Investigating the uncertainty in multi-fiber estimation in High Angular Resolution Diffusion Imaging

Liang Zhan[1] , Alex D. Leow[1] , Marina Barysheva[1], Albert Feng[1], Arthur W. Toga[1], Guillermo Sapiro[2], Noam Harel[2], Kelvin O. Lim[2], Christophe Lenglet[2], Katie L. McMahon[3], Greig I. de Zubicaray[3], Margaret J. Wright[4], Paul M. Thompson[1]

[1] Laboratory of Neuroimaging, Department of Neurology, University of California, Los Angeles, USA
[2] Center for Magnetic Resonance Research, University of Minnesota, USA
[3] Functional MRI Laboratory, Centre for Magnetic Resonance, University of Queensland, Brisbane, Australia
[4] Queensland Institute of Medical Research, Brisbane, Australia

Abstract. In this paper, we investigated the reconstruction accuracy and information uncertainty in multi-fiber estimation to better understand the trade-off between scanning time and angular precision in High Angular Resolution Diffusion Imaging (HARDI). Reconstruction accuracy was measured using the Kullback-Leibler divergence (sKL) on the orientation density functions (ODFs) first in simulations with varying b-values and variable additive Rician noise. ODFs were computed analytically from tensor distribution functions (TDFs) which model the HARDI signal at each point as a unit-mass probability density on the 6D manifold of symmetric positive definite tensors. Reconstruction accuracy rapidly increased with additional gradients at lower SNR. The information uncertainty was quantified by the Exponential Isotropy (EI), a TDF-derived measure of fiber integrity that exploits the full multidirectional HARDI signal. Simulations and empirical results both found that information uncertainty decreased as angular resolution increased, and plateaued at around 70~80 gradients. Furthermore, in high magnetic field (7 Tesla) HARDI, the reconstruction accuracy and information uncertainty index decreased at higher b-values.

Keywords: High Angular Resolution Diffusion Imaging, Tensor Distribution Function, multi-fiber reconstruction, Kullback-Leibler divergence, Exponential Isotropy

1 Introduction

Diffusion weighted MR imaging is a powerful tool to study water diffusion in tissue, providing vital information on white matter microstructure, such as fiber connectivity and composition in the healthy and diseased brain. To date, most diffusion imaging studies (especially in clinical applications) still employ the diffusion tensor imaging (DTI) model [1], which describes the anisotropy of water diffusion in tissues by estimating, from a set of K diffusion-sensitized images, the 3x3 diffusion tensor (the covariance matrix of a 3-dimensional Gaussian distribution). Each voxel's signal

intensity in the *k*-th image is attenuated, by water diffusion, according to the Stejskal-Tanner equation [2]: $S_k = S_0 \exp[-b\mathbf{g_k}^T\mathbf{D}\mathbf{g}_k]$, where S_0 is the non-diffusion weighted signal intensity, **D** is the 3x3 diffusion tensor, $\mathbf{g_k}$ is the direction of the diffusion gradient and b is Le Bihan's factor containing information on the pulse sequence, gradient strength, and physical constants. Although 7 gradients are mathematically sufficient to determine the diffusion tensor, MRI protocols with higher angular and radial resolutions, such as the high-angular resolution diffusion imaging (HARDI) or diffusion spectrum imaging (DSI) techniques, can resolve more complex diffusion geometries that a single-tensor model, as employed in standard DTI, fails to capture, e.g., fiber crossings and intermixing of tracts.

Recent technical advances have made HARDI more practical. A 14-minute scan can typically sample over 100 angles (with 2 mm voxels at 4 Tesla). HARDI's improved signal-to-noise ratio may be used to reconstruct fiber pathways in the brain with extraordinary angular detail, identifying anatomical features, connections and disease biomarkers not seen with conventional MRI. If more angular detail is available, fiber orientation distribution functions (ODFs) may be reconstructed from the raw HARDI signal using the Q-ball imaging technique [3]. Deconvolution methods [4,5] have also been applied to HARDI signals, yielding mathematically rich models of fiber geometries as probabilistic mixtures of tensors [6], fields of von Mises-Fisher mixtures [7], or higher-order tensors (i.e., 3x3x…x3 tensors) [8,9]. Stochastic tractography [10, 11] can also exploit HARDI's increased angular detail, and fluid registration methods have also been developed to align HARDI ODFs using specialized Riemannian metrics [12]. In most deconvolution-based methods, however, restrictive prior assumptions are typically imposed on the allowable fibers, e.g., all fiber tracts are considered to have the same anisotropy profile.

A novel approach, the Tensor Distribution Function (TDF), was recently proposed by Leow et al. in [13] to model multidirectional diffusion at each point as a probabilistic mixture of all symmetric positive definite tensors. The TDF models the HARDI signal more flexibly, as a unit-mass probability density on the 6D manifold of symmetric positive definite tensors, yielding a TDF, or continuous mixture of tensors, at each point in the brain. From the TDF, one can derive analytic formulae for the orientation distribution function (ODF), tensor orientation density (TOD), and their corresponding anisotropy measures. Because this model can accurately resolve sharp signal peaks in angular space where fibers cross, we studied how many gradients are required in practice to compute accurate orientation density functions - as more gradients require longer scanning times. In this paper, we assessed how many diffusion-sensitized gradients were sufficient to (1) accurately resolve the diffusion profile, measured by the Kullback-Leibler divergence (sKL) and (2) achieve a satisfactory information uncertainty index, quantified by the exponential isotropy (EI), a TDF-derived measure of fiber integrity that exploits the full multidirectional HARDI signal. We used simulation data generated from two-fiber systems crossing at 90 degrees with varying Rician noise, as well as 4T human HARDI94 data.

2 Methods

2.1 Image acquisition

Three datasets were used in this study. The first one was simulated: we created various models of two-fiber systems, crossing at 90 degrees with equal volume fractions (w_1=w_2=0.5). Here we chose λ_1=10x10^{-4} mm^{2}s^{-1} and λ_2=2x10^{-4} mm^{2}s^{-1} as the eigenvalues for each individual component tensor (with FA=0.77, typical for white matter in the brain) and we added Rician noise of different amplitudes (with signal-to-noise ratio, SNR=5, 15, 25, and with a standard deviation of S(0)/SNR) to generate simulations using discrete mixtures of Gaussian distributions. The simulated data were sampled at 94 points evenly distributed on the hemisphere with an angular distribution computed from a partial differential equation (PDE) based on electrostatic repulsion [14]; we chose 94 as it was the same as the number of gradients used in the human 4T HARDI experiment, which was the source of the second dataset analyzed in this study.

One young healthy human subject was scanned using a diffusion-sensitized MRI protocol on a Bruker Medspec 4 Tesla MRI scanner, with a transverse electromagnetic (TEM) headcoil. The timing and angular sampling of the diffusion sequence was optimized for SNR [14, 15]. The protocol used 94 diffusion-sensitized gradient directions, and 11 baseline scans with no diffusion sensitization (b-value: 1159 s/mm^2; TE/TR: 92.3/8250 ms; FOV=230x230; in-plane resolution: 1.8mmx1.8mm; 55 x 2mm contiguous slices; acquisition time: 14.5 minutes).

Finally, a third HARDI dataset came from a monkey scanned using diffusion imaging on a 7 Tesla MRI scanner at the Center for Magnetic Resonance Research, at the University of Minnesota, using 100 gradients and 3 different b-value settings (1000, 2000, 3000 s/mm^2), TR/TE of 4600/65 ms, and an imaging matrix of 128x128x50 with isotropic voxels of 1 mm^3 (acquisition time: 23.5 minutes).

2.2 Data Processing

Several angular sampling schemes, with 20 to 94 directions, were sub-sampled from the original 94 angular locations to maximize a measure of the total angular energy. The angular distribution energy between point *i* and point *j* is denoted by E_{ij}, and defined as the inverse sum of the squares of the least spherical distance between point *i* and point *j* and the squares of the least spherical distance between point *i* and point *j*'s antipodally symmetric point *J* (Eq. 1):

$$E_{ij}^{-1} = \mathrm{dist}^2(i,j) + \mathrm{dist}^2(i,J) \quad (1)$$

Here, i,j are two different points in the spherical surface, J is the antipodally symmetric point to j, and dist(i,j) is the least spherical distance between point i and point j (see **Figure 1**). The total angular distribution energy for one gradient subset with *N* diffusion-sensitized gradients was defined as the summation of angular

distribution energy between all points in all pairs, using geodesic distances on the sphere (Eq. 2):

$$E(N) = \sum_{i=1}^{N} \sum_{j=1}^{N} E_{ij} \quad (i \neq j) \tag{2}$$

We first chose one seed point from the original 94 points, which (without loss of generality) was chosen to be (1, 0, 0) in our study. We then found another 5 points from the remaining 93 points to maximize E (6), since 6 diffusion-sensitized gradients are the minimum required for tensor estimation (so long as a non-diffusion-sensitized reference signal is also collected). In this way, the first subset with 6 diffusion-sensitized gradients was produced. After this initial subset, we artificially increased the angular sampling one gradient at a time, by maximizing E (N) (where N is the total number of diffusion sensitized gradients).

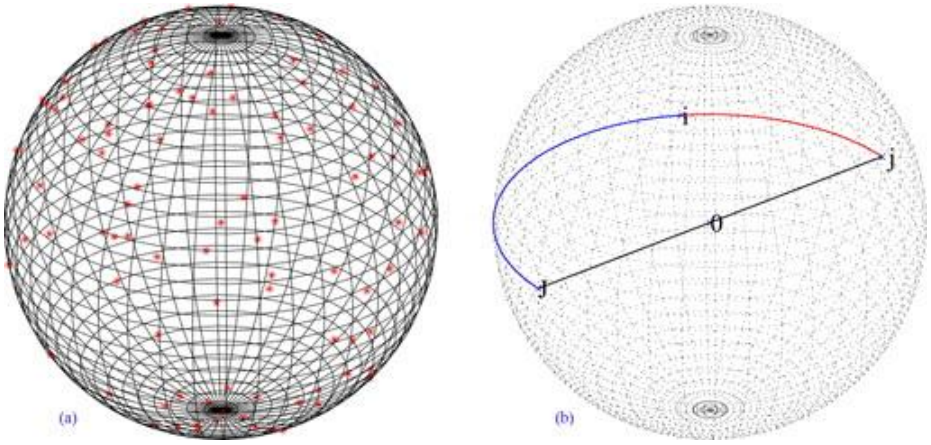

Figure 1. (a) Spherical distribution of diffusion gradient encoding angles. Red points on the sphere indicate the spherical distribution of angles at which diffusion-sensitized gradient images were collected for the 105-gradient HARDI sequence, which consists of 94 diffusion-sensitized gradients and 11 non-sensitized gradients. Each red dot in this figure represents one gradient direction, so there are 94 points in total on the unit sphere. In areas that appear to be relatively sparsely sampled, there is typically a sampled point on the opposite side of the sphere. Also, equidistribution problems sometimes lead to apparent clusters of points in some regions (see e.g., Friedman E. "Circles in Circles." http://www.stetson.edu/~efriedma/cirincir/), as the minimum point separation is only the same for all points for certain specific sample sizes. (b) Angular distribution energy calculation. In this figure, O is the original point, i, j represent two different points on the spherical surface, J is the antipodally symmetrical point to j. Based on Equation 1, the angular distribution energy between i and j is contributed based on the least spherical distance between point i and point j - denoted by dist(i,j) - and the least spherical distance between point i and point J - denoted by dist(i, J). dist(i,j) is illustrated by the red curve while dist(i,J) is represented by blue curve on the sphere on the right.

Using these optimized subsets of angular points, we sub-sampled the original HARDI94 data, and applied the framework in [13] to all these sub-samples. We denote the space of symmetric positive definite 3x3 matrices by $\mathbb{D}$. The probabilistic

ensemble of tensors, as represented by a tensor distribution function (TDF) P, is defined on the tensor space $\mathbb{D}$ that best explains the observed diffusion-weighted images (Eq. 3):

$$S_{calculated}\ (q) = \int_{D\in\bar{\mathbb{D}}} P(D)\exp(-tq^t Dq)dD \tag{3}$$

To solve for an optimal TDF P^*, we use the multiple diffusion-sensitized gradient directions q_i and arrive at P^* using the least-squares principle (Eq. 4):

$$P^* = argmin_p \sum_i (S_{obs}(q_i) - S_{calculated}\ (q_i))^2 \tag{4}$$

From the TDF, the orientation density function (ODF) may be analytically computed from Eq. 5. These ODFs were rendered using 642 points, determined using a seventh-order icosahedral approximation of the unit sphere.

$$ODF(\tilde{x}) = C\int_{r=0}^{\infty} p(r\tilde{x})dr = C\int_{D\in\bar{D}} P(D)(\det(D)\,\tilde{x}^t D^{-1}\tilde{x})^{-\frac{1}{2}}dD \tag{5}$$

To assess how accurately the diffusion profiles could be reconstructed from subsampled data based on different angular sampling schemes, the Kullback-Leibler (sKL) divergence, a commonly used measure from information theory, was used to measure the discrepancy between the reconstructed and ground truth ODFs. Reconstruction error was calculated from Eq. 6, in which *p(x)* is the ODF derived from the subsampled schemes with additive Rician noise of various amplitudes, while *q(x)* is the noise-free ODF derived from the ground truth data.

$$sKL(p,q) = \frac{1}{2}\int_{\Omega} \left\{p(x)\log[\frac{p(x)}{q(x)}] + q(x)\log[\frac{q(x)}{p(x)}]\right\}dx \tag{6}$$

We also computed another measure of fiber integrity proposed in the original TDF framework, the exponential isotropy (EI; Eq.7). Given any TDF P, EI quantifies the overall isotropy of diffusion at any given voxel, and highlights the gray matter instead of white matter as in FA (since gray matter voxels tend to have low anisotropy, or high isotropy, and thus high EI values). EI is defined as the exponential function of the Shannon Entropy, so EI can also be used to quantify the information uncertainty:

$$EI\big(P(D)\big) = e^{Shannon\ \ Entropy} = e^{-\int_{D\in\bar{D}} P(D)logP(D)dD} \tag{7}$$

3 Results and Discussion

3.1 How reconstruction accuracy was affected by angular resolution in the presence of variable additive Rician Noise.

Figure 2 shows several characteristic ways in which the additive Rician noise affected the reconstructed ODFs. The effects of image noise on the reconstructed

HARDI ODF included combinations of (1) local diffusion coefficient swelling, (2) incorrect rotations of the dominant fiber directions, and (3) mixing or omission of maximum diffusivity peaks in the radial fiber profile.

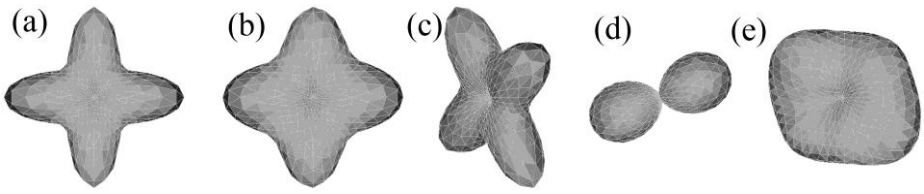

Figure 2. Noise effects on the HARDI ODF.
These glyphs show characteristic types of reconstruction errors that resulted from adding Rician noise to a simulated 2-fiber system, followed by deriving an ODF from the fitted tensor distribution function. (a) Ground truth ODF; (b) swelling of the local diffusion coefficient; (c) incorrect rotations of the dominant fiber directions (this is a rotation out of the plane of the page); (d) total omission of a dominant fiber direction; (e) mixing of the dominant directions. All these ODF are calculated based on Eq. 5 in the TDF framework without any regularization. Overall, the effect of noise on the HARDI ODF will most likely induce combinations of each of these types of distortion.

Next we assessed how the angular resolution affects the HARDI ODF reconstruction. **Figure 3** shows that, as expected, the higher the angular resolution, the more accurately the ODF can be recovered; even so, reconstruction errors vary from angular smearing and coalescing of the ODF peaks between 30 and 60 gradients to incorrect recovery of the dominant fiber direction at 20 gradients, which could be problematic for ODF-based tractography.

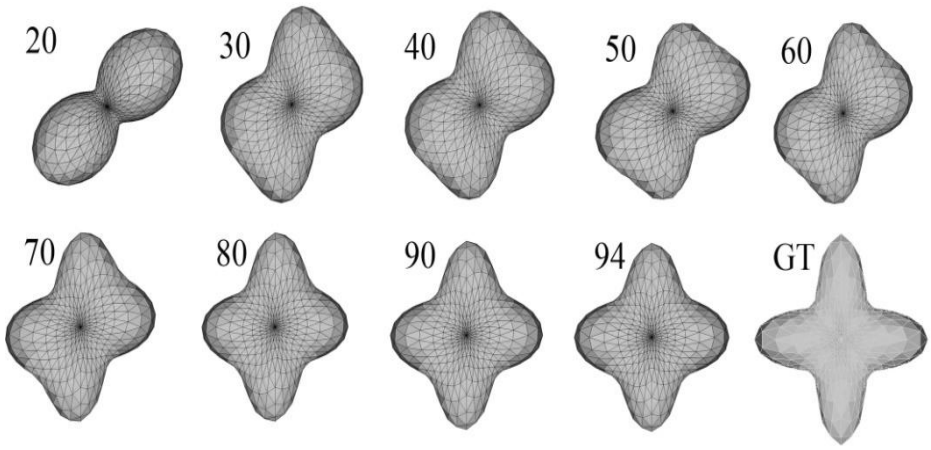

Figure 3. Angular Resolution effects on the HARDI ODF
This figure illustrates how angular resolution affects the HARDI ODF, which was calculated in the TDF framework based on Eq. 6, without any regularization. ODFs are reconstructed from sets of progressively more gradients, in directions that optimize the angular distribution energy (Eq. 3): the number on the upper left of each panel is the number of diffusion-sensitized gradients used to reconstruct the ODF. GT denotes ground truth.

In **Figure 3,** it is not immediately clear why the smaller number of gradient directions always coalesces the two peaks in the same direction (bottom-left to top-right); this most likely occurs because we use induction to define the gradient sets, so there cannot be perfect symmetry in the gradient set for all *n*, and some subsets will have a net excess of gradients in one quadrant (i.e., the point set will have a well-defined principal axis), which may lead the 2 dominant ODF peaks to coalescence into one in a specific quadrant, as the angular detail is reduced.

To quantify the accuracy of ODF recovery at different SNR levels and at different angular resolutions, we calculated the reconstruction error, represented by the sKL divergence between the recovered and the ground truth signal. As expected, the sKL error decreased with increasing SNR, and when more scanning directions were used (**Figure 4(a)**). The reconstruction accuracy of a 90-direction low-SNR sequence was about the same as a 30-direction sequence with five times the SNR. Our simulation studies showed that when SNR is low, adding directions has greater benefit. Moreover, higher angular resolution is needed for low SNR sequences to achieve reconstruction accuracy comparable to those obtained with higher SNR.

3.2 How the information uncertainty index was affected by the angular resolution

Information uncertainty was quantified here by EI which is a measure of fiber integrity related to FA (but avoiding the limitations of the single-tensor model).

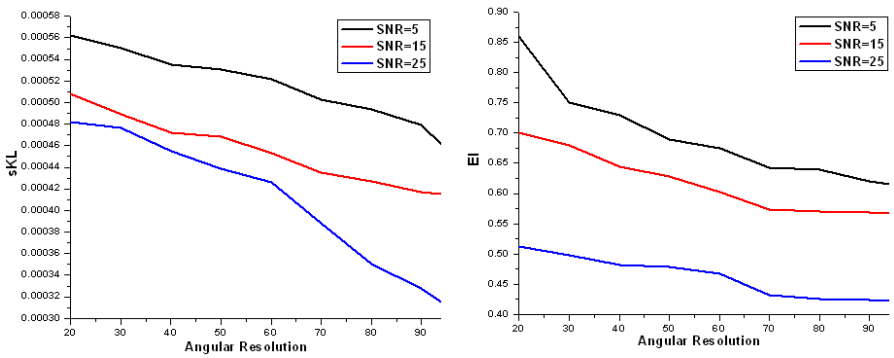

(a) sKL vs. Angular resolution (b) EI vs. Angular resolution

Figure 4 Simulation Studies with different SNR level

For each angular resolution scheme, 1000 simulations of two-tensor systems (equal volume fractions; 90º crossing) were computed with different SNR. (a) sKL divergence (reconstruction error) decreases with increasing SNR level, and with higher angular resolution sampling schemes. This means that the accuracy of the computed ODF improves as SNR increases and angular resolution increases. (b) EI decreases as the SNR level increases, and with more detailed angular sampling. In this figure, EI values have been normalized by the corresponding isotropic term, so that all EI values lie between 0 and 1 (which is the range for the more common anisotropy measure, FA). Also, EI tends to stabilize when the angular resolution reaches ~70 gradients. This is in line with the observation that standard FA measures are biased (too low) in regions where fibers mix or cross.

Other common measures of HARDI diffusion would also be used, such as generalized FA or total diffusion, but here we used EI as it has a direct link to the information content of the signal as defined by information theory.

As expected, EI decreased with increasing angular resolution. Simulation results show the EI stabilized by ~70 directions (**Figure 4(b)**). This is in line with the finding that fractional anisotropy, derived from DTI, is generally underestimated when fibers cross. Also, this result is consistent with **Figure 3,** which shows that HARDI70 has satisfying results when reconstructing a two-fiber system crossing at 90 degrees; these diagrams make it clearer why the isotropy falls (i.e., anisotropy rises) when the two fiber peaks no longer coalesce.

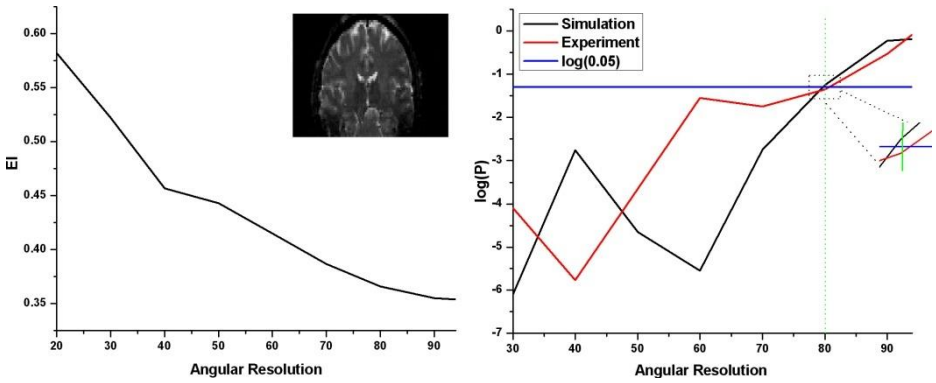

(a) EI *vs.* Angular Resolution (b) Paired *t* test result

Figure 5. 4 Tesla Human HARDI results

(a) EI *vs*. Angular Resolution in 4 T human HARDI94 data. We computed the average EI at different angular resolutions for one brain slice (the inset image is the corresponding T2-weighted slice). All EI values were normalized with respect to an isotropic diffusion profile to ensure that the EI values are between 0 and 1. We chose the average EI value in the cerebrospinal fluid (CSF) as the normalization constant since CSF has the highest diffusion isotropy in the brain. (b) Paired *t* test results. In this simulation data, the probability exceeds the threshold (p=0.05) when N is increases from 70 to 80, while for the empirical data, the probability exceeds the threshold when N is increases from 80 to 90. This answers the question, "does adding 10 more gradients improve the information in the signal?" Although these thresholds are to some extent arbitrary, they show that the information content converges within the standard range of gradients used in a HARDI study (~100).

Figure 5(a) shows how EI is affected by the angular resolution in the 4 Tesla human data. EI indicates the information uncertainty, so the smaller the EI value is, the less uncertainty there is in the multi-fiber estimation. The EI decrease with increasing angular resolution does slow down, but we did not find a plateau at around 70 gradients in **Fig. 5(a)**, as was seen in **Figure 4(b).** To better understand whether EI has converged, we performed a paired Student's *t* test on the EI values, assessing the effect of adding additional gradients, in increments of 10, (e.g. 40 *vs* 30, 30 *vs* 20) at all voxels in the brain. If the *t* test result is significant ($p<0.05$) then this test confirms that adding 10 more gradients does indeed lead to lower EI (i.e., lower information uncertainty). When this *t* test is not significant, there is no evidence that adding 10

more gradients to the acquisition protocol is helpful, so the signal may be said to be saturated. The result of this test clearly depends on the number of voxels (here 1000 for simulation and 3255 for the experiment), Even so, this is a reasonable and intuitive operational definition of saturation for practical purposes. We note that this test could be slightly improved by incorporating a multiple comparisons correction into the *p*-value, to control the false discovery rate, but we did not do so as the tests were intended as a heuristic to compare successive increments in gradient numbers.

Figure 5(b) shows paired *t* test results assessing whether the EI significantly decreases when adding more gradients (i.e., $EI_{N+10}<EI_N$) for both simulations and the 4-Tesla human data. In the simulation data, EI decreases as angular resolution increases; this progressive decrease is also statistically significant when initially adding batches of 10 additional gradients, then after 70 gradients are reached, the probability exceeds the threshold (p=0.05) and the information uncertainty no longer shows a stastistically significant improvement, consistent with Figure 4(b). Here, we may refer to this ceiling effect on EI, at 70, as the Statistical Saturation (SS) number (i.e., SS=70). We defined the meaning of this number to be that successively increasing the angular resolution always leads to statistically significant improvements in EI until the statistical saturation number of gradients is reached. This definition of incremental information gain clearly depends on the batch size (adding 10 gradients each time). For our empirical HARDI data collected at 4-Tesla, this statistical saturation number was 80. As a qualification, we note that our simulation is based on only two fibers crossing at 90 degrees with equal weighting. In the more complex case of human brain data, the voxels in each slice have varying numbers of crossing fibers, varying numbers of detectable dominant fibers, and inevitably, a different weighting for each single component fiber within each voxel. Thus, for the experimental data, more gradient directions may be needed to cause statistical saturation in the EI (our information uncertainty index).

3.3 How reconstruction accuracy was affected by multiple b values with Rician Noise.

Similarly, reconstruction accuracy was assessed with simulations with b values varying from 1000 s/mm^2 to 3000 s/mm^2 - which would be within a typical range used in diffusion spectrum imaging (DSI) studies. Rician noise was added at a SNR of 10, a level similar to real MRI images. From **Figure 6a**, we note that the sKL-divergence (reconstruction error) increases with increasing b values, but it decreases with increasing angular resolution. The explanation for this is that $S(q)/S(0)=\exp(-b\mathbf{q}^T\mathbf{D}\mathbf{q})$, so the higher b-value is, the smaller the value of S(q)/S(0) will be. This value will therefore be more greatly affected by noise, if the noise characteristics are set independently of the b-values. So increasing the b-value leads to an increasing effect of noise in the final composite data, and thus higher reconstruction error. Even so, the additional b-value shells may be used to provide additional information on the diffusion propagator that would not be obtainable using only a single b-value.

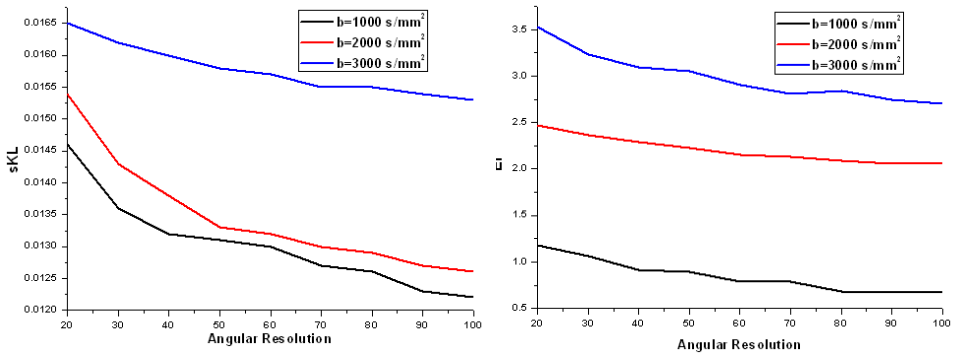

(a) sKL vs. Angular resolution , as b is varied (b) EI vs. Angular resolution, as b is varied

Figure 6. Simulation Study using different b- value settings

For each angular resolution scheme, 1000 simulations of two-tensor systems (equal volume fractions; 90º crossing) were computed with different b values (b=1000, 2000 and 3000 s/mm^2). Rician noise was added at a SNR level of 10. (a) sKL divergence (reconstruction error) increases with increasing b values, while decreasing using higher angular sampling schemes. This means that the accuracy of the computed ODF improves as b value decreases and angular resolution increases. (b) The EI behaves in the same fashion as sKL in (a).

3.4 How the information uncertainty index was affected by the b-values

To investigate the effect of b-value settings at high magnetic field (7 Tesla), we analyzed a 100-direction 7 T monkey HARDI dataset, exactly as in Section 3.2.

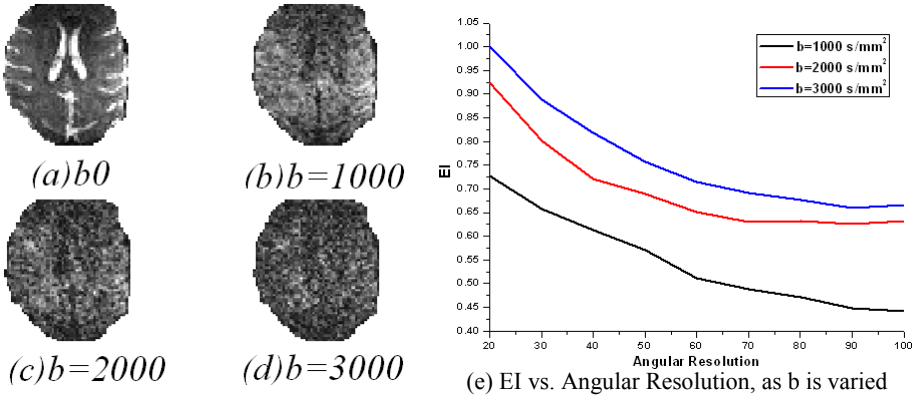

(e) EI vs. Angular Resolution, as b is varied

Figure 7. 7 Tesla HARDI scanning results

In this figure, 100-direction 7 Tesla HARDI data from a monkey was analyzed using EI to measure the information uncertainty. Illustrative slices are shown from the (a) T2 reference image (b) DWI at b=1000 s/mm^2 (c) DWI at b=2000 s/mm^2 (d) DWI at b=3000 s/mm^2, and (e) EI plot vs. Angular Resolution at the different b-values. As expected, higher b-value shells give noisier data.

Figure 7 shows a T2 image and DWI images taken at three separate b-values. The plots show how EI was affected by the angular resolution at the three different b-

value settings. EI was also affected by changing the angular resolution, at different b-values. Visualizations in **Figures 7(b)-(d**) show that the diffusion weighted images at higher b-values are quite noisy, which may be due to the combination of the higher magnetic field and the higher b value (suppressing the diffusion-weighted signal relative to the noise). Here, the EI versus angular resolution plot, **Figure 7(e)**, exhibits the same pattern as that seen in simulations (**Figure 6(b)**), suggesting that b-values higher than 2000 may be suboptimal when acquiring ultra-high magnetic field strength DWI images.

4 Conclusion

HARDI scanning allows better diffusion reconstruction than DTI, and provides new insight into fiber architecture and connectivity that cannot be achieved, even in principle, using a smaller number of diffusion-sensitized gradients. These advantages come at the expense of longer scanning times, but the trade-off may be worth it for studies assessing fiber connectivity and for fine-scale mapping of anatomy, and to avoid errors in routine clinical studies. We identified several types of ODF reconstruction errors that are typical when smaller numbers of gradients are used, and studied their asymptotics in optimized angular sets. To improve diffusion reconstruction accuracy and remove bias from the derived anisotropy measures, it is more effective to acquire additional angular samples than to repeatedly sample the same directions for purposes of signal averaging [16-18]. We found that, with a reasonable intuitive definition of saturation, the information uncertainty cannot be statistically improved when the number of diffusion-sensitized gradients exceeds 80. Also, from our preliminary study at 7 Tesla, the b-values should not be set too high, in order to obtain satisfactory EI values. Thus, our study may be of interest in designing future DTI and HARDI acquisition protocols for assessing fiber integrity in the living brain.

References

1. Basser PJ, Pierpaoli C. Microstructural and physiological features of tissues elucidated by quantitative diffusion tensor MRI. **J. Magn. Reson.,** vol. B 111, no. 3, pp.209-219 (1996).
2. Stejskal, EO, Tanner JE. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. **J. Chem. Phys.** 42:288-292 (1965).
3. Tuch DS. Q-Ball Imaging. **Mag. Res. Med.** 52:1358-1372 (2004).
4. Tournier JD, Calamante F, Gadian DG, Connelly A. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. **NeuroImage**: 23: 1176-1185 (2004).
5. Kaden E, Knosche TR, Anwander A. Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging. **Neuroimage**: 37: 474-486 (2007).
6. Jian B, Vemuri BC, Özarslan V, Carney PR, Mareci TH. A novel tensor distribution model for diffusion-weighted MR signal. **NeuroImage** 37: 164-176. (2007)

7. McGraw T, Vemuri BC, Yezierski B, Mareci T. Von Mises-Fisher mixture model of the diffusion ODF. Biomedical Imaging: Nano to Macro, **3rd IEEE International Symposium on Biomedical Imaging**: 65-68 (2006).
8. Ozarslan E, Vemuri BC, Mareci TH. Higher rank tensors in diffusion MRI, **Visualization and Image Processing of Tensor Fields**, (2005).
9. Barmpoutis A, Vemuri BC. Exponential Tensors: A framework for efficient higher-order DT-MRI computations. In Proceedings of ISBI07: **IEEE International Symposium on Biomedical Imaging**: Page(s): 792-795 (2007).
10. Perrin M, Poupon C, Rieul B. Validation of q-ball imaging with a diffusion fibre-crossing phantom on a clinical scanner. **Phil Trans R Soc Lond B Biol Sci** 360:881–891 (2005).
11. Jbabdi S, Woolrich MW, Andersson JLR et al. A Bayesian framework for global tractography. **NeuroImage** 37:116-129. (2007).
12. Kaden E, Knösche TR, Anwander A. Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging. **Neuroimage**. 37: 474-488 (2007).
13. Leow AD, Zhu S, Zhan L, McMahon KL, de Zubicaray GI, Meredith M, Wright MJ, Toga AW, Thompson PM (2009). The Tensor Distribution Function, **Magnetic Resonance in Medicine**, 2009 Jan. 18; 61(1):205-214 (2009).
14. Jones DK, Horsfield MA, Simmons A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. **Magnetic Resonance Imaging** 42(3): 515-525 (1999).
15. Pend H, Arfanakis K. Diffusion tensor encoding schemes optimized for white matter fibers with selected orientations. **Magnetic Resonance Imaging** 25: 147-153 (2007).
16. Goodlett C, Fletcher PT, Lin W, Gerig G. Quantification of measurement error in DTI: theoretical predictions and validation. **MICCAI** 2007; 10(Pt 1): 10-17 (2007).
17. Jones DK, Basser PJ. Squashing peanuts and smashing pumpkins: How noise distorts diffusion-weighted MR data. **Magnetic Resonance in Medicine** 52(5), 979–993 (2004).
18. Jones DK. The Effect of Gradient Sampling Schemes on Measures Derived From Diffusion Tensor MRI: A Monte Carlo Study. **Magnetic Resonance in Medicine** 51(4), 807–815 (2004).