}^{\ell} \tilde{Y}_\ell^m(\Theta, \Phi) Y_\ell^m(\theta, \phi)$, where ℓ denotes the order of the SH.

In our case, we use the 3D Funk–Hecke theorem to analytically evaluate integrals of functions on the sphere. Before describing this new derivation, we need to review the main steps of some of our previous work (13) in which we developed a simple and efficient way to estimate the continuous function on a q-ball that best approximates discrete and sparse measurements using SHs. This SH formulation is crucial for our analytical simplification of QBI and allows the introduction of the Laplace–Beltrami regularization criterion for the solution.

Signal Approximation with the Spherical Harmonics

The SH, normally indicated by Y_ℓ^m (ℓ denotes the order and m the phase factor), are a basis for complex functions on the unit sphere. Explicitly, they are given as follows

$$Y_\ell^m(\theta, \phi) = \sqrt{\frac{2\ell + 1}{4\pi} \frac{(\ell - m)!}{(\ell + m)!}} P_\ell^m(\cos \theta) e^{im\phi}, \quad [2]$$

where (θ, ϕ) obey physics convention ($\theta \in [0, \pi]$, $\phi \in [0, 2\pi]$) and P_ℓ^m is an associated Legendre polynomial. For $k = 0, 2, 4, \dots, \ell$ and $m = -k, \dots, 0, \dots, k$, we define the new index $j := j(k, m) = (k^2 + k + 2)/2 + m$ and define our modified basis \mathbf{Y} with elements Y_j such that

$$Y_j = \begin{cases} \sqrt{2} \cdot \text{Re}(Y_k^m), & \text{if } -k \leq m < 0 \\ Y_k^0, & \text{if } m = 0 \\ \sqrt{2} \cdot \text{Im}(Y_k^m), & \text{if } 0 < m \leq k \end{cases}, \quad [3]$$

where $\text{Re}(Y_\ell^m)$ and $\text{Im}(Y_\ell^m)$ represent the real and imaginary parts of Y_ℓ^m respectively. The basis is designed to

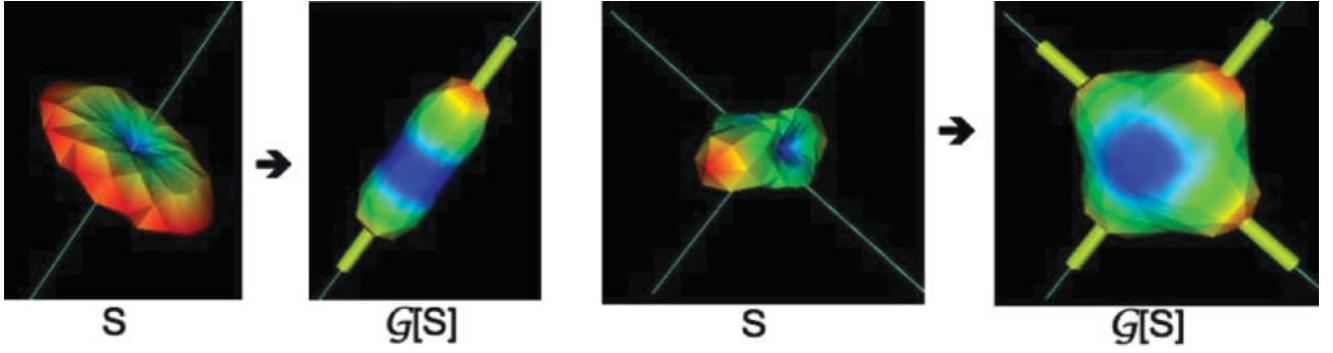


FIG. 1. Funk–Radon Transform illustrated for the input diffusion attenuation signal $S(b = 1,000 \text{ s/mm}^2)$ with 1 fiber (left) and two orthogonal fibers (right). The thin lines are the true underlying fiber directions and the thicker tubes are the detected maxima. One must imagine these functions as living on the surface of the sphere. Here, for visualization purposes, the radius of the respective spheres are scaled by the corresponding value on the surface. Blue to red colors represent low to high spherical values.

be symmetric, real and orthonormal. Symmetry is ensured by choosing only even order SH and the ratios in front of each term also ensure that the modified basis is real and orthonormal with respect to the inner product $\langle f, g \rangle = \int_{\Omega} \bar{f} g d\Omega$, where Ω denotes integration over the unit sphere and \bar{f} is the complex conjugate of f . We thus approximate the signal at each of the N gradient directions i as

$$S(\theta_i, \phi_i) = \sum_{j=1}^R c_j Y_j(\theta_i, \phi_i), \quad [4]$$

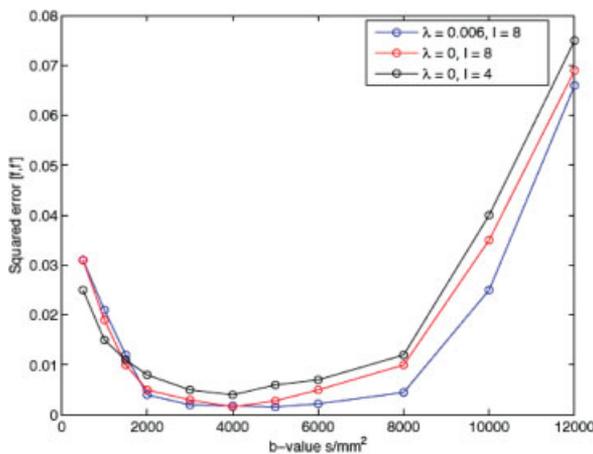
where $R = (\ell + 1)(\ell + 2)/2$ is the number of terms in the modified SH basis \mathbf{Y} of order ℓ . Letting \mathbf{S} be the $N \times 1$ vector representing the input signal for every encoding gradient direction, \mathbf{C} the $R \times 1$ vector of SH coefficients c_j and \mathbf{B} is the $N \times R$ matrix constructed with the discrete modified

SH basis

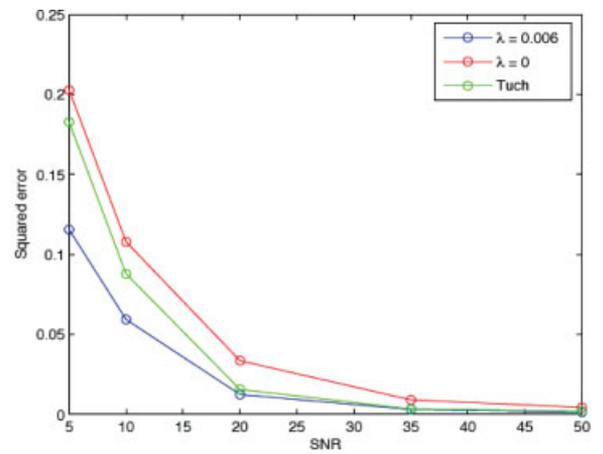
$$\mathbf{B} = \begin{pmatrix} Y_1(\theta_1, \phi_1) & Y_2(\theta_1, \phi_1) & \cdots & Y_R(\theta_1, \phi_1) \\ \vdots & \vdots & \ddots & \vdots \\ Y_1(\theta_N, \phi_N) & Y_2(\theta_N, \phi_N) & \cdots & Y_R(\theta_N, \phi_N) \end{pmatrix}, \quad [5]$$

we can write the set of equations as an over-determined linear system $\mathbf{S} = \mathbf{BC}$. We want to solve for the SH series coefficients c_j , where $c_j = \int_{\Omega} S(\theta, \phi) Y_j(\theta, \phi) d\Omega$.

At this point, instead of simply evaluating the integrals directly as done by Frank in (8) or performing a straightforward least-squared minimization as in (7,18), we add local regularization directly into our fitting procedure. This is to be able to use a high order estimation without over-modeling the small perturbations because of noise in the input diffusion MRI signal. We thus define a measure, E , of the deviation from smoothness of a function f defined



(a)



(b)

FIG. 2. Laplace–Beltrami regularization reduces ODF estimation errors. Plotted is the point-wise Euclidean squared error $[f, f']$ of Eq. [19] between the ground truth f and estimated normalized ODFs f' . (a) In the first simulation, the SNR was fixed to 35 while varying the b -value between 500 and 12,000 s/mm^2 for combinations of regularization parameter and estimation order of $(\lambda = 0.006, \ell = 8)$ in blue, $(\lambda = 0, \ell = 8)$ in red and $(\lambda = 0, \ell = 4)$ in black. (b) In the second simulation, we fix the b -factor to $b = 3,000 \text{ s/mm}^2$ and estimation order to $\ell = 8$ while varying the SNR between 5 and 50. Here, the point-wise Euclidean squared error $[f, f']$ is lower with $\lambda = 0.006$ in blue than with $\lambda = 0$ in red and Tuch's numerical method in green. In both tests, we randomly choose the number of fibers n per voxel to be 1, 2, or 3.

on the unit sphere as $E(f) = \int_{\Omega} (\Delta_b f)^2 d\Omega$, where Δ_b is the Laplace–Beltrami operator. The Laplace–Beltrami operator, which is the extension of the Laplacian operator to functions defined on surfaces, is a natural measure of smoothness for functions defined on the unit sphere. If the spherical function f is parameterized with SH, the Laplace–Beltrami operator is very simple to evaluate when acting on this parameterization as it satisfies the relation $\Delta_b Y_\ell^m = -\ell(\ell + 1)Y_\ell^m$. Note that this relation also holds for our modified SH basis \mathbf{Y} . Using the orthonormality of the modified SH basis, the above functional E can be rewritten straightforwardly as

$$\begin{aligned} E(f) &= \int_{\Omega} \Delta_b \left(\sum_p c_p Y_p \right) \Delta_b \left(\sum_q c_q Y_q \right) d\Omega \\ &= \sum_{j=1}^R c_j^2 \ell_j^2 (\ell_j + 1)^2 = \mathbf{C}^T \mathbf{L} \mathbf{C}, \end{aligned} \quad [6]$$

where \mathbf{L} is simply the $R \times R$ matrix with entries $\ell_j^2 (\ell_j + 1)^2$ along the diagonal (ℓ_j is the order associated with the j th coefficient). We thus obtain a closed-form expression for the regularization term. Therefore, the quantity we wish to minimize can be expressed in matrix form as

$$\mathbf{M}(\mathbf{C}) = (\mathbf{S} - \mathbf{B}\mathbf{C})^T (\mathbf{S} - \mathbf{B}\mathbf{C}) + \lambda \mathbf{C}^T \mathbf{L} \mathbf{C}, \quad [7]$$

where λ is the weight on the regularization term. The coefficient vector minimizing this expression can then be determined just as in the standard least-squares fit ($\lambda = 0$), from which we obtain the generalized expression for the desired SH series coefficient vector

$$\mathbf{C} = (\mathbf{B}^T \mathbf{B} + \lambda \mathbf{L})^{-1} \mathbf{B}^T \mathbf{S}. \quad [8]$$

From this SH coefficient vector we can recover the signal on the q-ball for any (θ, ϕ) as

$$S(\theta, \phi) = \sum_{j=1}^R c_j Y_j(\theta, \phi). \quad [9]$$

Intuitively, this approach penalizes an approximation function for having higher order terms in its modified SH series. Therefore, higher order terms will only be included in the fit if they significantly improve the overall accuracy of the approximation. This eliminates most of the high order terms due to noise while leaving those that are necessary to describe the underlying function. However, obtaining this balance depends on choosing a good value for the parameter λ . We use the *L-curve* numerical method (34) and experimental simulations to determine a good smoothing parameter.

Recently, Hess et al. (28,29) proposed to use a Tikhonov regularization method while solving for the SH coefficients when a large harmonic order approximation is wanted ($\ell > 4$) and standard $\ell = 4$ estimation otherwise. It is a simple technique used to numerically better condition the matrix involved in the pseudo-inverse of Eq. [8]. The solution is expressed in the same form as Eq. [8] but the Laplace–Beltrami matrix \mathbf{L} is replaced with the identity matrix \mathbf{I} . The regularization is thus uniform, that is lower

and higher order coefficients are weighted in the same way which is less desirable than a smoothing that minimizes perturbing effects occurring mainly at higher order harmonics. Moreover, the underlying assumption of Tikhonov regularization is that the space where the data lives is a flat manifold. This is a rough approximation in the ODF reconstruction problem since the data actually lives on a sphere, for which the right and appropriate tool to use is the Laplace–Beltrami operator. Finally, other works propose some postprocessing of the ODF to reduce errors involved in high frequency harmonics. In Tournier et al. (18), higher order terms of the SH fiber ODF reconstruction are attenuated by a heuristic amount. In Tuch (17), an isotropic spherical smoothing filter of a certain width is applied to the reconstructed ODFs.

Funk-Radon Transform Using Spherical Harmonics

We now prove a corollary to the Funk-Hecke theorem needed to solve the FRT integral. Here, we write the dot product between two vectors $\mathbf{u}, \mathbf{w} \in \mathcal{R}^3$ as $\mathbf{u}^T \mathbf{w}$. The key observation is that any continuous function f on the interval $[-1, 1]$ extends to a continuous function of two variables $g(\mathbf{u}, \mathbf{w})$ on the sphere defined by $g(\mathbf{u}, \mathbf{w}) = f(\mathbf{u}^T \mathbf{w})$. With this formulation, the Funk–Hecke formula is a theorem that relates the inner product of any SH with the projection on the sphere of any continuous function $f(t)$ defined on the interval $[-1, 1]$. The theorem was first published by Funk in 1916 and by Hecke in 1918 (35, chapter 9) and here, we give only its 3D version.

Funk-Hecke Theorem: *Let $f(t)$ be continuous on $[-1, 1]$ and H_ℓ any SH of order ℓ in C^3 , the space of 3D complex functions. Then, given a unit vector \mathbf{u}*

$$\int_{|\mathbf{w}|=1} f(\mathbf{u}^T \mathbf{w}) H_\ell(\mathbf{w}) d\mathbf{w} = \lambda(\ell) H_\ell(\mathbf{u}), \quad [10]$$

where

$$\lambda(\ell) = 2\pi \int_{-1}^1 P_\ell(t) f(t) dt$$

with P_ℓ the Legendre polynomial of degree ℓ .

In our case, recall that the Funk-Radon transform of the signal in a unit direction \mathbf{u} is the integral over the great circle perpendicular to \mathbf{u} , as stated in Eq. [1]. Hence, replacing the signal by the SH series approximations of Eq. [9], we have

$$\begin{aligned} \mathcal{G}[S](\mathbf{u}) &= \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) S(\mathbf{w}) d\mathbf{w} \\ &= \int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) \sum_{j=1}^R c_j Y_j(\mathbf{w}) d\mathbf{w} \\ &= \sum_{j=1}^R c_j \underbrace{\int_{|\mathbf{w}|=1} \delta(\mathbf{u}^T \mathbf{w}) Y_j(\mathbf{w}) d\mathbf{w}}_{I_j} \end{aligned} \quad [11]$$

Note that if the Dirac delta function δ was continuous on the interval $[-1, 1]$, I_j could be directly evaluated using the Funk–Hecke formula of Eq. [10]. However, $\delta(t)$ is discontinuous at zero. To overcome this problem, we approximate the delta function with a delta sequence $\delta_n(x)$ and take the

the modified SH series approximation with the Laplace–Beltrami smoothing was established in Ref. 13 for sparse measurements on the sphere. The analytical ODF estimation method offers the advantages that the discrete interpolation over many equators is eliminated and the solution for all directions is obtained in a single step (Eq. [17]). We now validate our new analytical QBI method and analyze the computational complexity gain. We also point out the different and complementary results obtained in Hess et al. (28).

METHODS

Numerical QBI Implementation

The implementation (17) involves several numerical computations such as a regridding to find points outside the actual measurements required to compute the discrete points on each great circle. Our implementation of Tuch’s numerical QBI is adapted from Campbell et al.’s. (22) QBI computation which is successfully used for multiple fiber characterization and tracking. Given N points spaced approximately uniformly on the surface of the sphere, the interpolation on each unit great circle can be done at every $(1/2)\sqrt{2\pi/N}$ radians to take full advantage of the N discrete measurements. This gives $k = \sqrt{8\pi N}$ points per equator. We also set the width parameter (angular width) of the spherical Gaussian interpolation kernel automatically to three times the angle between equator points, i.e. $\sigma = (3/2)\sqrt{2\pi/N}$.¹ This heuristic choice gives a good trade-off between the accuracy and stability in our experiments.

Synthetic Data Generation

We generate synthetic data using the multi-tensor model (7, 17,28,36,36), which leads to an analytical computation of the exact ODF. For a given b -factor, noise level and encoding direction i , we generate the diffusion-weighted signal $S(\mathbf{u}_i) = \sum_{k=1}^n p_k e^{-b\mathbf{u}_i^T \mathbf{D}_k \mathbf{u}_i} + \text{noise}$, where $i = 1, \dots, N$ for $N = 81$ or $N = 321$ gradient directions on the hemisphere (for 3rd or 7th order tessellation of the icosahedron respectively), n is the number of fibers, p_k is the proportion of the k th fiber and \mathbf{D}_k the k th diffusion tensor profile with eigenvalues $[300, 300, 1,700] \times 10^{-6}$ mm²/s (Fractional Anisotropy (FA) = 0.8) oriented randomly. In practice, we impose a minimum random angle between fibers of 45° and relative random weights between 0.3 and 0.7 for 2-fiber and between 0.2 and 0.4 for 3-fiber distributions. The noise is generated with a complex Gaussian noise with a standard deviation of σ , producing a S_0 signal with SNR = $1/\sigma$, that is we define SNR as the ratio of maximum signal intensity of S_0 to the standard deviation σ of the complex Gaussian noise. As in (28), we vary SNR values between 5 and 50. This SNR range covers expected low to high quality of in vivo HARDI data. The step by step procedure to generate this synthetic data is extensively detailed in (13).

The exact ODF is given in Tuch (17) for a single Gaussian fiber and by linearity, we can easily obtain the exact ODF

for n fibers and normalization constant Z (17),

$$\Psi(\mathbf{u}_i) = \sum_{k=1}^n \frac{p_k}{Z} \sqrt{\frac{\pi b}{\mathbf{u}_i^T \mathbf{D}_k^{-1} \mathbf{u}_i}}. \quad [18]$$

Diffusion MRI Data Acquisitions

A Biological Phantom

The rat phantom was produced by Campbell et al. at the McConnell Brain Imaging Center and Montreal Neurological Institute (22) on a 1.5T Sonata MR scanner using a knee coil. It was created from two excised Sprague-Dawley rat spinal cords embedded in 2% agar. The acquisition was done with a single-shot spin-echo planar sequence with twice-refocused balanced gradients, designed to reduce eddy current effects. The dataset is acquired with 90 pairs of points generated using electrostatic repulsion algorithm (37). The q-ball protocol was used with $b = 3,000$ s/mm², $q = 0.35$ μm⁻¹, TR = 6.4 s, TE = 110 ms, FOV 360 × 360 mm², 128 × 128 matrix, 2.8 mm isotropic voxels and four signal averages per direction. The SNR of the $b = 0$ image S_0 was estimated to be ~70 for the averaged phantom and around 10 for the cord at $b = 3,000$ s/mm².

Human Brain Data

We performed ODF reconstructions on two real human brain datasets, one with a high b -value of 3,000 s/mm² and the other with a more standard clinical b -value of 1,000 s/mm². The first dataset, brain#1, was acquired using a Siemens 3T MR scanner 8 channel head coil, using 99 pairs of points generated with the electrostatic repulsion algorithm (37), has 63 slices of 2 mm each, covering the entire cerebrum. The FOV was 256 × 256 mm², TR = 11.1 s, TE = 121 ms, $b = 3,000$ s/mm² ($q = 0.35$ μm⁻¹), BW 1346 Hz/Pixel, 128 × 128 matrix and phase partial Fourier 7/8. Ten $b = 0$ images were acquired and averaged to produce the S_0 image. The SNR in the white matter of this S_0 image was then estimated to be ~41. In the second dataset, brain#2, diffusion weighted images were acquired at the Leipzig Max Planck Institute on a whole-body 3 Tesla Magnetom Trio scanner (Siemens, Erlangen) equipped with an 8-channel head array coil (38). The spin-echo EPI sequence, TE = 100 ms, TR = 12 s, 128 × 128 image matrix, FOV = 220 × 220 mm², consists of 60 evenly distributed diffusion encoding gradients (37) with a b -value of 1000 s/mm² and 7 images without any diffusion weightings. The measurement of 72 slices with 1.7 mm thickness (no gap), which covered the whole brain, was repeated three times, resulting in an acquisition time of about 45 min. Hence, the S_0 image is the average of 21 $b = 0$ images. The SNR in the white matter of this S_0 image was estimated to be ~37. Additionally, fat saturation was employed, 6/8 partial Fourier imaging, Hanning window filtering, and parallel GRAPPA imaging with a reduction factor of 2.

Validation

Computational Complexity Analysis

We refer to Table 1 for the analytical QBI and to (17, Table 1) for the numerical QBI. Assuming that the input diffusion

¹In Tuch (17), k is set to 48° and σ to 5° although it is argued that k and σ can be selected for a desired level of numerical precision.

MRI volume is of size $X \times Y \times Z \times N$, we have a diffusion signal vector of $N \times 1$ at each voxel, where N is the number of gradient directions taken. We let k be the number of points on each equator over which the numerical Funk–Radon integral is computed. In our technique, we have defined $R = (1/2)(\ell + 1)(\ell + 2)$ to be the number of elements in the SH basis. It is straightforward to see that the analytical ODF reconstruction is $O(XYZNR)$ because of the $O(NR)$ matrix multiplication at every voxel, while the original ODF reconstruction is $O(XYZNk)$ because of the integration of k equator points for each sampling direction N at every voxel. Therefore, the difference in computational complexity between the two methods can be explained by the difference between R and k , where R is generally smaller than k . For example, at orders $\ell = 4, 6, 8$, for our method $R = 15, 28, 45$ respectively; in Tuch (17) k is set to be 48.

Robustness to Noise

To evaluate the ODF reconstruction methods, we generate HARDI signal profiles using the multi-tensor model described earlier for a single voxel. We perform two simulations to evaluate the ODF estimation. In the first simulation, SNR was fixed to 35 while varying the b -value between 500 and 12,000 s/mm² for estimation order $\ell = 4$ and $\ell = 8$. In the second simulation, we fix the b -factor to $b = 3,000$ s/mm² and estimation order to $\ell = 8$ while varying the SNR between 5 and 50. In both tests, we randomly choose the number of fibers n per voxel between 1, 2, and 3. The optimal regularization λ parameter can be obtained from the L -curve numerical method (34). To avoid having to compute the optimal λ for each HARDI profile at every iteration, we set $\lambda = 0.006$ for the rest of the article, a value shown to provide good separation of 1-fiber from 2-fiber distributions over a large range of SNR and b -values (13). We apply the transformation given in Eq. [17] to obtain the estimated SH coefficients of the ODF and use Eq. [9] to obtain the corresponding discrete function on the sphere for the N sampling directions.

ODF Shape Comparison

Letting f represent the exact ODF and f' the estimated ODF, we compute the average Euclidean squared error between f and f' over all tests and N samplings of the sphere, i.e.

$$[f, f'] = \frac{1}{N} \sum_{i=1}^N (f'_i - f_i)^2. \quad [19]$$

This method is used to compare ODF shapes throughout the Results section, including experiments using synthetic data, data from the biological phantom and human brain data. However, other appropriate distances such as Kullback–Leibler divergence could also be considered.

Fiber Detection and Angular Resolution

It is generally assumed that the fiber directions are simply given by the local maxima of the normalized $[0, 1]$ ODF,

where the function surpasses a certain threshold (here, we use 0.5). This relation can be used to extract local fiber orientation estimates for comparison with an estimated ground truth as done in (31) and is also useful for visualization purposes, overlaying the maxima over the ODF, e.g. (18,25,30).

To evaluate fiber detection differences between Laplace–Beltrami and Tikhonov regularization, we test on synthetic HARDI data generated using a 3rd order tessellation of the icosahedron that gives 81 samplings on the hemisphere, a SNR of 10 and 2 orthogonal fibers. We vary estimation order ℓ and use two b -factors of 3,000 and 1,000 s/mm². (1) We generate 1,000 such HARDI data separately, (2) we estimate ODFs with/without Laplace–Beltrami and Tikhonov regularization, (3) we count the number of times we correctly detect 2 ODF maxima, and (4) we report the percentage and average angular error \pm standard deviation in degrees over all 1,000 trials.

Finally, we perform a numerical experiment to evaluate angular resolution limitations of the ODF reconstruction with/without Laplace–Beltrami and Tikhonov regularization. We generate noise-free synthetic HARDI profiles for 2 fibers, for b -factors of 3,000 and 1,000 s/mm² and for spherical sampling densities $N = 81$ and $N = 321$. Then, we vary the crossing angle between fibers to determine the critical angle at which only a single maxima starts to be detected instead of two. We report this critical angle as the angular resolution of the estimation.

RESULTS

We show four contributions of our regularized analytical QBI method; (1) it is up to 15 times faster than Tuch’s numerical method, (2) it is robust to noise, (3) it improves accuracy in ODF maxima detection at the cost of slightly reducing angular resolution, and (4) it recovers fiber crossings from synthetic data, from a biological phantom and from real human brain data.

Running Time Comparison

Table 2 shows that analytical QBI is up to 15 times faster than Tuch’s numerical QBI in practice. Computation is performed on a Dell single processor, 3.4 GHz, 2 GB RAM machine. Given N samples on the sphere and a SH basis of order ℓ , the theoretical speed-up factor for order $\ell = 4, 6, 8$ is $\sim 3, 2, 1$ respectively. However, in practice the running time of the analytical QBI is nearly 15 times faster. The factor of 5 gap between the theoretical and experimental speed-up is because of the constant time operation hidden and not accounted for in the “big O” analysis of the running time in theory. This is mainly due to the interpolation kernel width of Tuch’s approach which is not considered in the complexity analysis and which adds a constant number of operations (4–7 in practice) at every equator point. The reading/writing of 4D volumes with 4th dimension N instead of R is also slower. Thus, as mentioned in (28), there are interesting potential data compression applications offered by the SH basis, since only a few harmonic coefficients need to be stored per voxel.

Table 2
Analytical QBI is up to 15 Times Faster than Tuch's Numerical QBI in Practice

Dataset	Size $X \times Y \times Z \times N$	Analytical QBI			Tuch QBI	
		$\ell = 4, R = 15$ Time (min:sec)	$\ell = 6, R = 28$ Time (min:sec)	$\ell = 8, R = 45$ Time (min:sec)	k	Time (min:sec)
Rat	$35 \times 17 \times 35 \times 91$	0:01.27	0:01.61	0:02.09	48	0:13.59
Brain#1	$128 \times 128 \times 63 \times 100$	0:40.42	0:53.06	1:49.37	50	13:27.12
Brain#2	$93 \times 116 \times 93 \times 61$	0:27.07	0:35.76	0:58.38	45	5:43.86

Tuch's numerical QBI is $O(XYZNk)$ and our analytical QBI is $O(XYZNR)$. Computation experiments are performed on a Dell single processor, 3.4 GHz, 2 GB RAM machine.

ODF Shape Comparison and Robustness to Noise

Figure 2 shows that the Laplace–Beltrami regularization reduces ODF estimation errors.

First, Fig. 2a shows that the estimation is precise. As expected, we observe that for optimal b -value, the error is less than 1%. Note also that the best results are not for the highest b -values because in this case, the diffusion weighted signal decreases sharply in SNR. This is because the signal intensity of the Gaussian-like profile of the signal varies more rapidly with ADC. At the extremity of very high b -values, most of the signal is lost. One would need to choose a much higher λ to prevent the regularized curve from approaching the unregularized curve. It is also expected that for low b -values, the accuracy of the estimation is reduced. This is mostly due to the Bessel function averaging effect (17). The signal is too smooth and there is very small contrast. Regularization is unnecessary in these cases. It is also interesting to note that a lower estimation order ($\ell = 4$) can outperform the regularized and unregularized estimation of order $\ell = 8$ for low b -values ($b < 1,500$ s/mm²). the higher angular resolution afforded by higher estimation order is not possible. Finally, lowest error is observed in a plateau of relatively high b -values between 2,000 and 6,000 s/mm², which agrees with reported results in the literature (6,18). How to choose the optimal b -value for a particular HARDI acquisition is still an open question but in our particular synthetic experiment, we find the smallest error occurring for b -factors between $b = 4,000$ and 5,000 s/mm².

Last, Fig. 2b shows that the error decreases when the noise level decreases, from more than 12% for a noisy signal (small SNR) to less than 1% for high quality data (large SNR). It is also important to compare the green curve (numerical QBI) with the analytical QBI blue/red curves with/without regularization that we have added in this test to report robustness results of the different approaches. We note that for high quality data, numerical QBI and analytical QBI with/without regularization are almost identical whereas for noisy data, the analytical QBI with regularization $\lambda = 0.006$ performs best, while numerical QBI is better than analytical QBI without regularization ($\lambda = 0$). Finally, as in Fig. 2a, we have analyzed the behavior of the unregularized estimation of lower order $\ell = 4$. The unregularized estimation errors of order 4 overlap the regularized estimation errors of order 8 for low quality data with SNR less than 20 and for high quality data with SNR larger than 20, the unregularized estimation errors of order 4 overlap the unregularized estimation errors of order 8. Because of

these overlaps and for clarity of Fig. 2b, we do not plot the curve.

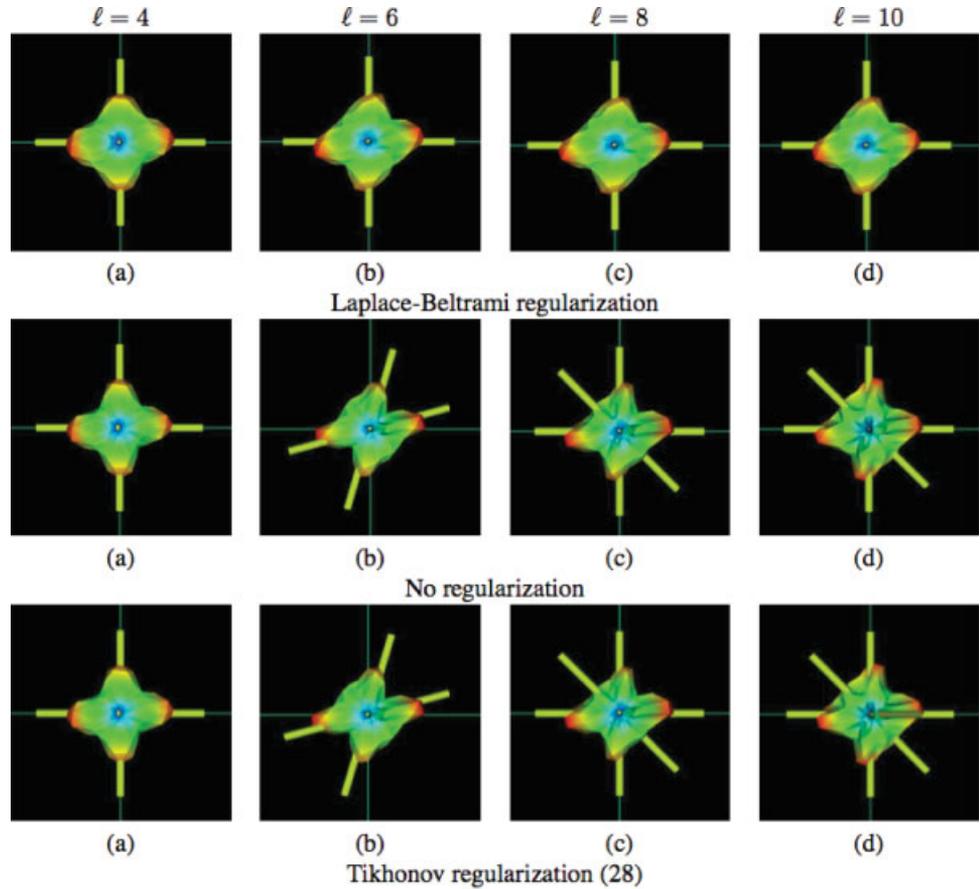
Fiber Detection and Angular Resolution

In this section, we show four results: (1) Laplace–Beltrami regularized ODFs decrease small perturbations due to noise that can create false maxima, (2) Laplace–Beltrami regularized ODFs improve the detection of crossing fibers while reducing angular error as calculated from the maxima of the ODF, (3) ODF maxima agree with the known underlying fiber configurations under different signal parameters, and (4) Angular resolution of the QBI method depends on regularization technique, estimation order ℓ , acquisition b -factor and spherical sampling density N .

First, Fig. 3 shows that Laplace–Beltrami regularized ODFs remove small perturbations due to noise that can create false maxima in ODFs estimated without regularization and with Tikhonov regularization (last two rows of Fig. 3). In particular, in this test, there are 3 maxima detected for Tikhonov regularization at $\ell = 8$ and for no regularization at $\ell = 8, 10$ and up to 4 maxima (one is less obvious and comes out of the page) for Tikhonov regularization at order $\ell = 10$. Hence, not only there is a danger of over-modeling the data when using a high order ($\ell > 6$) with/without Tikhonov regularization but also, there is an angular error made on the detected maxima at $\ell = 6$. One can potentially tweak the threshold in the fiber detection method to remove some of the spurious maxima but cannot correct for the angular error made. The Tikhonov regularization results shown in Fig. 3 were obtained with parameter $\lambda = 2$.

Second, Table 3 confirms the observation made in Fig. 3. Table 3 shows that Laplace–Beltrami regularized ODFs improve the detection of crossing fibers while reducing angular error as calculated from the maxima of the ODF, as seen in the example of $\ell = 6$ of Fig. 3. For these tests, we recorded the percentage of correct 2-fiber ODF maxima detected and noted the average angular error \pm standard deviation made in degrees. In the case there were more than 2 maxima detected, the error was estimated on the two closest ODF maxima to ground truth. With Laplace–Beltrami regularization the detection is nearly perfect at $b = 3,000$ s/mm² and above 88% at $b = 1,000$ s/mm² for all orders whereas the detection dramatically decreases for high order estimations $\ell = 6, 8$, and 10 for Tikhonov regularization and without regularization. Table 3 shows also that Laplace–Beltrami smoothing reduces the average

FIG. 3. ODF regularization decreases small perturbations due to noise that can create false maxima. Thin lines are the true underlying fiber directions and thicker tubes are the detected ODF maxima. Simulations are done with harmonic order between $\ell = 4, 6, 8, 10$ without regularization ($\lambda = 0$), with Tikhonov regularization ($\lambda = 2$) and with our Laplace–Beltrami regularization ($\lambda = 0.006$). ODFs are generated with 81 gradient directions, SNR = 10, $b = 3,000$ s/mm² and orthogonal fibers crossing.



angular error as calculated from the ODF maxima detected and their ground truth. Overall, orthogonal fibers are detected accurately by all methods for $\ell = 4$ even at the lower b -value. The approximation is smooth enough that effects due to noise are reduced. However, Laplace–Beltrami regularization is necessary to obtain good results for higher order ($\ell > 4$).

The behavior of ODFs in Fig. 3 and better performance in Table 3 because of the Laplace–Beltrami regularization in the ODF estimation are expected. High-order modeling error due to noise is avoided while minimally altering the lower order coefficients involved in the description of the

ODF. The Tikhonov regularization used in Hess et al. (28) is an approach mainly used to improve the numerical conditioning of the matrices. However, it is not designed to smooth the spherical functions as it perturbs the diagonal elements uniformly which has the effect of adding λ to every eigenvalue. This does not change the overall shape of the ODF and does not eliminate spurious peaks. Moreover, in our experiments, we found that both Hess et al. SH basis and our basis are well conditioned. That is, the ratio of the largest over the smallest eigenvalue of the $\mathbf{B}^T \mathbf{B}$ matrix involved in the least-square expression of Eq. [8] remains small, even when varying density

Table 3

Laplace–Beltrami Regularization Improves the Percentage of Detecting Crossing Fibers While Reducing Angular Error as Calculated from the Maxima of the ODF

b -factor	Order ℓ	Laplace–Beltrami ($\lambda = 0.006$)	No regularization ($\lambda = 0$)	Tikhonov (28)
3,000 s/mm ²	4	99.9%, 2.1° ± 5.4°	99.6%, 1.6° ± 4.7°	99.6%, 1.6° ± 4.7°
	6	99.6%, 2.8° ± 6.1°	95.8%, 4.4° ± 7.2°	95.9%, 4.2° ± 7.0°
	8	99.4%, 2.5° ± 5.8°	62.9%, 4.6° ± 7.4°	63.1%, 4.5° ± 7.2°
	10	99.6%, 2.6° ± 5.8°	31.5%, 6.0° ± 7.5°	31.1%, 6.6° ± 7.7°
1,000 s/mm ²	4	96.2%, 8.6° ± 10.6°	96.1%, 7.1° ± 8.9°	96.2%, 7.0° ± 8.8°
	6	90.3%, 10.4° ± 10.8°	69.4%, 11.9° ± 10.1°	71.0%, 11.1° ± 9.7°
	8	88.5%, 10.8° ± 11.4°	23.4%, 11.3° ± 10.9°	24.3%, 11.1° ± 10.5°
	10	88.0%, 10.8° ± 11.3°	4.5%, 12.6° ± 10.7°	4.4%, 12.8° ± 10.2°

Percentage reflects if ODF were correctly detected with 2 maxima. We report success %, average angular error in degrees ± standard deviation in degrees for the ODF estimated with Laplace–Beltrami regularization, with Tikhonov regularization and without regularization. The simulations are HARDI signal of 1,000 separate pairs of orthogonal fibers, with a SNR of 10 and spherical sampling density of $N = 81$, while varying estimation order ℓ and the b -factor. For $\ell > 4$, Laplace–Beltrami regularization is necessary to obtain good fiber detection.

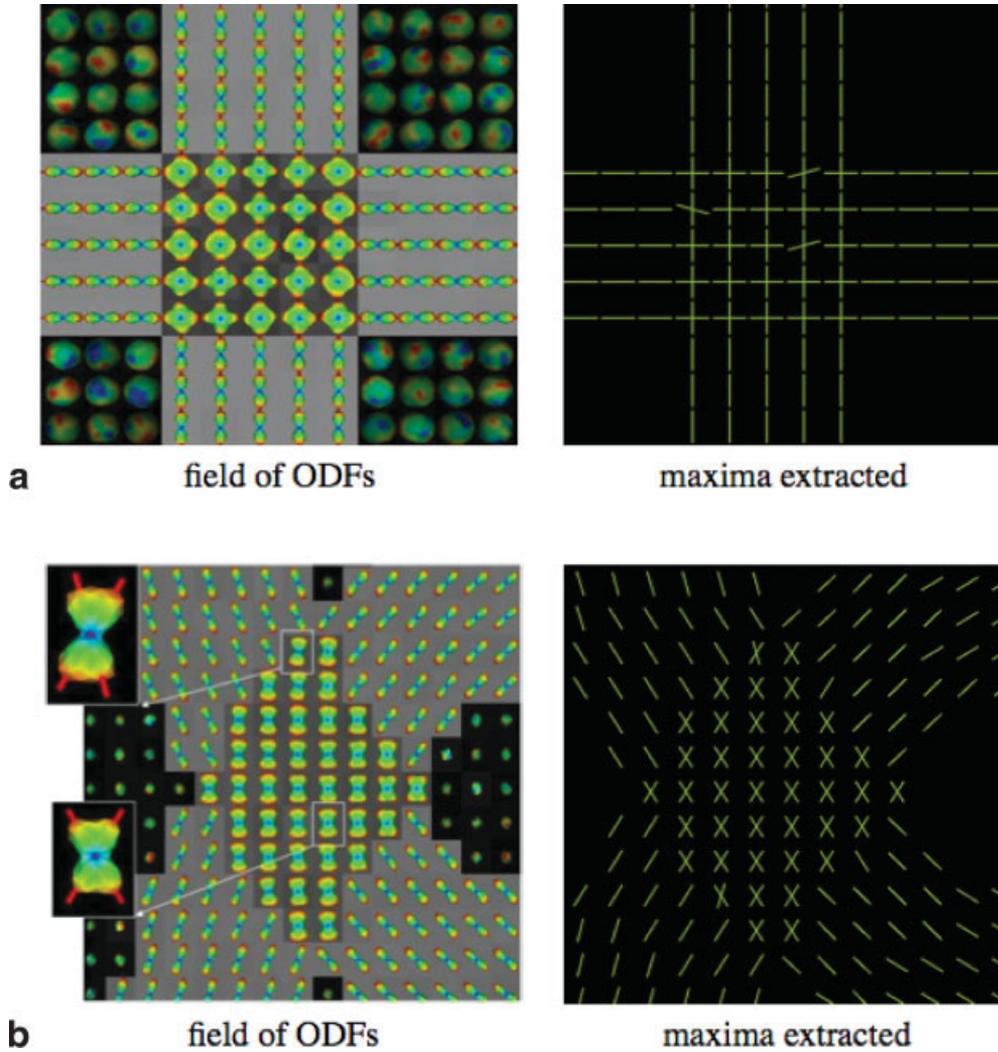


FIG. 4. ODFs recover multiple fiber crossing in synthetic data generated with $b = 1,500$ s/mm², SNR = 15. (a) 90° crossing and (b) 60° crossing. An order 8 estimation with $\lambda = 0.006$ was used.

N for $N = 30, 66, 81, 99, 130$, and 321 samplings on the hemisphere.

Next, Fig. 4 qualitatively shows that we are able to reliably recover the underlying fiber population for imaging parameters similar to clinical applications. The example is generated with b -factor of 1,500 s/mm², with sampling density $N = 81$, with noise level of SNR = 15 and with a 2-fiber population crossing at 90° and 60° respectively. Other examples for b -factors (500, 1,500, 3,000 s/mm²) and other SNR (5, 15, 34) can be found (25). Note that the detected ODF maxima qualitatively agree with the underlying fibers with some small angular error due to noise level, even though a relatively low b -value is used. The average angular error between ground truth directions and detected maxima was calculated to be $\sim 10^\circ$ for both datasets. Simulation with a 3rd order tessellation gives a 16° difference between each ODF reconstruction point. The error is thus less than an angular sampling unit and separating both fiber configurations is done with the same quantitative precision. As one would expect, the smaller the angle between fibers, the harder it is to distinguish them.

Last, Table 4 illustrates the angular resolution of the analytical QBI technique with respect to the SH order

ℓ and the regularization method used in the ODF estimation as well as the b -factor and spherical sampling density used to generate the synthetic signal. The critical separation angle is reported, that is the angle between 2 fibers under which only a single ODF maximum starts to be detected. Five expected observations can be made from Table 4. (1) It is harder to distinguish crossing fibers for lower b -values. In fact, there is approximately a 15° gain in angular resolution when going from b -value 1,000 s/mm² to b -value 5,000 s/mm² and higher. (2) For $\ell > 4$ there is an improvement in angular resolution of roughly 3° to 10° for all methods because higher order estimation include higher order frequencies in the approximation of the signal. The largest angular resolution improvement occurs for high b -value and high spherical sampling density ($N = 321$). (3) Tikhonov and no regularization have a better angular resolution than using Laplace–Beltrami regularization. In the Laplace–Beltrami case, the critical angle remains reasonable and about 0°–5° higher. This decrease in angular resolution is due to smoothing of the higher frequency information in Laplace–Beltrami regularization. Hence, there is a trade-off between accuracy of the fiber detection and angular resolution. (4) There is an improvement of a few degrees in angular

Table 4

Angular Resolution Limitations of the ODF Reconstruction With/Without Laplace–Beltrami (LB) ($\lambda = 0.006$) and Tikhonov (TK) (28) Regularization

		$b = 10,000 \text{ s/mm}^2$			$b = 5,000 \text{ s/mm}^2$			$b = 3,000 \text{ s/mm}^2$			$b = 1,000 \text{ s/mm}^2$		
Order ℓ		LB	TK	$\lambda = 0$	LB	TK	$\lambda = 0$	LB	TK	$\lambda = 0$	LB	TK	$\lambda = 0$
$N = 81$	4	57°	55°	55°	59°	56°	56°	63°	60°	60°	75°	71°	71°
	6	53°	49°	49°	55°	52°	53°	59°	53°	54°	74°	68°	68°
	8	52°	46°	48°	55°	50°	50°	58°	53°	53°	74°	68°	68°
	10	52°	46°	47°	55°	49°	49°	58°	52°	53°	74°	68°	68°
$N = 321$	4	54°	54°	54°	56°	55°	55°	60°	59°	59°	72°	71°	70°
	6	44°	42°	43°	47°	45°	46°	52°	50°	50°	69°	67°	67°
	8	39°	35°	36°	44°	40°	41°	50°	46°	47°	69°	66°	67°
	10	36°	31°	31°	43°	38°	39°	50°	46°	46°	69°	66°	67°

Noise-free synthetic HARDI profiles are generated for 2 fibers, for b -factors of 10,000, 5,000, 3,000 and 1,000 s/mm^2 and for spherical sampling density using $N = 81$ and $N = 321$. The crossing angle between fibers is varied to report the critical angle under which only a single maxima starts to be detected instead of two.

resolution when increasing spherical sampling density used to generate the synthetic data. (5) For higher b -values and higher sampling density, there is an important angular resolution improvement of about 10° when comparing the Laplace–Beltrami regularized estimation of order $\ell > 4$ with the Tikhonov and $\lambda = 0$ estimations at order $\ell = 4$. This increase in angular resolution is even more apparent at higher b -values.

In a different angular resolution study, Hess et al. (28) have studied the theoretical relationship between the approximation order ℓ and angular resolution, as calculated from the full width half max of the main lobe of the spherical point spread function. In particular, for $\ell = 4, 6, 8, 10$, the angular resolution is approximately $65^\circ, 45^\circ, 35^\circ, 30^\circ$ respectively (described in Ref. 28, Fig. 1). This study mainly looks at the impact of SH estimation order without considering the HARDI signal parameters or the regularization technique used in the estimation. In practice, we observe critical angles that are higher because there are many other parameters than order ℓ that may influence angular resolution of the solution. Not only have we studied the SH order ℓ but we have also studied

the regularization method and λ parameter as well as the b -factor and the spherical sampling density N .

This section showed that there is a trade-off between angular resolution and accuracy of the fiber detection. Overall, Laplace–Beltrami regularization reduces errors in fiber detection while keeping a reasonable angular resolution, and in the case of high b -value acquisitions with high number of sampling directions on the sphere, the Laplace–Beltrami regularized solution for high orders ($\ell > 4$) has better angular resolution than the unregularized solution at order $\ell = 4$.

Biological Phantom

Figure 5 shows that ODFs recover multiple fiber crossing in the rat biological phantom. Note that the ODFs have multiple peaks that agree with the known underlying fiber population, which we have emphasized with the subfigure showing only ODF maxima. This is not the case when looking at the DTI ellipsoids which are flat and sphere-like in the crossing. Moreover, ODF shapes shown in the second row of Fig. 5 are qualitatively nearly identical. In fact, when

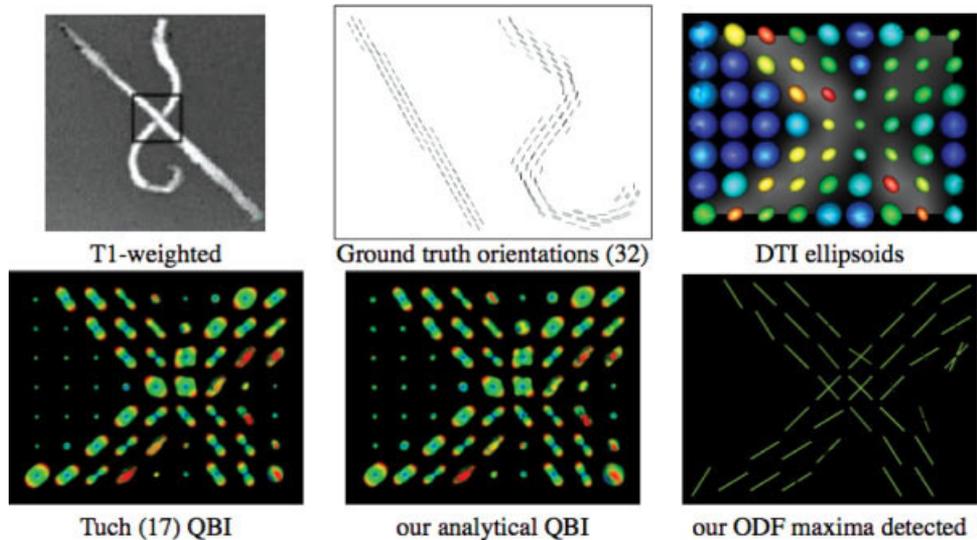


FIG. 5. ODFs recover multiple fiber crossing in the rat biological phantom. Ground truth orientations were obtained (31). An order 8 estimation with $\lambda = 0.006$ was used.

Table 5
Analytical and Numerical QBI Methods Yield Essentially the Same Results on the Rat Phantom While Reducing Errors Obtained from DTI

	Comparison with ground truth directions	
	Median	Mean \pm std
Analytical QBI	12.20°	15.94° \pm 15.32°
Numerical QBI	12.19°	15.94° \pm 15.40°
DTI	15.2°	19.4° \pm 16.2°

Ground truth orientations and comparison are done as in Ref. 31.

computing the mean and standard deviation of the average Euclidean squared difference (Eq. [19]) between the ODFs at every voxel of the volume, we obtain a $0.55\% \pm 0.17\%$ difference, demonstrating the strong agreement between the methods. That is, ODFs computed from the analytical and numerical QBI (17) are more than 99% in agreement.

Table 5 shows that analytical and numerical QBI (17) methods yield essentially the same results on the rat phantom while reducing errors obtained from DTI. To perform a quantitative evaluation of the ODF maxima, we used the “ground truth” orientations from Ref. 31 (illustrated in Fig. 5). The orientations are in fact more like a silver standard (a gold standard does not exist in a biological phantom) as they were determined by extracting the centerlines of each super-sampled rat cord and then smoothly extended to the center of the boundaries. For each ODF dataset and for the DTI ellipsoids, the maxima are extracted. Then, at each voxel, the smallest angular difference between the available maximum(a) and ground truth

orientation(s) is recorded. The median and mean \pm standard deviation (std) orientation errors in degrees are shown in Table 5. The analytical and numerical QBI methods yield essentially the same results while significantly reducing the errors obtained from the DTI profiles. It is important to note that because of the discrete sampling of the sphere, even perfect ODF data will be expected to have some error with respect to the ground truth orientations. This minimal expected error is related to the solid angle subtended by one facet of the sphere tessellation induced by the sampling. For example, it can be found to equal 7.2° for a uniform sampling of the hemisphere by 100 directions (31).

Human Brain Data

As for the rat biological phantom, if we compare ODFs on the brain#1, the overall shapes of the ODFs are nearly the same for the analytical and numerical QBI (17) methods. We record a small mean and standard deviation of the average Euclidean squared difference (Eq. [19]) between ODFs from the two methods of $0.68\% \pm 0.23\%$.

Qualitatively, Fig. 6 shows that ODFs recover multiple fiber crossing in brain#2 where DTI profiles are limited. Diffusion tensors and ODFs are overlaid on the classical FA anisotropy measure and its high order generalization GFA (17) respectively. We zoom on the ODFs in two ROI. Again, ODF maxima agree with our knowledge of the crossings between the cortical spinal tract (cst) and superior longitudinal fibers (coming out of the plane) in ROI (a) and crossings between the cst and corpus callosum (in the plane) in ROI (b). Figure 6 also emphasizes the limitations of DTI and the ability of the ODF to recover multiple fiber

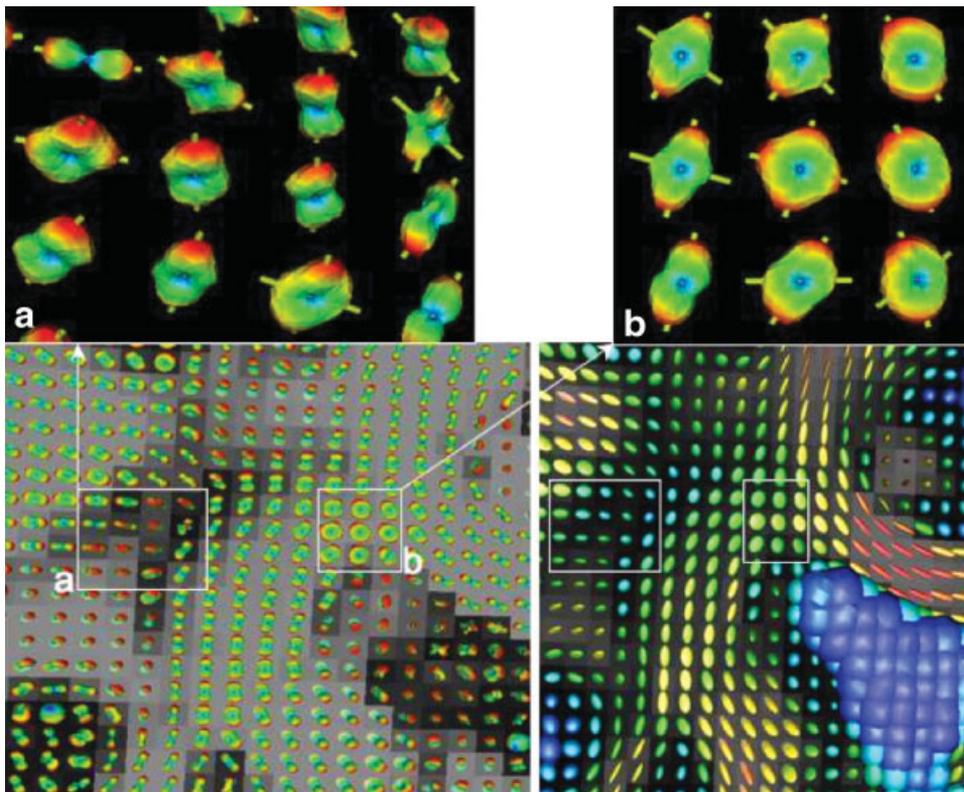


FIG. 6. ODFs recovering multiple fiber crossing in a ROI of brain#2 where DTI profiles are limited. The ODFs and diffusion tensors are overlaid on the GFA and FA measure respectively and we zoom on the ODFs of two smaller ROI. The ROI(a) shows crossing fibers between the cortical spinal tract (cst) and superior longitudinal fibers (coming out of the plane) and the ROI(b) shows crossing between the corpus callosum (in the plane) and the cst. ROI(a) is tilted to see the fiber crossing better. An order 8 estimation with $\lambda = 0.006$ was used.

orientations that are nearly orthogonal even from a dataset with relatively low b -value of $1,000 \text{ s/mm}^2$.

DISCUSSION

We have proposed a regularized, fast and robust analytical solution for the ODF estimation problem in QBI.

The main focus of the paper was on the ODF reconstruction. We have derived an analytical solution for QBI that agrees with previous solutions (23,28) recently developed in parallel in the literature. As already pointed out in the article, the Hess et al. (28) paper is very close in spirit to our article but both papers complement each other. Each paper makes separate contributions that support the strength of the analytical solution using spherical harmonics. Hess et al. have shown interesting properties of the analytical solution that we have not focused on, such as, the theoretical angular resolution limitation with respect to the harmonic order ℓ of the basis, the better performance of the analytical solution for low sampling density of q-ball ODF reconstruction compared to Tuch's numerical technique and the data compression potential of the approach.

Further, we have made additional contributions: (1) we have introduced a new closed-form Laplace–Beltrami regularization in the signal-fitting step, (2) we have proved a novel corollary to the Funk–Hecke theorem to derive the analytical solution for ODF reconstruction, (3) we have done a complexity analysis with a practical running time experiment that shows an increase in speed by a factor of 15 over numerical QBI (17), (4) we have shown the advantages of Laplace–Beltrami regularization theoretically and experimentally at the cost of slightly reducing angular resolution, and (5) we have performed a careful validation of both analytical and numerical QBI (17) techniques on noisy synthetic data and real data where ground truth is known.

The new analytical solution comes from the fitting of the signal with SH, which allowed us to impose the Laplace–Beltrami regularization criterion. The Funk–Hecke corollary was proved using a delta sequence so that the Funk–Hecke formula could be used to solve the Funk–Radon integral and obtain a simple regularized expression for the ODF reconstruction. Without this derivation, the FRT can only be computed with a more complicated numerical scheme. This solution eliminates the discrete numerical integration step over each equator needed in Tuch's numerical QBI implementation which speeds up computation by a factor up to 15, while solutions stay in close agreement. As imaging techniques are improved and the number of gradient directions are increased, this can potentially be an important speed-up factor even with the optimal numerical QBI implementation. It can also have potential data compression applications (28).

We have also shown the better performance and robustness of the ODF reconstruction in the presence of noisy synthetic data. Theoretically, we have argued that the Laplace–Beltrami is the right regularization criterion to use. Experimentally, at the cost of slightly reducing angular resolution, we have shown that it incorporates less noise in the high order SH coefficients describing the ODF and thus better describes voxels with multiple fibers, especially in the presence of high noise level. Therefore, it is possible to use a high order approximation while limiting the

over-modeling of perturbations due to noise. This is not the case without regularization or when using Tikhonov regularization.

Determining the angular resolution of the analytical QBI method is a difficult problem that depends on the SNR of the signal, b -value and spherical sampling density N used in the QBI acquisition, and regularization parameter λ and estimation order ℓ used in the ODF reconstruction. Some answers to this angular resolution problem have been proposed in the literature. Tuch (17) reported the angular limitation of the QBI protocol due to the intrinsic beam smoothing of the FRT whereas Hess et al. (28) studied the theoretical angular limitation arising from the choice of SH order ℓ . In general, it would be of great interest to know the achievable angular resolution given the specific sampling N , b -value and SNR in the diffusion-weighted data and given the order ℓ and regularization λ used in the ODF reconstruction.

Overall, the article showed some of the powerful properties of the SH representation, which have interesting postprocessing applications. First, Gaussian/Laplacian smoothing on the sphere by the Laplace–Beltrami operator extension is trivial to compute and allows one to impose a regularization criterion on the solution. Second, derivatives and integrals on the sphere have analytical expressions. This could allow for an automatic maxima extraction on the ODFs. ODF sharpening is also an application currently under investigation and was introduced in Ref. 25. Using both the Laplace–Beltrami and the Funk–Hecke theorem, we can define sharpening operations that enhance maxima of the underlying fiber distribution. This shows promising results especially when the dataset is acquired with low b -values (25). The underlying fibers are more easily detected from the sharpened diffusion ODF. Finally, it is now important to integrate this analytical ODF reconstruction in a high order tracking and segmentation algorithm.

ACKNOWLEDGMENTS

The authors thank the McConnell Brain Imaging Center and Center for Intelligent Machines of McGill University and especially, J.S.W. Campbell, K. Siddiqi, V. V. Ryman, and B.G. Pike for the rat cord dataset and the human brain dataset brain#1. We are also grateful to J.S.W. Campbell and P. Savadjiev for the numerical QBI implementation and rat spinal ground truth computation and comparison. Moreover, thanks to the Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany and especially, A. Anwender and T. Knoesche for the brain#2 dataset. This part was supported by PAI Procope. Thanks also to M. de La Gorce and C. Lenglet for their work. Finally, thanks to reviewers for comments and suggestions that have considerably improved the manuscript.

REFERENCES

1. LeBihan D, Breton E, Lallemand D, Grenier P, Cabanis E, Laval-Jeantet M. MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. *Radiology* 1986;161:401–407.
2. Beaulieu C. The basis of anisotropic water diffusion in the nervous system—A technical review. *NMR Biomed* 2002;15:435–455.

3. Basser PJ, Mattiello J, LeBihan D. Estimation of the effective self-diffusion tensor from NMR spin echo. *J Magn Reson B* 1994;123:247–254.
4. Callaghan PT. Principles of nuclear magnetic resonance microscopy. Oxford: Oxford University Press; 1993.
5. Wedeen VJ, Reese TG, Tuch DS, Weigel MR, Dou J-G, Weiskoff RM, Chessler D. Mapping fiber orientation spectra in cerebral white matter with fourier-transform diffusion MRI. In: Proceedings of the 8th Annual Meeting of ISMRM. Philadelphia, 2000. p 82.
6. Tuch DS. Diffusion MRI of Complex Tissue Structure. Ph.D. thesis, Harvard University and Massachusetts Institute of Technology, 2002.
7. Alexander DC, Barker GJ, Arridge SR. Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magn Reson Med* 2002;48:331–340.
8. Frank LR. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magn Reson Med* 2002;47:1083–1099.
9. Ozarslan E, Mareci TH. Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution imaging. *Magn Reson Med* 2003;50:955–965.
10. Zhan W, Gu H, Xu S, Silbersweig DA, Stern E, Yang Y. Circular spectrum mapping for intravoxel fiber structures based on high angular resolution apparent diffusion coefficients. *Magn Reson Med* 2003;49:1077–1088.
11. Zhan W, Stein EA, Yang Y. Mapping the orientation of intravoxel crossing fibers based on the phase information of diffusion circular spectrum. *NeuroImage* 2004;23:1358–1369.
12. Chen Y, Guo W, Zeng Q, Yan X, Huang F, Zhang H, He G, Vemuri BC, Liu Y. Estimation, smoothing and characterization of apparent diffusion coefficient profiles from high angular resolution DWI. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, volume 1, Washington, DC, USA; 2004. pp. 588–593.
13. Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Apparent diffusion coefficients from high angular resolution diffusion imaging: Estimation and applications. *Magn Reson Med* 2006;56:395–410.
14. Tuch DS, Reese TG, Wiegell MR, Makris NG, Belliveau JW, Wedeen VJ. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn Reson Med* 2002;48:577–582.
15. Behrens TEJ, Woolrich MW, Jenkinson M, Johansen-Berg H, Nunes RG, Clare S, Matthews PM, Brady JM, Smith SM. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med* 2003;50:1077–1088.
16. Jansons KM, Alexander DC. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Probl* 2003;19:1031–1046.
17. Tuch DS. Q-ball imaging. *Magn Reson Med* 2004;52:1358–1372.
18. Tournier J-D, Calamante F, Gadian DG, Connelly A. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution *NeuroImage* 2004;23:1176–1185.
19. Alexander DC. Maximum entropy spherical deconvolution for diffusion MRI. In: Proceedings of the 19th International Conference on Information Processing in Medical Imaging (IPMI), Glenwood Springs, CO, USA; 2005. pp. 76–87.
20. Ozarslan E, Vemuri BC, Mareci TH. Fiber orientation mapping using generalized diffusion tensor imaging. In: IEEE International Symposium on Biomedical Imaging (ISBI), Arlington, Virginia, USA; April 2004. pp. 1036–1038.
21. Chen Y, Guo W, Zeng Q, He G, Vemuri BC, Liu Y. Recovery of intravoxel structure from HARD DWI. In: IEEE International Symposium on Biomedical Imaging (ISBI), Arlington, Virginia, USA; April 2004. pp. 1028–1031.
22. Campbell JSW, Siddiqi K, Rymar VV, Sadikot A, Pike BG. Flow-based fiber tracking with diffusion tensor q-ball data: Validation and comparison to principal diffusion direction techniques. *NeuroImage* 2005;27:725–736.
23. Anderson AW. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magn Reson Med* 2005; 54:1194–1206.
24. Perrin M, Poupon C, Cointepas Y, Rioul B, Golestani N, Pallier C, Riviere D, Constantinesco A, Bihan DL, Mangin JF. Fiber tracking in q-ball fields using regularized particle trajectories. In: Proceedings of the 19th International Conference on Information Processing in Medical Imaging (IPMI), Glenwood Springs, CO, USA; 2005. pp. 52–63.
25. Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. A linear and regularized ODF estimation algorithm to recover multiple fibers in q-ball imaging. Technical Report 5768, INRIA, November 2005.
26. Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. A fast and robust ODF estimation algorithm in q-ball imaging. In: IEEE International Symposium on Biomedical Imaging (ISBI), Arlington, Virginia, USA; April 2006. pp. 81–84.
27. Ozarslan E, Shepherd T, Vemuri BC, Blackband S, Mareci TH. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *NeuroImage* 2006;31:1086–1103.
28. Hess CP, Mukherjee P, Han ET, Xu D, Vigneron DB. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magn Reson Med* 2006;56:104–117.
29. Hess CP, Mukherjee P, Han ET, Xu D, Vigneron DB. A spherical harmonic approach to q-ball imaging. In Proceedings of the 13th Annual Meeting of ISMRM 2005, Miami, Florida, USA; p. 389.
30. Lin C, Wedeen V, Chen J, Yao C, Tseng WI. Validation of diffusion spectrum magnetic resonance imaging with manganese-enhanced rat optic tracts and ex vivo phantoms. *NeuroImage* 2003;19:482–495.
31. Savadjiev P, Campbell JSW, Pike BG, Siddiqi K. 3D curve inference for diffusion MRI regularization and fibre tractography. *Med Image Anal* 2006;10:799–813.
32. Tuch DS, Wisco JJ, Khachaturian MH, Ekstrom LB, Kotter R, Vanduffell W. Q-ball imaging of macaque white matter architecture. *Phil Trans Roy Soc Lond B* 2005;360:869–879.
33. Backus GE. Geophysical interpolation of measurements of average phase velocities of surface waves over great circular and semi-circular paths. *Bull Seismol Soc Am* 1964;54:571–610.
34. Hansen PC. The L-curve and its use in the numerical treatment of inverse problems. In: Johnston P, editor. Computational inverse problems in electrocardiology. Southampton: WIT Press; 2001. p 119–142.
35. Andrews G, Askey R, Roy R. Special functions. Cambridge: Cambridge University Press, 1999.
36. Alexander DC, Barker GJ. Optimal imaging parameters for fiber-orientation estimation in diffusion MRI. *NeuroImage* 2005;27:357–367.
37. Jones DK, Horsfield MA, Simmons A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn Reson Med* 1999;42:515–525.
38. Anwander A, Tittgemeyer M, von Cramon DY, Friederici AD, Knösche TR. Connectivity-based parcellation of broca's area. *Cerebral Cortex*; 2007;17:816–825. (available at <http://cercor.oxfordjournals.org/cgi/content/abstract/bhk034v1>)