

Visualization Tools for High Angular Resolution Diffusion Imaging*

David W. Shattuck¹, Ming-Chang Chiang¹, Marina Barysheva¹,
Katie L. McMahon², Greig I. de Zubicaray², Matthew Meredith²,
Margaret J. Wright³, Arthur W. Toga¹, and Paul M. Thompson¹

¹ Laboratory of Neuro Imaging, Dept. of Neurology, UCLA, Los Angeles, CA, USA

² Centre for Magnetic Resonance, University of Queensland, Brisbane, Australia

³ Queensland Institute of Medical Research, Brisbane, Australia

Abstract. There is a major effort in medical imaging to develop algorithms to extract information from DTI and HARDI, which provide detailed information on brain integrity and connectivity. As the images have recently advanced to provide extraordinarily high angular resolution and spatial detail, including an entire manifold of information at each point in the 3D images, there has been no readily available means to view the results. This impedes developments in HARDI research, which need some method to check the plausibility and validity of image processing operations on HARDI data or to appreciate data features or invariants that might serve as a basis for new directions in image segmentation, registration, and statistics. We present a set of tools to provide interactive display of HARDI data, including both a local rendering application and an off-screen renderer that works with a web-based viewer. Visualizations are presented after registration and averaging of HARDI data from 90 human subjects, revealing important details for which there would be no direct way to appreciate using conventional display of scalar images.

1 Introduction

Diffusion magnetic resonance imaging [1] applies gradient fields during image acquisition, allowing local measures of water diffusion preferences to be made. In diffusion tensor imaging (DTI), a reference image and at least six diffusion images are acquired with non-collinear gradients and used to reconstruct a volume of rank-2 tensor models [2]. The shapes of these tensors provide information about the underlying structure within the biological sample. For brain imaging, DTI provides valuable insight into the white matter architecture, with a distinct advantage over traditional MRI structural imaging where this architecture is not well-resolved.

While the 2nd-order tensor model is capable of describing regions of the brain that have consistent structure, the model breaks down in voxels that contain multiple fibers. To address this problem, Tuch et al. proposed High Angular

* Funding: Grants P41-RR013642, RO1-HD050735, and NHMRC496682 (Australia).

Resolution Diffusion Imaging (HARDI) [3]. HARDI uses larger numbers of directions, often 100 or more, to acquire measurements of the diffusion. The sampled data can be used to estimate the orientation distribution function (ODF), which gives the probability of directional water diffusion within each voxel. One representation of the ODF is in the form of a spherical harmonic basis [4], allowing the shape at each voxel to be compactly represented by a truncated set of spherical harmonic coefficients (SHCs). While HARDI presents several advantages over DTI, it significantly raises the complexity of data processing and visualization.

A number of other packages support display and processing of DTI data, e.g., SCIRun (<http://www.sci.utah.edu>), TrackVis (<http://www.trackvis.org>), DTIStudio (<http://www.dtistudio.org>), Slicer (<http://www.na-mic.org/Wiki/index.php/Slicer>), and MedINRIA (<http://www-sop.inria.fr/asclepios/software/MedINRIA>). Camino (<http://www.cs.ucl.ac.uk/research/medic/camino>) can fit 2 and 3-tensor models to HARDI data, and it exports models for viewing in other programs; however it has limited capabilities for display. None of these tools specifically support visualization of spherical-harmonic-based ODF models. While there have been various published works [5,6] that make use of these types of data, the customized tools used to explore these data do not appear to have reached the neuroimaging community yet.

There is a clear need for visual methods to evaluate the results of image processing operations on HARDI data or to investigate features in the data. One important application for visualization of ODFs is in the interpretation of HARDI registration results. Techniques have been developed for the registration of DTI [7,8,9] that account for the inherently non-Euclidean nature of DTI, which requires that the tensors be reoriented during spatial normalization and averaging. HARDI further complicates this problem, as it allows for more detailed shapes that may represent, for example, the crossing of two nerve tracts. The rotational approach, which can be determined if we compute a registration using a tensor model, may be insufficient to account for the distortions required to match the ODF across subjects.

In this paper, we present a set of tools for producing interactive visualization of HARDI ODF models based on the modified spherical harmonic basis presented by Descoteaux et al. [4]. We demonstrate their capabilities using individual and averaged HARDI data.

2 Methods

HARDI data from 90 subjects were acquired on a 4T Bruker Medspec MRI scanner using an optimized diffusion tensor sequence [10]. These data included both 30-direction (HARDI30) and 105-direction (HARDI105) diffusion imaging. For the HARDI30 data, imaging parameters were: 21 axial slices (5 mm thick), FOV = 23 cm, TR/TE 6090/91.7 ms, 0.5 mm gap, with a 128x100 acquisition matrix. The reconstruction matrix was 128x128, yielding a 1.8x1.8 mm² in-plane resolution. The total scan time was 3.05 minutes. For the HARDI105 data, 105 images were acquired: 11 baseline (b0) images with no diffusion sensitization

and 94 diffusion-weighted images (b-value 1159 s/mm²) in which gradient directions were evenly distributed on the hemisphere [10]. Imaging parameters were: TR/TE 92.3/8250 ms, 55 x 2mm contiguous slices, FOV = 23 cm. The reconstruction matrix was 128x128, yielding a 1.8x1.8 mm² in-plane resolution. The total scan time was 14.5 minutes. For each subject, rank-2 diffusion tensor images were computed from the HARDI signals using MedINRIA software. The first component of the diffusion tensor (D_{xx}) was manually stripped of non-brain tissues to produce a brain mask for the subsequent processing and display.

Orientation distribution functions (ODF) for water diffusivity were estimated voxel-wise from the HARDI signals using the Funk-Radon Transform (FRT). The HARDI measurements were expanded as a spherical harmonic (SH) series, which simplified the FRT to a linear matrix operation [5]. The estimated ODF was normalized to unit mass, creating a diffusion probability density function (PDF) parameterized by spherical angle. The representation for the ODF followed the model described in [4], using the modified spherical harmonic basis

$$Y_j = \begin{cases} \frac{1}{\sqrt{2}}((-1)^m(Y_l^m) + Y_l^{-m}) & -l < m < 0 \\ Y_l^0 & m = 0 \\ \frac{i}{\sqrt{2}}((-1)^{m+1}(Y_l^m) + Y_l^{-m}) & 0 < m \leq l \end{cases}, \quad (1)$$

where $\theta \in [0, \pi]$, $\phi \in [0, 2\pi]$, $j(l, m) = (l^2 + l + 2)/2 + m$. Y_l^m is the standard spherical harmonic basis function

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\phi} \quad (2)$$

where P_l^m is the associated Legendre polynomial. The modified SH basis takes advantage of the symmetries in the spherical harmonics to produce a real-valued set of basis functions, simplifying computation for data pre-processing and rendering. The diffusion ODF at each voxel is then represented by

$$ODF[k](\theta, \phi) = \sum_{j=1}^N ODFC_j[k] Y_j(\theta, \phi) \quad (3)$$

where $ODFC_j[k]$ is the j -th coefficient of the spherical harmonic transform for the k -th voxel. $N = (l+1)(l+2)/2$, and can be truncated at different levels. Higher orders of coefficients correspond to higher frequencies on the sphere, thus truncating at higher orders allows the shape of the ODF to be more complicated, while truncating at lower levels will produce much smoother shapes.

We aligned and averaged the HARDI30 datasets for 90 subjects using a nonlinear fluid registration approach. An initial affine transformation was computed to produce a suitable initial alignment for the fluid method. The masked D_{xx} image was registered to the ICBM53 average brain template using an affine 9-parameter transformation produced by the software FLIRT [11] and resampled to isotropic voxel resolution (dimension: 128x128x93 voxels, resolution: 1.7x1.7x1.7 mm³). The resulting transformation parameters were used to rotationally reorient the

tensor at each voxel and then affine align the tensor-valued images based on trilinear interpolation of the log-transformed tensors. All affine-registered DT images were then registered to a randomly selected subject's image, using an inverse-consistent fluid registration algorithm that minimized the symmetrized Kullback-Leibler divergence of the two tensor-valued images [12]. This minimization allows us to achieve better correspondence in the white matter structures than would be achieved with scalar image matching. We note that DT data are not available from the ICBM53 template, hence we are not able to apply tensor-matching directly to the template. The images of the diffusion ODFs were resampled to the target space by applying the corresponding DTI mapping. To keep the direction of the diffusion ODFs oriented with the direction of the underlying fibers, the ODFs were reoriented using the Preservation of Principal Direction (PPD) method [7], where the principal direction of the ODF was determined based on principal component analysis of the ODF [12].

We developed a Windows-based interactive software package for 3D visualization of the ODF shapes using OpenGL and C++. Each shape is represented as a glyph based on a sampling of the values of $ODF[k](\theta, \phi)$. The software can load a series of SHC volumes and display orthogonal views of slices through the coefficient data, as well a 3D view of the reconstructed glyph shapes. Its graphical user interface (GUI) provides controls to adjust the number of coefficients (N), the number of samples for θ and ϕ , as well as a window size for how many voxel ODFs are displayed at once. When the parameters are adjusted, a new series of glyphs are computed and the display is updated. The software can also display additional volume files in register with the HARDI data, as well as surface models of other anatomical structures, such as the cortex or streamline models of fiber tracts, and tensor glyphs for analyzing DTI data. Because each voxel is represented by a detailed shape model, we note that the rendering time for high levels of detail can be prohibitive for interaction. Thus, the software allows the user to navigate using a lower-fidelity set of glyphs and produce the detailed rendering when an appropriate view has been selected.

The software provides user-selectable color models for the ODF glyphs, including the R_2 measure proposed in [13],

$$c = \frac{\sum_{j:l=2} |c_j|}{\sum_{\forall j} |c_j|} \quad (4)$$

where c_j is the j -th SHC. This represents the fraction of energy in the $l = 2$ components and is an indicator of structural content. We implemented other measures described in [4] and a color model that shades each glyph according to the direction of the maximum in the ODF. The latter uses an encoding similar to DTI, where red corresponds to left-right directions, green corresponds to anterior-posterior directions, and blue corresponds to superior-inferior directions.

We also developed an off-screen renderer that can generate slice views of the ODFs. This renderer provides a command-line interface with all the parameters of the GUI, but it can produce images with dimensions as high as 16K by 16K pixels. Images can be produced in JPEG format, as well as the Zoomify format

(Zoomify, Inc., Santa Cruz, CA), which produces a multi-resolution pyramid decomposition of the image into tiles of smaller images. We applied this rendering package using a clustered-computing environment, allowing us to send each slice in the HARDI volume to a different compute node. The resulting images were stored directly into a web-accessible directory, with an automatically constructed web-interface that allowed navigation through each slice of data, zooming to any level of detail, and switching among axial, coronal, and sagittal views of the HARDI data. This allows for convenient sharing of research results online, obviating the need for special hardware and allowing for rapid display of results without downloading the entire set of volumetric data.

3 Results

We applied the methods described in the previous section. The results were displayed using a notebook computer with an Intel Centrino Duo 2.16Ghz processor and an NVidia 7950GTX graphics card. We observed that frame rates suitable for interaction (approximately 10 frames / second) were achieved for whole-slice masked data with 15×15 samples in (θ, ϕ) , which were represented by approximately 2 million triangles per slice.

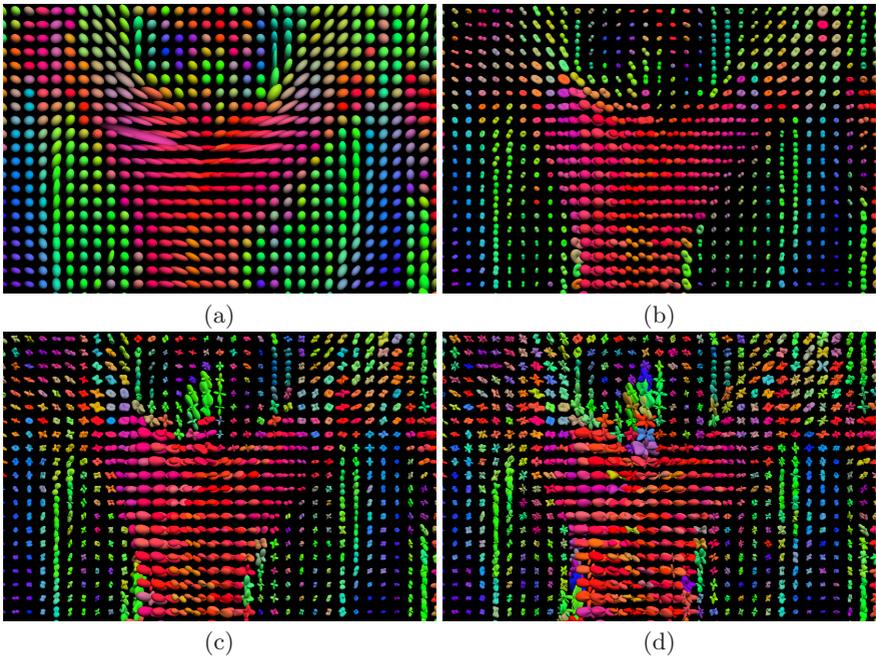


Fig. 1. Glyphs computed from a HARDI105 data set. (a) DTI ellipsoid glyphs (b) HARDI ODF (6 SHCs) (c) HARDI ODF (15 SHCs) (d) HARDI ODF (28 SHCs).

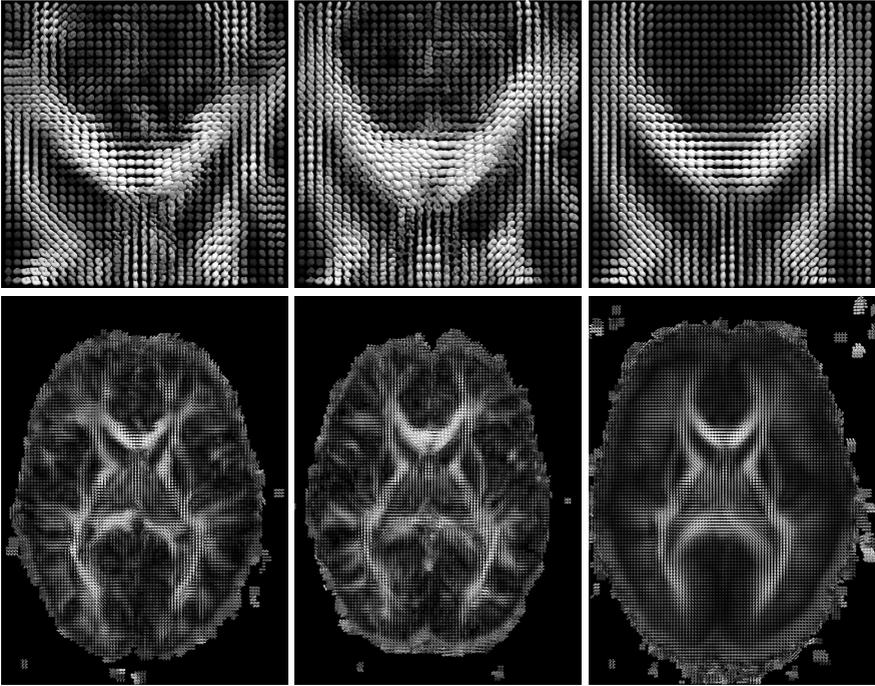


Fig. 2. Visualization of orientation distribution functions. (left) Data from a single subject. (middle) Data from a second subject, resampled after alignment to match the first subject. (right) An average map of orientation distribution functions for 90 subjects. Each image is colored using the R_2 measure. The top row is a set of zoomed views of the images beneath them. The white matter structure is readily apparent in these images. We note that the group average image demonstrates visually that the averaged ODF for each voxel retains significant shape characteristics.

Figure 1 shows renderings of a region near the corpus callosum in a slice of DTI and ODF data from a single subject. The first view shows the DTI glyph model, normalized to unit mass to be consistent with the ODFs. The other views show the SH representation of the ODF, with $N = 6$, $N = 15$, and $N = 28$. These correspond to increasing values of the l -parameter, i.e., $l \in \{2, 4, 6\}$. The sampling in (θ, ϕ) was 50×50 . The color encoding scheme for each frame indicates the maximum direction for the tensor or ODF. We observe that, as expected, the ODF representation provides significantly more detail than the DTI version. This is seen, e.g., in the upper left corner of the images where the corticospinal tracts (blue) and the corpus callosum (red) meet. The tensors in this region are fairly isotropic, making their direction ambiguous. In the ODF models, the crossings become more apparent as there are multiple local maxima. As we move to higher numbers of coefficients ($N = 28$), we see what appears to be an increase in the noise in the image. Selection of an appropriate level at which to truncate

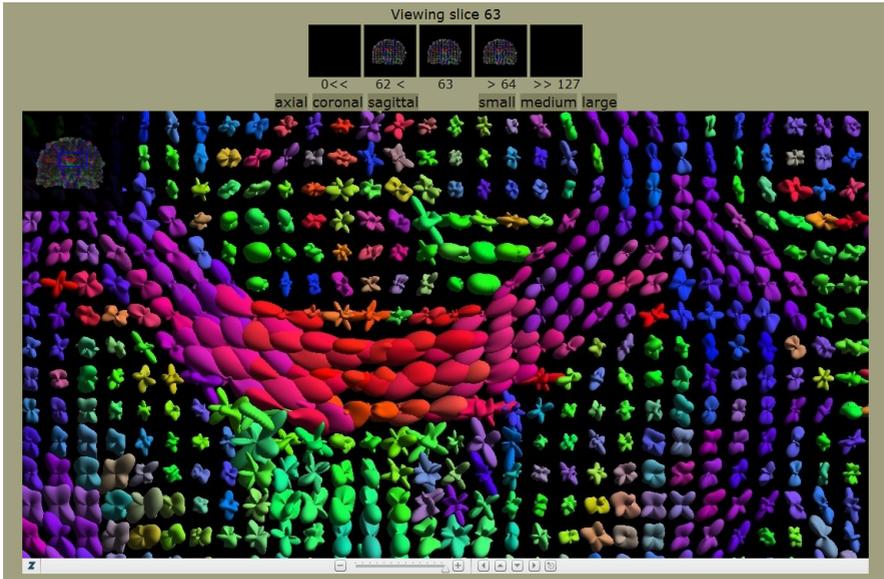


Fig. 3. The zoomify-enabled web-view of a HARDI dataset, produced automatically by our software. The user can navigate to different slices and orientations, and zoom to high levels of detail. The image being viewed is 8192×8192 pixels.

the SH series remains an open area of research; visualizations such as these may assist in better understanding the impact of the coefficients.

Figure 2 shows the results of the registration of the HARDI30 data for 90 human subjects, as described in the previous section. These visualizations indicate the transformations that are induced upon the ODF model and demonstrate the process that occurs as the spherical harmonic representations of the ODF are averaged. Orientations of the second subject (middle panel) can be seen to closely match those of the first subject (left panel) in many regions. The 90-subject average (right panel) clearly preserves many of the features of the individual brains, though much of the higher order detail has been smoothed by the averaging process. We note that the shapes of the ODFs in the average map are similar to those in Fig. 1.b., where $N = 6$ SHCs were used.

We also applied the off-screen renderer to the HARDI105 data. Figure 3 shows the resulting web interface, which provides access to 2GB of pre-rendered image data. The images were computed with a 75×75 sampling of (θ, ϕ) space with an image resolution of 8192×8192 pixels. The total wall-clock time to compute this collection was less than 10 minutes using a 111-node 2.4GHz Opteron cluster. We note that the cluster is a shared resource, and its load may vary greatly. Renderings of individual slices at this resolution required less than one minute of compute time. We were able to use the web-based interface to share results directly among collaborators, including ones on different continents. In

the rendering shown in Fig. 3, we observe features that would not normally be detectable with DTI. For example, on the upper right that corpus callosum fibers become interspersed with other white matter tracts orthogonal to them. We also observe spatially varying contrast inside the striatum.

4 Discussion

We have introduced a new set of tools for the visualization of ODF models computed from HARDI. These tools allow the display of intricate details, and we applied them to produce novel views of averaged multi-subject HARDI data. The visualizations reveal that the continuum-mechanical HARDI registration method applied in this paper, based on fluid convection of HARDI functions, does in fact preserve, and even enhance, key features in population studies of HARDI. We believe these tools will be of practical use to the neuroimaging community, and we will be distributing them online. We will be applying these visualization tools to understanding aspects of HARDI imaging to examine disease characteristics. We will use these visualization tools to understand features of the ODF data acquired at 3T and 7T in patients with Alzheimer's disease. We will also explore improving the rendering capabilities through the use of GPU programming.

References

1. LeBihan, et al.: MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology* 161(2), 401–407 (1986)
2. Basser, et al.: Estimation of the effective self-diffusion tensor from the NMR spin echo. *J. Magn. Reson. B* 103(3), 247–254 (1994)
3. Tuch, et al.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn. Reson. Med.* 48(4), 577–582 (2002)
4. Descoteaux, et al.: Apparent diffusion coefficients from high angular resolution diffusion images: Estimation and applications. Research Report 5681, INRIA (2005)
5. Descoteaux, et al.: A linear and regularized odf estimation algorithm to recover multiple fibers in q-ball imaging. Research Report 5768, INRIA (2005)
6. McGraw, et al.: Von Mises-Fisher mixture model of the diffusion ODF. In: ISBI 2006, April 6, 2006, pp. 65–68 (2006)
7. Alexander, et al.: Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Trans. Med. Imaging* 20(11), 1131–1139 (2001)
8. Xu, et al.: Spatial normalization of diffusion tensor fields. *Magn. Reson. Med.* 50(1), 175–182 (2003)
9. Zhang, et al.: Deformable registration of diffusion tensor MR images with explicit orientation optimization. *Med. Image Anal.* 10(5), 764–785 (2006)
10. Jones, et al.: Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn. Reson. Med.* 42(3), 515–525 (1999)
11. Smith, et al.: Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* 23 (Suppl. 1), S208–S219 (2004)
12. Chiang, M.C., et al.: Fluid registration of diffusion tensor images using information theory. *IEEE Trans. Med. Imaging* 27(4), 442–456 (2008)
13. Chen, et al.: Estimation, smoothing, and characterization of apparent diffusion coefficient profiles from high angular resolution DWI. In: CVPR, pp. 588–593 (2004)