

Estimating Orientation Distribution Functions with Probability Density Constraints and Spatial Regularity*

Alvina Goh¹, Christophe Lenglet², Paul M. Thompson³, and René Vidal¹

¹ CIS and Dept. of Biomedical Engineering, Johns Hopkins University

² CMRR and Dept. of Electrical and Computer Engineering, University of Minnesota

³ LONI and Dept. of Neurology, University of California at Los Angeles

Abstract. High angular resolution diffusion imaging (HARDI) has become an important magnetic resonance technique for in vivo imaging. Current techniques for estimating the diffusion orientation distribution function (ODF), i.e., the probability density function of water diffusion along any direction, do not enforce the estimated ODF to be nonnegative or to sum up to one. Very often this leads to an estimated ODF which is not a proper probability density function. In addition, current methods do not enforce any spatial regularity of the data. In this paper, we propose an estimation method that naturally constrains the estimated ODF to be a proper probability density function and regularizes this estimate using spatial information. By making use of the spherical harmonic representation, we pose the ODF estimation problem as a convex optimization problem and propose a coordinate descent method that converges to the minimizer of the proposed cost function. We illustrate our approach with experiments on synthetic and real data.

1 Introduction

Diffusion magnetic resonance imaging (MRI) is a technique that produces in vivo images of biological tissues by exploiting the constrained diffusion properties of water molecules. An important area of research in diffusion MRI is the development of methods for reconstructing the orientation distribution function (ODF) – a *probability density function* (pdf) that characterizes the distribution of water diffusion along different directions on the sphere. A very successful reconstruction technique is high angular resolution diffusion imaging (HARDI) [1], which measures water diffusion along N uniformly distributed directions on the sphere. Given these signals, several reconstruction techniques can be used to characterize diffusion. Higher-order tensors leverage the work done in diffusion tensor imaging (DTI) [2] by using higher-order polynomials to model diffusivity [3,4]. [5] fits the HARDI signals with a mixture of tensors model whose weights are specified by a probability function defined on the space of symmetric positive definite matrices. Another approach is to construct the ODF directly from HARDI signals. One of the earliest methods, known as Q-ball imaging (QBI), uses the Funk-Radon transform to estimate ODFs [6]. ODFs have also been approximated with different basis functions

* Work supported by startup funds from JHU, by grants NSF CAREER IIS-0447739, NIH R01 HD050735, NIH R01 EB007813, NIH P41 RR008079, NIH P30 NS057091, ONR N00014-05-10836 and ONR N00014-09-1-0084.

such as spherical harmonics [7,8,9,10,11] and the Poussin kernel [12]. Such methods are typically very fast because the ODF can still be computed analytically.

A first important limitation of existing QBI methods is that they can give large diffusion estimates outside the major fiber directions. [8] addresses this by assuming that there is a distribution of fiber orientations in each voxel and using a sharpening spherical deconvolution method to transform the diffusion ODF into a sharp fiber ODF (fODF). A second limitation of existing QBI methods is that they do not enforce the estimated ODF to be nonnegative. When the diffusion MR signal is corrupted by noise, this can cause the estimated ODF to have negative values, a situation that does not obey the underlying principle of diffusion. [13] attempts to alleviate this problem by using a constrained spherical deconvolution (CSD) method to estimate the fODF. Even though CSD reduces the occurrence of negative values, it does not completely eliminate them. A more recent method [14] eliminates the negative values by minimizing a nonnegative least-squares cost function. A third limitation of existing QBI methods is that the ODF at each voxel is estimated independently of the information provided in the spatial neighborhood. This results in noisy estimates of the ODF field. While regularization methods have been developed [15,16], we are not aware of any work addressing all three issues for HARDI.

We present an estimation method that gives sharp diffusion ODFs, constrains the estimated ODF to be a proper pdf, and incorporates spatial regularization. Our algorithm is based on the ODF reconstruction scheme in [11], which derives the ODF taking into account the solid angle consideration and is able to give naturally sharp ODFs. This is different from existing works [8,10], where the computed ODF is the linear projection of the actual diffusion probability and gives an artificial weight to points according to their distances from the origin. Our method represents the ODF as a linear combination of spherical harmonic (SH) functions, whose coefficients are found by minimizing an energy that incorporates a regularization term and nonnegativity constraints. This results in a convex optimization problem whose global minimizer can be found using coordinate descent. We illustrate our method with experiments on synthetic and real data.

2 Analytical Computation of ODFs with Spherical Harmonics

We first review the ODF reconstruction scheme in [11]. Let S_0 be the baseline signal and $S(\theta, \phi)$ be the HARDI signal acquired at the gradient direction (θ, ϕ) . The ODF is $\mathbf{p}(\theta, \phi) = \frac{1}{4\pi} + \frac{1}{16\pi^2} FRT\{\nabla_b^2 \ln(-\ln(\frac{S(\theta, \phi)}{S_0}))\}$, where FRT is the Funk-Radon transform and ∇_b^2 is the Laplace-Beltrami operator independent of the radial component. Notice that the first term integrates to 1 over the sphere, and the second term integrates to 0 [11]. The (modified) SH basis [8] of order l contains $R = \frac{(l+1)(l+2)}{2}$ terms defined for $j(k, m) = \frac{k^2+k+2}{2} + m, k = 0, 2, 4, \dots, l$ and $m = -k, \dots, 0, \dots, k$, as

$$Y_j = \sqrt{2} \operatorname{Re}(Y_k^{|m|}) \text{ if } -k \leq m < 0; \quad Y_k^0 \text{ if } m = 0; \quad \sqrt{2} \operatorname{Im}(Y_k^m) \text{ if } 0 < m \leq k;$$

where $Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\phi}, \theta \in [0, \pi], \phi \in [0, 2\pi], P_l^m$ is a Legendre polynomial, and $\operatorname{Re}(\cdot)$ and $\operatorname{Im}(\cdot)$ are the real and imaginary parts, respectively. Notice that $Y_1(\theta, \phi) = \frac{1}{2\sqrt{\pi}}$ is a constant function on the sphere that integrates to a constant, whereas the integral of $Y_j(\theta, \phi), j > 1$ over the sphere is always 0.

In order to estimate the ODF, let $S(\theta_i, \phi_i)$ be the HARDI signal acquired at each of the N gradient directions, $(\theta_i, \phi_i)_{i=1}^N$, and define the $N \times 1$ vector $\mathbf{s} = [\ln(-\ln(\frac{S(\theta_1, \phi_1)}{S_0})) \dots \ln(-\ln(\frac{S(\theta_N, \phi_N)}{S_0}))]^\top$. The signal \mathbf{s} is first approximated as $\mathbf{s} \approx \mathbf{B}\mathbf{c}$, where \mathbf{B} is the $N \times R$ SH basis matrix whose i -th row of \mathbf{B} is given as $\mathbf{B}_i = [Y_1(\theta_i, \phi_i) \dots Y_R(\theta_i, \phi_i)]$, and \mathbf{c} is the $R \times 1$ vector of SH coefficients that parametrize the signal \mathbf{s} . Given \mathbf{s} and \mathbf{B} , the unknown vector \mathbf{c} is found by solving the least-squares problem

$$\min_{\mathbf{c} \in \mathbb{R}^R} f(\mathbf{c}) = \frac{1}{2} \|\mathbf{B}\mathbf{c} - \mathbf{s}\|^2. \tag{1}$$

Assume now that the ODF is reconstructed using a tessellation scheme with M gradient directions, $(\theta_i^r, \phi_i^r)_{i=1}^M$. It is common for N , the number of gradient directions with which the HARDI signal is acquired, to be less than M . The reconstructed ODF is

$$\mathbf{p} = \mathbf{C}\mathbf{d}, \tag{2}$$

where \mathbf{C} is an $M \times R$ SH basis matrix whose i -th row is $\mathbf{C}_i = [Y_1(\theta_i^r, \phi_i^r) \dots Y_R(\theta_i^r, \phi_i^r)]$, and \mathbf{d} is the vector of SH coefficients of the ODF, which is given by [11]

$$\mathbf{d} = \left[\frac{1}{2\sqrt{\pi}} \quad \mathbf{0}_{1 \times (R-1)} \right]^\top + \frac{1}{16\pi^2} \mathbf{L}\mathbf{P}\mathbf{c}. \tag{3}$$

\mathbf{L} is the $R \times R$ diagonal Laplace-Beltrami eigenvalues matrix with $L_{jj} = -l_j(l_j + 1)$, where l_j is the order of the j -th term, and \mathbf{P} is the $R \times R$ diagonal Funk-Radon transform matrix, where $P_{jj} = 2\pi P_{l_j}(0)$ and $P_{l_j}(0)$ is the Legendre polynomial of degree l_j at 0.

3 Nonnegative and Spatially Regularized ODF Estimation

Notice that while the least-squares estimation method in §2 enforces the sum of \mathbf{p} to be one, it does not restrict \mathbf{p} to be nonnegative. In addition, the ODF reconstruction at a voxel is done independently of the information contained in the spatial neighborhood of that voxel. In this section, we present our estimation method that constrains the estimated ODF to be a proper probability density function and incorporates spatial regularity.

Let V denote the HARDI volume and $|V|$ the number of voxels in V . At each voxel $\mathbf{x}_i = (x_i, y_i, z_i)$, we have the base-line signal $S_{0,i}$ and the $N \times 1$ HARDI signal S_i . Thus, we can define the signal vector $\mathbf{s}_i = [\ln(-\ln(\frac{S_i(\theta_1, \phi_1)}{S_{0,i}})) \dots \ln(-\ln(\frac{S_i(\theta_N, \phi_N)}{S_{0,i}}))]^\top$ and its corresponding vector of SH coefficients \mathbf{c}_i . In order to enforce that the ODF \mathbf{p}_i at \mathbf{x}_i is nonnegative, we need to enforce the additional constraint $\mathbf{p}_i = \mathbf{C}\mathbf{d}_i \geq 0$. Making use of Eqns. (2) and (3), we rewrite the constraint as $-\mathbf{C}\mathbf{L}\mathbf{P}\mathbf{c}_i \leq 4\pi\mathbf{1}$.

To solve the ODF estimation problem in a way that accounts for the nonnegativity of \mathbf{p} and incorporates spatial regularization, we define the following optimization problem

$$\begin{aligned} \min_{\mathbf{c}_1, \dots, \mathbf{c}_{|V|}} g(\mathbf{c}_1, \dots, \mathbf{c}_{|V|}) &= \frac{1}{2} \sum_{i=1}^{|V|} \|\mathbf{B}\mathbf{c}_i - \mathbf{s}_i\|^2 + \lambda \sum_{\|\mathbf{x}_i - \mathbf{x}_j\| < r} w_{ij} \|\mathbf{c}_i - \mathbf{c}_j\|^2, \\ \text{subject to } -\mathbf{C}\mathbf{L}\mathbf{P}\mathbf{c}_i &\leq 4\pi\mathbf{1}, \quad i = 1, \dots, |V|, \end{aligned} \tag{4}$$

The first term corresponds to the data term and the second term corresponds to the regularization term. The following parameters need to be set: 1) λ is the nonnegative regularization factor and marks the tradeoff between the data term and the regularization term.

When $\lambda \rightarrow 0$, spatial regularity is ignored whereas if $\lambda \rightarrow \infty$, the data term is disregarded. 2) r is the spatial radius defining the neighborhood of each voxel. Examples of defining the voxel connectivity in the 3D volume include choosing the 6 nearest voxels or those that lie within a certain distance. 3) w_{ij} are nonnegative weights that measure the similarity of the data at \mathbf{x}_i and \mathbf{x}_j within the local neighborhood. A common way of defining these weights is to use the Gaussian kernel and define $w_{ij} = \exp(-\frac{\|\mathbf{s}_i - \mathbf{s}_j\|^2}{\sigma^2})$.

It is possible to restate the optimization problem in Eqn. (4) as a large quadratic optimization problem and one could theoretically attempt to solve for $[\mathbf{c}_1, \dots, \mathbf{c}_{|V|}]$ simultaneously. However, adopting such a strategy will require a tremendous amount of memory and intensive computational power as a typical HARDI volume contains $|V| \approx 10^6$ voxels, the signal at each voxel is acquired at $N \approx 100$ gradient directions, and the ODF reconstruction is done with a few hundred tessellation directions. Instead, we adopt an iterative algorithm, specifically the coordinate descent method, and we show that coordinate descent will converge to the minimizer of Eqn. (4).

Theorem 1. [17] Consider minimizing functions of the form

$$\phi(\beta_1, \dots, \beta_p) = \kappa(\beta_1, \dots, \beta_p) + \sum_{k=1}^p \chi_k(\beta_k), \tag{5}$$

where β_k is a vector, $\kappa(\cdot)$ is a differentiable and convex function, and $\chi_k(\cdot)$ are convex functions. When the different vectors β_k 's do not have overlapping entries and $\sum_{k=1}^p \chi_k(\beta_k)$ is separable, coordinate descent converges to the minimizer of $\phi(\cdot)$. The coordinate descent method is formally described as

1. Initialization: Set $t = 0$ and choose any $\beta^0 = (\beta_1^0, \dots, \beta_p^0) \in \text{domain}(\phi)$.
2. At each iteration $t + 1, t \geq 1$: Given $\beta^t = (\beta_1^t, \dots, \beta_p^t) \in \text{domain}(\phi)$, choose an index $s \in \{1, \dots, p\}$ and compute a new estimate $\beta^{t+1} = (\beta_1^{t+1}, \dots, \beta_p^{t+1}) \in \text{domain}(\phi)$ such that

$$\beta_s^{t+1} = \arg \min_{\beta_s} \phi(\beta_1^t, \dots, \beta_{s-1}^t, \beta_s, \beta_{s+1}^t, \dots, \beta_p^t), \quad \text{and} \quad \beta_k^{t+1} = \beta_k^t, \quad \forall k \neq s.$$

Now, the optimization problem in Eqn. (4) is equivalent to minimizing the Lagrangian

$$\min_{\mathbf{c}_1, \dots, \mathbf{c}_{|V|}} \phi(\mathbf{c}_1, \dots, \mathbf{c}_{|V|}) = g(\mathbf{c}_1, \dots, \mathbf{c}_{|V|}) - \sum_{i=1}^{|V|} \gamma_i^\top (\mathbf{CLP}\mathbf{c}_i - 4\pi\mathbf{1}), \tag{6}$$

where $\gamma_i \geq 0$. Since $g(\mathbf{c}_1, \dots, \mathbf{c}_{|V|})$ is the sum of two quadratic functions, it is differentiable and convex. In addition, since any affine function is convex (and concave), $\gamma_i^\top (\mathbf{CLP}\mathbf{c}_i - 4\pi\mathbf{1})$ is convex. Finally, the different \mathbf{c}_i 's, belonging to different voxels, do not overlap with each other and $\sum_{i=1}^{|V|} \gamma_i^\top (\mathbf{CLP}\mathbf{c}_i - 4\pi\mathbf{1})$ is separable. From Theorem 1, it is immediate to see that coordinate descent will converge to the minimizer for Eqn. (6) or equivalently, the minimizer for Eqn. (4). In addition, at each iteration, $\mathbf{c}_i^{t+1} = \arg \min_{\mathbf{c}_i} \phi(\mathbf{c}_1^t, \dots, \mathbf{c}_i, \dots, \mathbf{c}_{|V|}^t) = \arg \min_{\mathbf{c}_i} g(\mathbf{c}_1^t, \dots, \mathbf{c}_i, \dots, \mathbf{c}_{|V|}^t) - \gamma_i^\top (\mathbf{CLP}\mathbf{c}_i - 4\pi\mathbf{1})$. Therefore, we can solve for \mathbf{c}_i^{t+1} from the quadratic programming problem, $\mathbf{c}_i^{t+1} = \arg \min_{\mathbf{c}_i} g(\mathbf{c}_1^t, \dots, \mathbf{c}_i, \dots, \mathbf{c}_{|V|}^t) = \arg \min_{\mathbf{c}_i} \frac{1}{2} \|\mathbf{B}\mathbf{c}_i - \mathbf{s}_i\|^2 + \lambda \sum_{\|\mathbf{x}_i - \mathbf{x}_j\| < r} w_{ij} \|\mathbf{c}_i - \mathbf{c}_j^t\|^2$ subject to $-\mathbf{CLP}\mathbf{c}_i \leq 4\pi\mathbf{1}$. Algorithm 1 gives our ODF estimation method in detail. Even though Algorithm 1 estimates one \mathbf{c}_i in each iteration, it is possible to partition the problem to estimate a subset of \mathbf{c}_i 's simultaneously.

Algorithm 1. Nonnegative ODF Estimation with Spatial Regularization

Given the HARDI volume $\mathbf{s}_i = [\ln(-\ln(\frac{S_i(\theta_1, \phi_1)}{S_{0,i}})), \dots, \ln(-\ln(\frac{S_i(\theta_N, \phi_N)}{S_{0,i}}))]^\top$, $i = 1 \dots |V|$, and predefined parameters: spatial radius r defining the voxel connectivity, threshold ϵ and maximum number of iterations P .

1. Calculate the nonnegative weights w_{ij} .
2. First pass $p = 0$ through the entire volume:
 At each voxel \mathbf{x}_i , calculate using a quadratic programming solver,

$$\mathbf{c}_i^0 = \arg \min_{\mathbf{c}} \frac{1}{2} \|\mathbf{B}\mathbf{c} - \mathbf{s}_i\|^2, \quad \text{subject to } -\mathbf{CLP}\mathbf{c} \leq 4\pi\mathbf{1}.$$

3. Subsequent passes $p \geq 1$ through the entire volume:

a. Initialize $t = 1$. While $t \leq |V|$,

- i. Set $i = t$ and compute using a quadratic programming solver a new estimate for voxel \mathbf{x}_i ,

$$\mathbf{c}_i^{(p-1)|V|+t} = \arg \min_{\mathbf{c}} \frac{1}{2} \|\mathbf{B}\mathbf{c} - \mathbf{s}_i\|^2 + \lambda \sum_{\|\mathbf{x}_i - \mathbf{x}_j\| < r} w_{ij} \|\mathbf{c} - \mathbf{c}_j^{(p-1)|V|+(t-1)}\|^2,$$

$$\text{subject to } -\mathbf{CLP}\mathbf{c} \leq 4\pi\mathbf{1}.$$

- ii. For the remaining voxels, set $\mathbf{c}_j^{(p-1)|V|+t} = \mathbf{c}_j^{(p-1)|V|+(t-1)}$ for $\forall j \neq i$.
- iii. Set $t \leftarrow t + 1$.

- b. Set $p \leftarrow p + 1$. The stop criterion used is when $p = P$ or the decrease in cost function $g(\cdot)$ between the p^{th} and $(p - 1)^{th}$ iterations is less than ϵ .
-

4 Experiments

We present experiments on synthetic and real datasets using the proposed estimation method. We examine the quality of the estimated ODF given by: 1) the commonly used least-square (LS) estimate obtained by solving Eqn. (1), 2) the estimate obtained with nonnegativity constraint solely (QP) by solving Eqn. (4) with $\lambda = 0$, and 3) the estimate obtained with nonnegativity constraint and spatial regularization (QP-S) by solving Eqn. (4) with non-zero λ . ODFs are computed using a $l = 6^{th}$ order SH expansion in all our experiments. Note that if the least-square solution results in an ODF that have negative values, the common treatment is to set such values to a small positive number.

We first evaluate the performance on synthetic data generated using the multi-tensor method in [8]. The first synthetic experiment studies the improvement of the QP estimate over the LS estimate when the signal is corrupted by noise at a single voxel, in 100 trials. We first construct the true ODFs of 1, 2, and 3 fibers, as shown in Fig. 1(a), and the corresponding HARDI signals $\{S(\theta_i, \phi_i)\}_{i=1}^N$ at $N = 100$ gradient directions with $S_0 = 1$. Noisy versions of $\{S(\theta_i, \phi_i)\}_{i=1}^N$ are generated by adding complex Gaussian noise with zero mean and standard deviation $\sigma = \frac{S_0}{\zeta}$, where ζ is the signal-to-noise ratio (SNR). We use the Riemannian distance $\text{dist}_{Rie}(\cdot, \cdot)$ between probability density functions [18] to compare different ODFs. Fig. 1(b) shows the error $\text{dist}_{Rie}(\phi_t, \phi_e)$ between the true ODF ϕ_t and estimated ODF ϕ_e using LS and QP when the voxel contains 1, 2 or 3 fibers for varying levels of SNR. Notice that QP always gives a lower error than LS and does not give any negative values of p whereas LS results in 6%

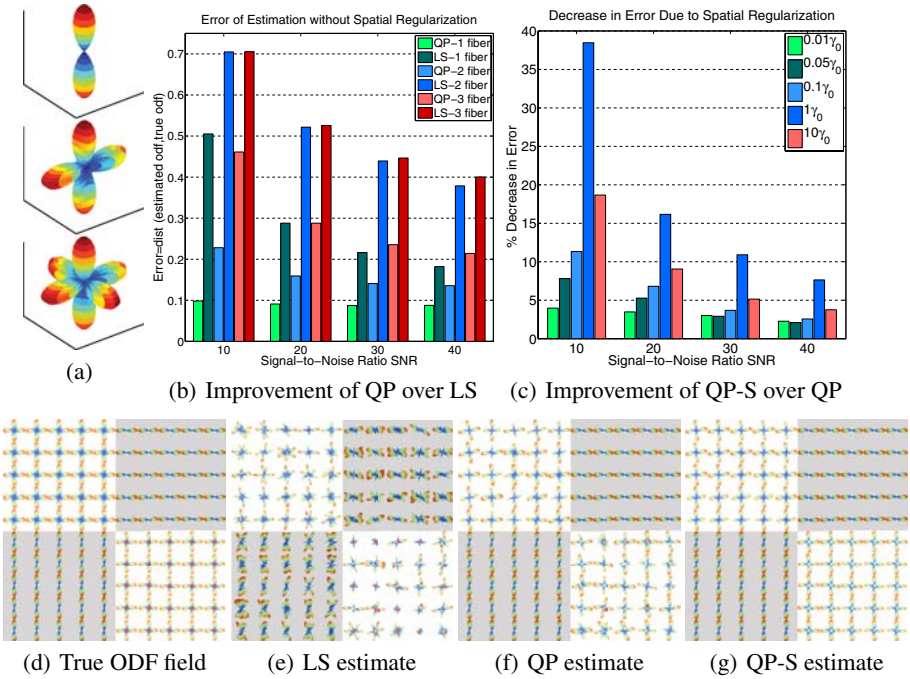



Fig. 1. Synthetic experiments. The first experiment (Figs. 1(a)-1(b)) studies the improvement of the QP estimate over the LS estimate when the signal is corrupted by noise. The second experiment (Figs. 1(c)-1(g)) studies the effect of having spatial regularization under varying values of regularization parameter λ for the ODF field shown in Fig. 1(d).

of the values being less than zero. In addition, the error increases when the number of fibers increases. This is expected as more acquisition directions are required to estimate an ODF with 3 fiber crossings compared to an ODF with a single fiber.


The second synthetic experiment shows the improvement of QP-S over QP in a ODF field over 100 trials. We construct an ODF field ϕ_t as shown in Fig. 1(d). The voxels in the 1st and 3rd quadrants contain 1 fiber, the 2nd quadrant 2 fibers and the 4th quadrant 3 fibers (with 1 fiber pointing out of the plane). We study the effects of varying the regularization parameter λ from $0.01\lambda_0$ to $10\lambda_0$, where $\lambda_0 = 1$. Fig. 1(c) shows the % decrease in error when estimation is done with QP-S compared to only doing QP, for varying SNR and λ . The error is measured as $\sum_x \text{dist}_{\text{Rie}}(\phi_t(x), \phi_e(x))$ where ϕ_e is the estimated ODF field. At low SNR, QP-S gives an estimated ODF field that is significantly closer to the true ODF field. When λ is too small or too large, the resulting ϕ_e is about the same quality of the LS estimate. Figs. 1(e)-1(g) show the estimated ODF fields for LS, QP, and QP-S with $p = 5$ passes of one trial at SNR= 10.

Finally, we apply our estimation method to a HARDI human brain dataset. Diffusion weighted MR images were obtained using the following imaging parameters: 55 axial slices (2mm thick), TR/TE=8250/92.3ms, with a 128×128 acquisition matrix (1.8mm in-plane resolution). 105 images were acquired, 11 with no diffusion




(a) LS estimate

(b) QP estimate



(c) QP-S with $p = 1$ pass estimate

(d) QP-S with $p = 2$ passes estimate



(e) Zoomed-out LS estimated ODF field

(f) Zoomed-out QP-S with $p = 2$ estimated ODF field

Fig. 2. Estimation results for real brain data. Fig. 2(a) is the LS estimate, Fig. 2(a) the QP estimate, Figs. 2(b) and 2(c) are the QP-S estimates with $p = 1$ and $p = 2$ passes through the volume, respectively. Figs. 2(e)-2(f) show the zoomed-out LS ODF field and QP-S with $p = 2$ estimated ODF field with the results of the red box shown in Figs. 2(a)-2(d). The ODFs are superimposed on top of the generalized fractional anisotropy maps.

sensitization and $N = 94$ diffusion weighted images at $b = 1159 \text{ s/mm}^2$. A portion of the results are shown in Fig. 2, where the ODFs are superimposed on top of the generalized fractional anisotropy maps. Figs. 2(a), 2(b), 2(c), and 2(d) show the estimates given by LS, QP, QP-S with $p = 1$ pass, and QP-S with $p = 2$ passes, respectively. Figs. 2(e) and 2(f) show the LS and QP-S with $p = 2$ estimated ODF fields of one slice where the zoomed-in results of the red box are shown in Figs. 2(a)-2(d). Notice that the ODF field estimated by LS is the noisiest, whereas QP-S with $p = 2$ gives a smoother ODF field and still preserves the discontinuities between different regions. In addition, the generalized fractional anisotropy map of QP-S shown in Fig. 2(f) is significantly cleaner than that of LS in Fig. 2(e) and the different regions of the brain can be seen more clearly.

5 Conclusion

We have presented an ODF estimation method that gives sharp diffusion ODFs, constrains the estimated ODF to be a proper pdf, and incorporates spatial regularization. Results on synthetic and real data demonstrate the advantage of working with our proposed algorithm. Future work will extend to multiple q-shell reconstruction method in [19].

References

1. Tuch, D.S.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine* 48, 577–582 (2002)
2. Basser, P., Mattiello, J., LeBihan, D.: Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance B* 103, 247–254 (1994)
3. Barmpoutis, A., Jian, B., Vemuri, B.C., Shepherd, T.M.: Symmetric positive 4th order tensors & their estimation from diffusion weighted MRI. In: Karssemeijer, N., Lelieveldt, B. (eds.) *IPMI 2007*. LNCS, vol. 4584, pp. 308–319. Springer, Heidelberg (2007)
4. Ghosh, A., Descoteaux, M., Deriche, R.: Riemannian framework for estimating symmetric positive definite 4th order diffusion tensors. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I*. LNCS, vol. 5241, pp. 858–865. Springer, Heidelberg (2008)
5. Leow, A.D., Zhu, S., Zhan, L., McMahon, K., de Zubicaray, G.I., Meredith, M., Wright, M.J., Toga, A.W., Thompson, P.M.: The tensor distribution function. *MRM* 61, 205–214 (2009)
6. Tuch, D.S.: Q-ball imaging. *Magnetic Resonance in Medicine* 52(6), 1358–1372 (2004)
7. Özarslan, E., Mareci, T.: Generalized diffusion tensor imaging and analytical relationships between DTI and high angular resolution diffusion imaging. *MRM* 50, 955–965 (2003)
8. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine* 58, 497–510 (2007)
9. Frank, L.R.: Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine* 47(6), 1083–1099 (2002)
10. Hess, C.P., Mukherjee, P., Han, E.T., Xu, D., Vigneron, D.B.: Q-ball reconstruction of multi-modal fiber orientations using the spherical harmonic basis. *MRM* 56(1) (2006)
11. Aganj, I., Lenglet, C., Sapiro, G.: ODF reconstruction in Q-ball imaging with solid angle consideration. In: *IEEE Int. Symposium on Biomedical Imaging* (2009)
12. Rathi, Y., Michailovich, O., Bouix, S., Shenton, M.: Orientation distribution estimation for Q-ball imaging. *MMBIA* (2008)

13. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI. *NeuroImage* 35(4), 1459–1472 (2007)
14. Jian, B., Vemuri, B.: A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE TMI* 26(11), 1464–1471 (2007)
15. Neji, R., Azzabou, N., Paragios, N., Fleury, G.: A convex semi-definite positive framework for DTI estimation and regularization. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Paragios, N., Tanveer, S.-M., Ju, T., Liu, Z., Coquillart, S., Cruz-Neira, C., Müller, T., Malzbender, T. (eds.) *ISVC 2007, Part I. LNCS*, vol. 4841, pp. 220–229. Springer, Heidelberg (2007)
16. Assemlal, H.E., Tschumperl'e, D., Brun, L.: Robust variational estimation of PDF functions from Diffusion MR signal. In: *CDMRI* (2008)
17. Tseng, P.: Coordinate ascent for maximizing nondifferentiable concave functions. Technical Report LIDS-P 1840, MIT, Laboratory for Information and Decision Systems (1988)
18. Goh, A., Lenglet, C., Thompson, P., Vidal, R.: A nonparametric Riemannian framework for processing high angular resolution diffusion images (HARDI). In: *IEEE CVPR* (2009)
19. Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., Harel, N.: Multiple q-shell ODF reconstruction in q-ball imaging. In: Yang, G.-Z., et al. (eds.) *MICCAI 2009, Part II. LNCS*, vol. 5762, pp. 423–431. Springer, Heidelberg (2009)