Generalized Constant Solid Angle ODF and Optimal Acquisition Protocol for Fiber Orientation Mapping

Amith Kamath[1], Iman Aganj[3,4], Junqian Xu[2], Essa Yacoub[2], Kamil Ugurbil[2], Guillermo Sapiro[5], and Christophe Lenglet[2]

[1] Electrical and Computer Engineering, University of Minnesota.
[2] CMRR, Department of Radiology, University of Minnesota.
[3] Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School.
[4] Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.
[5] Electrical and Computer Engineering, Duke University.
{kama0149@umn.edu}

Abstract. Recent advances in diffusion MRI have allowed for improved understanding of the white matter connectivity. Models like the Diffusion Tensor, diffusion Orientation Distribution Function (ODF) with a mono-exponential signal decay have shown good fiber reconstruction accuracies. More complex radial signal decay models, like the bi-exponential model, have been shown to better approximate the in-vivo diffusion signal. In this paper, we generalize the Constant Solid Angle ODF (CSA-ODF) algorithm to handle any q-space sampling and exploit the bi-exponential model. Simulation results to optimize the reconstruction and acquisition parameters are described. Finally, the algorithm is validated on human brain data. Our generalized CSA-ODF model performs optimally with 200 q-space data points distributed over three shells acquired at $b = 1000, 2000 s/mm^2$ and in the range $[3000, 6000] s/mm^2$ for the third shell. Crossings up to about 30 degrees can be recovered, and fiber orientations can be detected with a precision of about 1 degree.

1 Introduction

In the past decade, there has been significant interest in diffusion Weighted Magnetic Resonance (MR) Imaging for applications related to brain connectivity mapping [1,2]. This has been facilitated by faster MR acquisitions and improved signal modeling. The fiber orientation can be approximated using models like Diffusion Tensor Imaging (DTI) [3], Ball and Stick model [4], Diffusion Orientation Transform (DOT) [5], diffusion ODF [6] based models, Constrained Spherical Deconvolution (CSD) [7] and compartment-based models as in Panagiotaki et al. [8]. Aganj et al [9] introduced a correction of the original ODF model [6], called the CSA-ODF model, which implicitly sharpens the peaks, and is mathematically accurate. In this case, the fiber orientations are estimated using maxima extraction [10, 11] methods.

The utility of acquiring diffusion-weighted data over multiple q-space shells has been discussed previously by many groups [9, 12–16]. Diffusion Spectrum Imaging (DSI) [17] achieves a similar purpose by sampling q-space on a uniform 3D Cartesian grid. Current q-ball methods typically rely on sampling diffusion signals on a single shell in q-space, based on which a radial decay model can be estimated. This places a restriction on the types of signal decay models that can be used. Acquiring multiple signal values along a single gradient direction allows estimating more complex decay models like the bi-exponential model, which is known to be more accurate [18], but open new computational challenges.

Contributions: We explore an approach where the acquisition is done over a finite set of b-values in *staggered* gradient directions, so as to allow a better angular resolution, along with providing radial information of the diffusion signal decay. Given a finite number of measurements on a particular shell, we investigate how to accurately interpolate such spherical data to be able to retrieve samples in any orientation. Computationally efficient methods in estimating the ODF have been explored in detail in [19]. An intelligent choice of Spherical Harmonic (SH) bases allows for a linear and regularized algorithm. We further refine previous investigations on the optimal choice of regularization for SH fitting.

The CSA-ODF model [9] has been shown to perform better than "artificially" sharpened models, and this work focuses on extending this method to support generic sampling of q-space. We study the optimal set of b-values for the multi-shell acquisition for this new generalized CSA-ODF algorithm. We also investigate the minimum number of gradient directions to use in order to minimize distortions on ODF reconstruction and uncertainty on fiber estimation. We finally investigate the minimum achievable separating angle between fibers and the uncertainty on the localization of fiber orientation with the proposed acquisition and reconstruction protocol.

The rest of the paper is organized as follows. A short overview of the CSA-ODF model is presented in Section 2. We then describe the extension of this algorithm to generic q-space sampling. Section 3 includes extensive simulation studies and results. This is followed by results to validate the algorithm on human brain data.

2 Methods

2.1 q-ball Estimation

The q-ball method [6] estimates fiber orientations from high angular resolution diffusion imaging (HARDI) data. q-ball imaging is a natural choice for fiber orientation estimates, as the data is acquired on a spherical shell with different orientations rather than on a grid in q-space. Using more complex radial decay models can provide a better understanding of the diffusion process and hence the axonal configuration of the white matter.

A modified SH basis was used in [19–21] for representing spherical distributions. This leads to an analytic computation of the Funk Radon Transform (FRT) where the ODF is obtained as a linear combination of SH bases. The CSA-ODF model [9] corrects a mathematical inaccuracy, due to which the ODF

was originally measured over a cylindrical profile rather than within the intended solid angle Ω. The SH coefficients usually denoted by Y_l^m (l is the order, m is the phase factor), are the basis functions for complex functions on the sphere. It is assumed that ODFs are antipodally symmetric, and consequently, even values of l are used to decompose the data. Increasing the order of approximation reduces the width of the ODF lobes and hence increases the angular resolution as well as, to some extent, sensitivity to noise. With this kind of basis functions, the normalized signal $S(\theta_i, \phi_i)$, for K points '$i = 1, 2, ...K$' on the q-shell is decomposed as

$$S(\theta_i, \phi_i) = \Sigma_{j=1}^{N} c_j Y_j(\theta_i, \phi_i) \tag{1}$$

where j indexes the even ordered SH bases, for '$l = 0, 2, 4, ...l_{max}$' and '$m = -l, ...0, ...l$'. c_j are the N SH coefficients. These coefficients completely determine the shape of the ODF.

For discrete samples on a single shell, this can be written as an overdetermined linear system, where the number of gradient directions along which the signal is acquired is higher than N. The SH coefficients are usually estimated using a least-squares minimization procedure, but this is known to be sensitive to noise, especially preponderant in high b-value datasets with lower SNR. Several regularization schemes have therefore been suggested, especially Tikhonov [20] and Laplace Beltrami regularizations [19], which smooth the reconstruction, at the expense of losing sharp features of the fiber distribution. This trade-off is essential in the interpolation of the signal values for multi-shell reconstruction, and a detailed analysis is presented in Section 3.2.

2.2 Multishell structure

Diffusion MR signal acquisitions are characterized by two important parameters: the b-value(s) (typically specified in s/mm^2) and the gradient table. For typical diffusion levels in the human brain ($0.0021mm^2/s$) [21], it has been observed that a single b-value in the order of $2200 - 2800s/mm^2$ is optimal to reconstruct crossing fibers [21]. The gradient table selects points on the shell and encodes the spatial coordinates of points where the signal is acquired.

This type of sampling technique can be extended to multiple concentric shells of different b-values [22], so that the diffusion signal is measured at different levels of decay. Keeping the orientation fixed, as the b-value is varied, signal samples can be measured, which allows a straightforward estimation of higher order diffusion models along each orientation independently. We refer to this acquisition scheme as the *aligned* multi-shell scheme. It is also possible to distribute points uniformly (and incrementally) over several shells in order to maximize angular coverage, as demonstrated in [23]. In this setting, a wider range of orientations is covered. We refer to this acquisition scheme as the *staggered* multi-shell scheme [23].

2.3 Generalized CSA-ODF Algorithm

We generalize the previous work [9] on CSA-ODF by i) optimally fitting data on each shell using SH series and adaptive regularization, and ii) handling any

combination of b-values to fit a bi-exponential model (when at least three shells are present). We limit the discussion in this section to three shells, but this approach can be used over any number of shells, as illustrated in the experimental section. Once the signal S, sampled in K gradient directions has been fitted to a SH series of a given order (4, 6, or 8 in general), yielding N coefficients c_i, it is straightforward to estimate the signal value at a new orientation $\hat{z}_j$ using Equation (1). We investigate in Section 3.2 optimal regularization choices at various SNR levels, SH orders and for various fiber configuration using Generalized Cross-Validation (GCV) [24].

With signal values available in any direction, it becomes possible to estimate the three parameters of a bi-exponential model. Aganj et al. [9] derive an analytic formula to estimate these parameters, under the constraint that the shells must have b-values in arithmetic progression. We generalize this to any set of b-values, using a trust region based optimization scheme (as implemented in MATLAB). The signal decay along a given direction is thus fitted to a bi-exponential model as follows: Along $\hat{z}_j$, the signal is estimated as

$$X_{\hat{z}_j}(b) = \alpha_{\hat{z}_j} e^{-\delta^{(1)}_{\hat{z}_j} b} + (1 - \alpha_{\hat{z}_j}) e^{-\delta^{(2)}_{\hat{z}_j} b} \tag{2}$$

where b is the b-value, and $\alpha_{\hat{z}_j}$ is the volume fraction of the first decay component. $\delta^{(1)}_{\hat{z}_j}$ and $\delta^{(2)}_{\hat{z}_j}$ are the decay rates for the corresponding compartment along $\hat{z}_j$ usually referred to as "slow" and "fast" compartment [18]. The ODF can now be calculated as

$$ODF(\hat{z}_j) = \frac{1}{4\pi} + \frac{1}{16\pi^2} \mathrm{FRT}\{\nabla^2_b (\alpha_{\hat{z}_j} ln(\delta^{(1)}_{\hat{z}_j}) + (1 - \alpha_{\hat{z}_j}) ln(\delta^{(2)}_{\hat{z}_j}))\} \tag{3}$$

The SH coefficients c_i of the double-logarithm of the signal are computed such that

$$\alpha_{\hat{z}_j} ln(\delta^{(1)}_{\hat{z}_j}) + (1 - \alpha_{\hat{z}_j}) ln(\delta^{(2)}_{\hat{z}_j}) = \Sigma_{i=1}^{N} c_i Y_i(\theta_j, \phi_j) \tag{4}$$

The SH coefficients c_{ODF} of the ODF are then obtained through a linear transformation of the c_i as described in [9]. We describe the method in detail in Algorithm 1. K signal acquisitions (S) are supposed available along K gradient directions encoded by *bvecs*. The i^{th} element of shellID encodes the b-value at which the i^{th} acquisition in S was made. T is the matrix containing values taken by each considered spherical harmonic along each direction in *bvecs*. S_{interp} is the interpolated signal on each q-shell. α and $\delta^{(1,2)}$ are the parameters of the bi-exponential model. X is the estimated bi-exponential decay for each direction, and c_{ODF} is the SH coefficients of the ODF thus estimated. In order to illustrate the performance of the bi-exponential model fit, we provide in Fig. 1 the signal decay profile as a function of b-values along various directions for a single- and two-fiber configurations. The cylinder-dot-zeppelin compartment model [8] was used to generate data at SNR 15, along 300 directions. Along the fiber direction of the single fiber model, when averaged over the closest 30 datapoints, we observe a fast diffusivity of $0.0017 mm^2/s$ with a volume fraction of 0.525 and a slow diffusivity of $0.0005 mm^2/s$ with a volume fraction of 0.474. Across the

Algorithm 1 : Generalized CSA-ODF

Input: $S_{(K,1)}$, bvecs$_{(K,3)}$, b-value$_{(3,1)}$ (for 3 shells), ShellID$_{(K,1)}$.
Output: c_{ODF}
for $i = 1$ to 3 **do**
  $S_i \leftarrow S$ where ShellID == b-value(i)
  $c_i = (T_i^T T_i + \lambda_i L)^{-1} T_i^T S_i$
  $S_{\text{interp}_i} \leftarrow T_J * c_i$ {where J indexes an arbitrary gradient table.}
end for
for $j = 1$ to number of directions in J **do**
  $X_j(\alpha_j, \delta_j^{(1)}, \delta_j^{(2)}) \leftarrow$ bi-expfit(S_{interp}) {bi-expfit as defined in Equation (2)}
end for
$c_{\text{ODF}} \leftarrow \text{ODF}((T_J^T T_J)^{-1} T_J^T X_J)$ {the ODF function as defined in Equation (3)}.

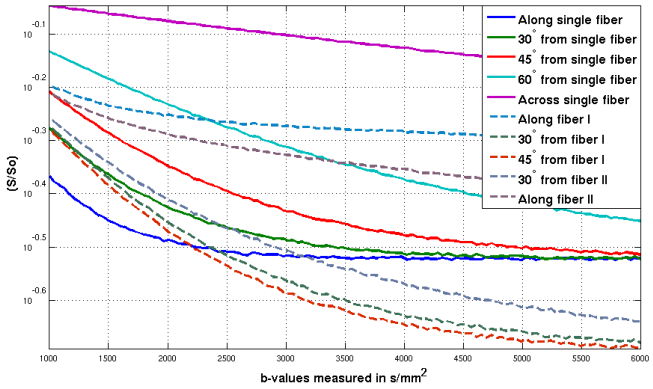

Fig. 1: Bi-exponential decay for the single- and two-fiber configurations.

fiber direction, when averaged over the closest 30 datapoints, we observe a fast diffusivity of $0.0029mm^2/s$ with a volume fraction of 0.296 and a slow diffusivity of $0.0002mm^2/s$ with a volume fraction of 0.68. These values are of the same order as the parameters we use to simulate the data. Moreover, we verify [8] that diffusion along the cylinder is bi-exponential, while largely monoexponential across the cylinder. Similar results were found in the two-fiber case along both cylinders (bi-exponential) and in the orthogonal direction.

3 Protocol Optimization

3.1 Simulation setup

We rely on realistic diffusion models available in Camino [25]. We use three compartments for single-fiber configurations and five compartments for two-fiber configurations, with volume fractions 0.6 for the intra-axonal, 0.1 for the extra-axonal, and 0.3 for the isotropic compartment, as described in [8]. This is scaled appropriately for the two-fiber case. The intra-axonal compartment is modeled as a cylinder of radius $0.004mm$ with a gaussian phase distribution (GPD) and diffusivity $0.0017mm^2/s$. The extra-axonal compartment is modeled

as a zeppelin with diffusivity $0.0017mm^2/s$ along the fiber, and $0.0002mm^2/s$ across it. Finally, the isotropic component is modeled as a dot. Eleven b-values: $[1000, 1500, 2000, ...6000]s/mm^2$ are calculated using the following pulse sequence parameters G = 0.05 T/m, TE = $0.1s$, $\delta = 0.02s$ and varying Δ. In all experiments, we choose a maximum number of 300 gradient directions divided across three shells. For ground truth signal and associated ODF, we choose 300 directions on eleven shells (from the 11 b-values), at SNR 40. As a measure of similarity of ODFs, the L^2 norm of the difference in the SH coefficients is used.

3.2 Optimal regularization analysis:

The first step of the generalized CSA-ODF algorithm creates a continuous representation of the diffusion signal on each shell. To that end, we fit SH series of order 4, 6 or 8. The regularized least squares solution for the inverse problem of estimating the SH coefficients is given by:

$$c = (T^T T + \lambda I)^{-1} T^T S \tag{5}$$

where c is the vector of SH coefficients, S is the normalized signal, I is the identity matrix, and λ is the regularization parameter.

Methods to estimate the optimal value for λ have been discussed in [24, 26, 27]. The need for regularization arises from the small singular values of T. Small variations in S leads to high variation in the SH coefficients c. Regularized inversion is parameterized by λ which ideally is between the highest and lowest singular value of T. Laplace-Beltrami regularization has been extensively used

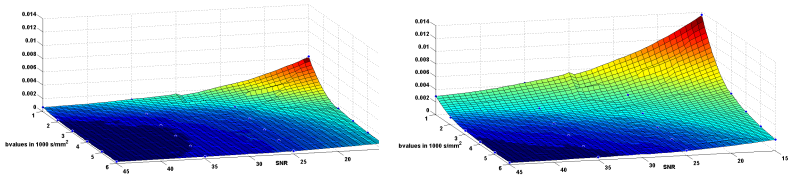

Fig. 2: Surface plot of λ for order 4 (left), and order 6 (right) for the two-fiber configuration as a function of SNR and b-value.

in the context of ODF reconstruction mostly because of the simple form of the Laplace-Beltrami operator for functions described by SH series [19]. It is controlled by replacing the identity matrix in Equation (5) by L (the square of the Laplace Beltrami matrix) whose diagonal elements are $(l+1)^2 l^2$, where l is the order of the SH at the corresponding column.

λ needs to be adjusted as a function of the acquisition and reconstruction parameters. We carried out an exhaustive investigation of the optimal choice of regularization for eleven b-values in the range $[1000, 1500, ...6000]s/mm^2$, five SNR=$5, 15, 25, 35$, and 45, three SH reconstruction orders of 4, 6 and 8, and finally, three fiber configurations with 1, 2 or 3 orientations. The optimal λ value is calculated using the GCV method [24], rather than the L-curve method, since the L-curve method assumes the discrete Piccard conditions [26], which is not

Table 1: Optimal λ values for three-fiber, order 6 reconstruction (b-value in s/mm^2).

b-value	SNR: 5	SNR: 15	SNR: 25	SNR: 35	SNR: 45
1000	0.083302	0.014496	0.011214	0.006020	0.004279
1500	0.038631	0.008837	0.004006	0.001763	0.002086
2000	0.047775	0.006534	0.002520	0.001391	0.001055
2500	0.017271	0.004824	0.001556	0.000823	0.000573
3000	0.021266	0.002089	0.001388	0.000670	0.000429
3500	0.026	0.002072	0.001057	0.000470	0.000318
4000	0.027045	0.001593	0.0006	0.00037	0.000319
4500	0.01063	0.002516	0.000865	0.000676	0.000243
5000	0.024093	0.001576	0.000795	0.000448	0.000453
5500	0.022867	0.001241	0.000788	0.000383	0.000391
6000	0.020809	0.001715	0.001317	0.000419	0.000353

true for Equation (5). Fig. 2 shows surface plots of λ, for varying b-values and SNRs. Two SH orders are represented in the case of a two-fiber configuration. Each data point is obtained as the average optimal λ for 100 repetitions with randomly selected angles between fibers. Table 1 provides numerical values in the three-fiber case. The trends show that lesser regularization is necessary for higher SNRs, and that the optimal λ can deviate quite significantly from the commonly used 0.006 value [19], especially at SNR=25 and b-values greater than $1500s/mm^2$. Interestingly, at all orders, regularization slightly decreases with increased b-value for a given SNR. This behavior is consistent across the three investigated configurations. Eight other tables similar to Table 1 are available, therefore making it possible to adjust the choice of λ to specific experimental conditions, as done in the rest of the paper. We note that SH parameterization of the ODF (Equation (3)) is done without regularization.

3.3 Optimal b-value analysis:

We study the effects of varying the sets of b-values to accurately estimate the ODF. We restrict ourselves to fourteen representative sets of three b-values. We will demonstrate in the following that adding a fourth shell only improve performance marginally. We restrict the range up to a b-value of $6000s/mm^2$ as most of the bi-exponential nature of the diffusion signal is concentrated in the range of $1500 - 4000s/mm^2$. We fix the total number of gradient directions to 300 (distributed as 21, 86, 193 over the shells) [23]. All SNRs, single- and two-fiber configurations and SH orders 4 and 6 are tested. The corresponding optimal λ is always chosen from Section 3.2. In each case, as for all experiments from now on, 20 repetitions are performed and the mean result reported. Results of this simulation are shown in Fig. 3. The set $[1000, 2000, 6000]s/mm^2$ shows lower error at SNR=15, 25 or 40 for an order 4 reconstruction and one-fiber configuration. Similar trends are observed at SNR 5 although, as expected, with much higher reconstruction error levels. Results are consistent with Fig. 1, where the signal decay exhibits the highest deviation from the mono-exponential behavior in the vicinity of b-value $2000s/mm^2$. b-value sets that skip this critical region show errors that are considerably larger as they do not capture the bi-exponential nature of the diffusion signal. b-value sets $[1000, 2000, 3000/4000/5000]s/mm^2$

exhibit similar performance but we focus on the set $[1000, 2000, 6000] s/mm^2$ to demonstrate the performance of the bi-exponential fitting algorithm in the most non-uniform sampling case. b-value $[1000, 2000, 6000]$ s/mm^2 show supe-

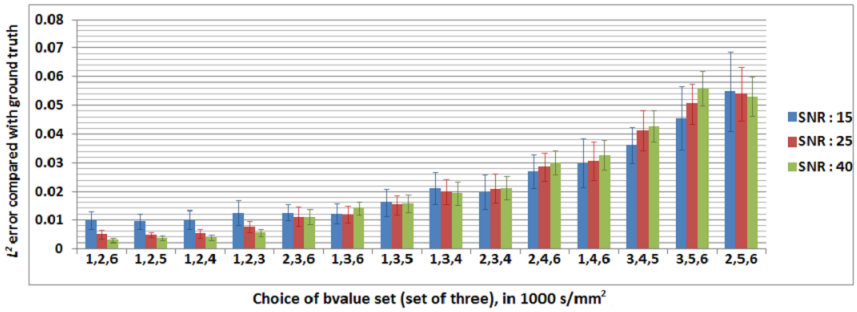

Fig. 3: Reconstruction error for order 6, single-fiber configuration.

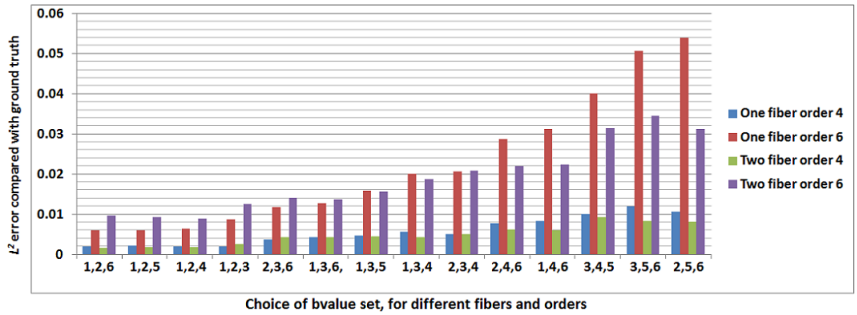

Fig. 4: Average reconstruction error for four combinations of SH order and fibers.

Table 2: Reconstruction error for varying order, SNR and fiber configurations for two representative b-value sets.

		Order: 4		Order: 6	
		one-fiber	two-fiber	one-fiber	two-fiber
SNR: 5	b=1,2,6	0.03383	0.02842	0.04709	0.03997
	b=1,4,6	0.03808	0.02938	0.05965	0.04820
SNR: 25	b=1,2,6	0.00138	0.00111	0.00498	0.00868
	b=1,4,6	0.00791	0.00577	0.03063	0.02152

rior reconstruction in all our tests as seen in Fig. 4. Clearly, b-value sets including 1000 and $2000 s/mm^2$ and with $3000, 4000, 5000$ or $6000 s/mm^2$ as third shell outperform other combinations. We will consider $[1000, 3000, 5000]$ and $[1000, 4000, 6000] s/mm^2$ respectively as "average" and "poor" candidates for subsequent comparisons. For completeness, Table 2 demonstrates, for "good" and "poor" choices of b-value, the relative insensitivity of the reconstruction to fiber configurations (a very desirable property for brain data) as well as the difference in error (one order of magnitude) between SNR 5 and 25.

3.4 Optimal number of gradient directions:

Reconstruction error is studied for a reduced number of total points, with all the other parameters remaining unchanged. This provides a quantitative measure of the loss in reconstruction accuracy, in comparison with the ground truth. Fig. 5 (for a two-fiber configuration at SNR 25) shows how the error decreases as the number of sample points increases. It is seen that the curves are nearly flat after around 200 directions. Moreover, b-value sets previously identified as "good", "average" and "poor" preserve their relative performance with increasing number of directions. Additionally, we compare the *aligned* sampling scheme

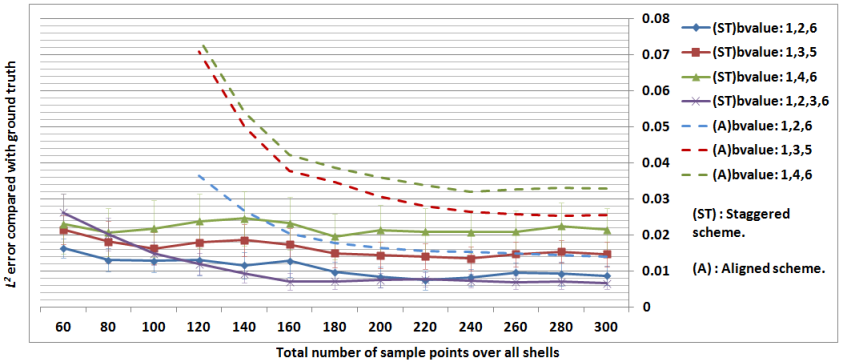

Fig. 5: Reconstruction error as a function of the number of directions acquired for different cases of b-value for two-fiber configuration at SNR 25, order 6 reconstruction.

with the *staggered* scheme for the same total number of data points over all shells, e.g: 300 *staggered* points divided as 21,86 and 193 are compared with 300 *aligned* points divided as 100, 100 and 100. Fig. 5 shows the reconstruction errors in both cases. The interpolation step in the *staggered* case does not introduce significant distortions. Moreover, this approach clearly provides a better angular resolution and improves ODF reconstruction. Fig. 5 also shows that a four-shell $[1000, 2000, 3000, 6000]s/mm^2$ acquisition does not provide an appreciable increase in accuracy. It therefore illustrates that three points are likely sufficient to estimate the three parameters α, $\delta^{(1)}$ and $\delta^{(2)}$.

3.5 Angular resolution and error

We investigate the minimum achievable separating angle between fiber orientations using the *staggered* acquisition scheme with 200 directions, b-values of $[1000, 2000, 6000]s/mm^2$, SH order 8 and SNR 40. It has been demonstrated that angles around 30 degrees can be recovered when using order 8 SH series [20]. Fig. 6 shows the reconstruction in steps of 10 degrees from orthogonal to a single-fiber configuration, with an order 8 reconstruction. The * indicates the minimum angle between which we resolve crossings, which precisely happens between 35 and 30 degrees. Table 3 summarizes the angular error achieved with the same sampling scheme but varying SH orders and SNR for an orthogonal two-fiber configuration. For each case, 100 repetitions were performed and the mean $\pm$ standard deviation of the angular error between the estimated ODF

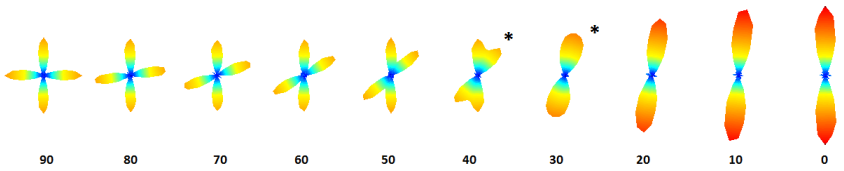

Fig. 6: Angular resolution in the two-fiber case, varying from orthogonal-fiber to a single-fiber configuration.

Table 3: Angular error performance (degrees) for an orthogonal-fiber configuration.

mean ± std	SNR: 5	SNR: 15	SNR: 25	SNR: 40
Order: 4	5.3759 ± 2.0902	1.5826 ± 0.6079	1.0886 ± 0.4610	0.7299 ± 0.2736
Order: 6	5.4046 ± 2.1040	1.5992 ± 0.6564	1.1093 ± 0.4642	0.7463 ± 0.2725
Order: 8	5.4309 ± 2.1006	1.6184 ± 0.6215	1.0920 ± 0.4685	0.7356 ± 0.2769

maxima [10] and the true orientation are reported. Even at low SNR, the maximum error is around 5 degrees. For realistic SNR, an average error of 1 degree is achieved which outperforms many existing ODF reconstruction technique [7,28].

4 Application to human brain data

Finally, our generalized CSA-ODF algorithm was applied to human brain data from a healthy volunteer [29], acquired on a 3T Siemens scanner with the following parameters: Voxel size $1.5mm$ isotropic, TR/TE=3200/77ms, 133 aligned DWI at b= $[1000, 2000, 3000]s/mm^2$ and 10 $b0$. ODF reconstructions are presented in Fig. 7. Fiber crossings in the centrum semiovale and various other complex white matter areas are successfully recovered. Moreover, we down-sampled the data by half (from 133 to 67 directions) and verified (Fig. 7) that the ODF reconstruction still performed very well. This is expected from our simulations, although a staggered scheme was not available for this dataset (in which case we would expect even better results).

5 Conclusion

We have generalized the CSA-ODF model to account for a flexible acquisition scheme, and provided an analysis of the optimal parameters for the reconstruction of the ODF. It is observed that specific sets of b-values (e.g. $[1000, 2000, 6000]$ s/mm^2) perform clearly better, which can be explained by looking deeper into the bi-exponential nature of the diffusion signal decay. For this particular reconstruction algorithm, three-shell acquisitions perform nearly as well as four-shell acquisitions, and around 200 gradient directions are sufficient for good angular resolution as well as accuracy.

Acknowledgements:

This work was partly supported by NIH grants R01 EB008432, P41 RR008079 (NCRR), P41 EB015894 (NIBIB), P30 NS057091, P30 NS5076408, and the Human Connectome Project (U54 MH091657) from the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research.

References

1. K. K. Seunarine and D. C. Alexander. Chapter 4 - multiple fibers: Beyond the diffusion tensor. In *Diffusion MRI*, pages 55 – 72. Academic Press, 2009.

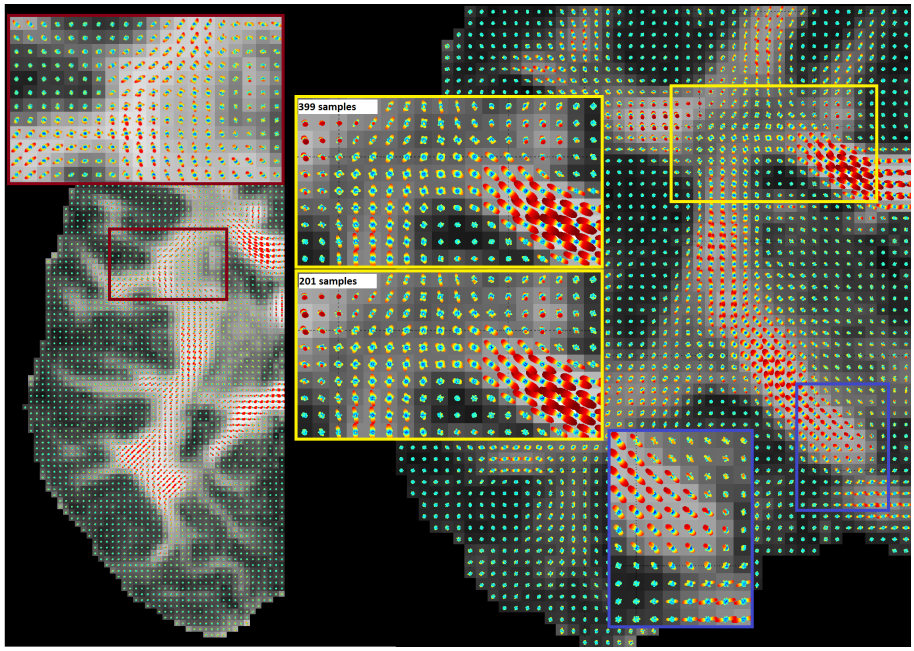

Fig. 7: The generalized CSA-ODF algorithm applied to brain data, superimposed on a Generalized Fractional Anisotropy (GFA) map with an *aligned* sampling scheme.

2. D.C. Alexander. An introduction to computational diffusion MRI: the diffusion tensor and beyond. *Visualization and processing of tensor fields*, pages 83–106, 2006.
3. P. J. Basser, J. Mattiello, and D. L. Bihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
4. T. E. J. Behrens, M. W. Woolrich, M. Jenkinson, H. Johansen-Berg, R. G. Nunes, S. Clare, P. M. Matthews, J. M. Brady, and S. M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn. Reson. Med.*, 50(5):1077–1088, 2003.
5. E. Ozarslan, T. M. Shepherd, B. C. Vemuri, S. J. Blackband, and T. H. Mareci. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *NeuroImage*, 31(3):1086 – 1103, 2006.
6. D. S. Tuch. Q-ball imaging. *Magn. Reson. Med.*, 52(6):1358–1372, 2004.
7. J.-D. Tournier, C.-H. Yeh, F. Calamante, K.-H. Cho, A. Connelly, and C.-P. Lin. Resolving crossing fibres using constrained spherical deconvolution: Validation using diffusion-weighted imaging phantom data. *NeuroImage*, 42(2):617 – 625, 2008.
8. E. Panagiotaki, T. Schneider, B. Siow, M. G. Hall, M. F. Lythgoe, and D. C. Alexander. Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. *NeuroImage*, 59(3):2241 – 2254, 2012.
9. I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel. Reconstruction of the orientation distribution function in single-and multiple-shell q-ball imaging within constant solid angle. *Magn. Reson. Med.*, 64(2):554–566, 2010.
10. I. Aganj, C. Lenglet, and G. Sapiro. ODF maxima extraction in spherical harmonic representation via analytical search space reduction. *MICCAI*, pages 84–91, 2010.

11. A. Ghosh and R. Deriche. Extracting geometrical features & peak fractional anisotropy from the ODF for white matter characterization. In *ISBI*, pages 266–271, 2011.
12. S. Jbabdi, S. N. Sotiropoulos, A. M. Savio, M. Grana, and T. E. J. Behrens. Model-based analysis of multishell diffusion MR data for tractography: How to get over fitting problems. *Magnetic Resonance in Medicine*, 2012.
13. M. Descoteaux, R. Deriche, D. L. Bihan, J.-F. Mangin, and C. Poupon. Multiple q-shell diffusion propagator imaging. *Medical Image Analysis*, 15(4):603 – 621, 2011. Special section on IPMI 2009.
14. M. H. Khachaturian, J. J. Wisco, and D. S. Tuch. Boosting the sampling efficiency of q-ball imaging using multiple wavevector fusion. *Magn. Reson. Med.*, 57(2):289–296, 2007.
15. H.-E. Assemlal, D. Tschumperlé, and L. Brun. Evaluation of q-space sampling strategies for the diffusion magnetic resonance imaging. MICCAI '09, pages 406–414, 2009.
16. A. Tristán-Vega, S. Aja-Fernández, and C.-F. Westin. On the blurring of the funk—radon transform in q—ball imaging. MICCAI '09, pages 415–422, 2009.
17. V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn. Reson. Med.*, 54(6):1377–1386, 2005.
18. Y. Assaf and Y. Cohen. Chapter 7 - inferring microstructural information of white matter from diffusion MRI. In *Diffusion MRI*, pages 127 – 146. 2009.
19. M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Regularized, fast, and robust analytical q-ball imaging. *Magn. Reson. Med.*, 58(3):497–510, 2007.
20. C. P. Hess, P. Mukherjee, E. T. Han, D. Xu, and D. B. Vigneron. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magn. Reson. Med.*, 56(1):104–117, 2006.
21. D.C. Alexander and G.J. Barker. Optimal imaging parameters for fiber-orientation estimation in diffusion MRI. *Neuroimage*, 27(2):357–367, 2005.
22. L. Zhan, A. D. Leow, I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, N. Harel, A. W. Toga, and P. M. Thompson. Differential information content in staggered multiple shell HARDI measured by the tensor distribution function. In *ISBI'11*, pages 305–309, 2011.
23. E. Caruyer, J. Cheng, C. Lenglet, G. Sapiro, T. Jiang, and R. Deriche. Optimal Design of Multiple Q-shells experiments for Diffusion MRI. In *MICCAI Workshop on CDMRI'11*, September 2011.
24. G. H. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):pp. 215–223, 1979.
25. P.A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D.C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. *14th scientific meeting of the ISMRM*, 2759, 2006.
26. P. C. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34(4):pp. 561–580, 1992.
27. E. Caruyer and R. Deriche. Optimal Regularization for MR Diffusion Signal Reconstruction. In *ISBI*, May 2012.
28. K.-H. Cho, C.-H. Yeh, J.-D. Tournier, Y.-P. Chao, J.-H. Chen, and C.-P. Lin. Evaluation of the accuracy and angular resolution of q-ball imaging. *NeuroImage*, 42(1):262 – 271, 2008.
29. J. L. R. Andersson, J. Xu, E. Yacoub, E. Auerbach, S. Moeller, and K. Ugurbil. A comprehensive gaussian process framework for correcting distortions and movements in diffusion images. *ISMRM Annual Meeting*, May 2012.