Registration of Spherical Functions from High Angular Resolution Diffusion Imaging using the Heat Kernel Signature and Möbius Transformation

L. Zhan[1], Y. Wang[2], P.M. Thompson[1]

1. Laboratory of Neuro Imaging, Department of Neurology, UCLA School of Medicine, Los Angeles, CA, USA
2. Computer Science and Engineering, School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ, USA

Abstract. High angular resolution diffusion imaging (HARDI) is a powerful variant of diffusion MRI, which images the 3D profile of diffusion at each imaged location in the brain. At each voxel, this leads to an orientation density function (ODF) expressing the probability density of water diffusion in each direction on the unit sphere. As diffusion is greatest along the brain's axons, these functions are used to map fiber trajectories (tractography) and fiber integrity. To average and compare this data across subjects, we developed a new method based on the heat kernel signature and 'Möbius voting' to identify and align peaks of diffusivity on the sphere. We compare our method to standard coordinate-based averaging, and it helps to reinforce consistent features in both synthetic data and real HARDI data. This scheme to compare and integrate HARDI data may be helpful in population-based analyses of brain integrity and connectivity.

Keywords: HARDI, ODF, Heat Kernel Signature, Möbius Transformation

1 Introduction

High angular resolution diffusion imaging (HARDI) offers several advantages for studying the local integrity, geometry and connectivity of white matter fibers in the living brain. When fibers cross or mix in the same voxel, measures based on the widely-used diffusion tensor can be biased (e.g., fractional anisotropy estimates tend to be too low), in the ~40% of white matter voxels where more than one dominant fiber direction is detectable [1, 2]. From HARDI, orientation distribution functions (ODF) may be defined as the radial projection of the spherical diffusion function. Despite the rich information in HARDI datasets – yielding an entire spherical function at each point in a 3D image – the statistical study of diffusion images have mainly focused on scalars derived from the tensor model. The vast majority of studies still examine fractional anisotropy (FA) as the main target of study. [3]

Precise voxel-level comparison of HARDI datasets cannot be performed across subjects without removing some of the confounding structural variability. The diffusion functions can be considered as defined on the anatomy of each individual, and if

nonlinear registration of the anatomy is used for normalization, there will be better agreement among the diffusion functions defined on the anatomy. Better structural alignment generally improves the power to detect changes related to disease, development, and aging, and improves the accuracy of segmentation and post-processing of HARDI data [4].

Thus, there has been a recent flurry of HARDI registration algorithms [5-10]. Some work has attempted to register fields of spherical functions, treating the spherical function as a probability density, and defining distances between them based on information theory, such as the Fisher-Rao metric, or symmetrized Kullback-Leibler distance [11]. Such metrics may be integrated over the whole brain, and their variational derivative may be computed with respect to tunable parameters of 3D vector fields used to register the images [12]. Chiang et al. [12] found that 3D fluid registration of diffusion images gave more accurate alignment when the tensors or ODFs were aligned using information theory, but the approach was limited because there was no attempt to define corresponding features (such as peaks of diffusivity) on the ODFs across subjects. Because of this, corresponding features were not reinforced in the ODFs across subjects, when data were averaged and compared.

In the standard, single-tensor, DTI model, the fitted diffusion tensor has only one dominant direction (its principal eigenvector), so registration is straightforward as it only involves aligning one direction to the other (although clearly there may be ambiguous or difficult cases if the principal eigenvector is not unique or if the 2 or 3 highest eigenvalues are close in magnitude). By contrast, the ODF model can often have multiple dominant directions. This leads to difficulties in peak matching, unless an explicit effort is made to match them. In this study, we propose a new ODF registration method, designed to reinforce common features in populations of diffusion images. As the ODF is defined on the unit sphere, we first use the heat kernel signature (HKS) to detect the peaks in the spherically-parameterized functional domain. We then use stereographic projection to map the 3D spherical surface to the 2D complex plane, and match the HKS peaks in complex plane using Möbius transformation. We then use the inverse stereographic projection to pull-back the mapping to the 3D space. Using this method, ODF multiple peaks may be matched very quickly, giving robust results in line with intuition.

2 Method description

2.1 ODF Computation using the Tensor Distribution Function

We adopted the Tensor Distribution Function (TDF) theory [13, 14] to calculate the ODF. The space of symmetric positive definite 3x3 matrices was denoted by $\bar{D}$. The probabilistic ensemble of tensors, represented by a tensor distribution function (TDF) P, is defined on the tensor space $\bar{D}$ that best explains the observed diffusion-weighted signals. We used the least-squares approach with the gradient descent defined in [13] to solve for an optimal TDF P^*. From the TDF P^*, the ODF may be computed analytically [15]. All ODFs were rendered using 642 point samples, determined using an icosahedral approximation of the unit sphere. The TDF-based ODF

calculation can be expressed in **Equation 1** (Please refer to [13] for explanation in details). In addition, 1280 faces were determined using Delaunay triangulation. Every local maximum of the 3D surface may be one of the possible dominant directions. To find the local maxima, we adapted Heat Kernel Signature theory.

$$\begin{cases} S_{cal}(q) = \int_{D \in \bar{D}} P(D) \cdot e^{(-bq^T Dq)} dD \\ P^* = argmin_p \sum_i \left(S_{obs}(q_i) - S_{cal}(q_i) \right)^2 \\ ODF(\tilde{x}) = \frac{1}{4\pi} \sum_{D \in \bar{D}} \frac{P^*(D)}{det(D)^{\frac{1}{2}}(\tilde{x}^T \cdot D^{-1} \cdot \tilde{x})^{\frac{3}{2}}} \end{cases} \qquad (1)$$

2.2   Peak detection using the Heat Kernel Signature

The Heat Kernel Signature (HKS) was first proposed by [16]. Let M be a compact Riemannian manifold, possibly with a boundary. The heat diffusion process over M is governed by the heat equation:

$$\Delta_M \mu(x,t) = -\frac{\Box \mu(x,t)}{\Box t} \qquad (2)$$

Here Δ_M is the Laplace-Beltrami operator of M. If M has boundaries, the Dirichlet boundary condition $\mu(x,t)=0$ for all $x \in \Box M$ and all t. Given an initial heat distribution $f: M \to \Box$, let $H_t(f)$ denote the heat distribution at time t, namely $H_t(f)$ satisfies the heat equation for all t, and $lim_{t \to 0} H_t(f) = f$. H_t is called the heat operator. It is well known [17] that for any M, there exists a function $k_t(x,y): \Box^+ \times M \times M \to \Box$ such that

$$H_t f(x) = \int_M k_t(x,y) f(y) dy \qquad (3)$$

Where dy is the volume form at $y \in M$. The minimum function $k_t(x,y)$ that satisfies Eq. (3) is called the heat kernel. Then, given a point x on the manifold M, the Heat Kernel Signature, HKS(x) is defined as a function over the temporal domain:

$$HKS(x): \Box^+ \to \Box, \; HKS(x,t) = k_t(x,x) \qquad (4)$$

The HKS has a very nice property: the surface HKS function of a point x is directly related to the Gaussian curvature on a surface at x [16], which is an intrinsic geometric property of the shape. This method offers the advantages of inelastic deformation invariance and is somewhat robust to topological noise. So it offers a reliable way for us to compute surface Gaussian curvature related statistics, e.g., for peak detection (**Figure 1**). HKS was computed using 1-100 (in unit increments) as the artificial time and then normalized by the maximum value. We chose the HKS values at t=1 as the HKS values for each ODF. Here, the local maxima of the ODFs were selected, based on the amplitude of the HKS value at each point in the spherical surface. For each ODF, we pick n points - let us assume p_i from ODF1 and q_i from ODF2 and i=1, 2, …, n. **Fig. 1** shows a few computed HKS examples on various ODFs. Once we define local maxima (p and q) for the ODFs, the next step is to find a diffeomorphism to match the ODFs. To match multiple peaks at the same time, we use Möbius transformation theory, as every Möbius transformation is a bijective conformal map of the Riemann sphere to itself.
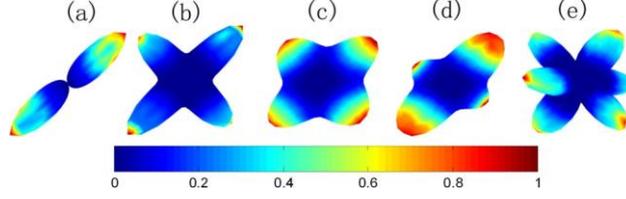
Fig. 1. Heat Kernel Signature for several typical kinds of orientation density functions (ODFs). (a) HKS for one fiber; (b) HKS for two sharply concentrated fibers crossing at 90 degrees; (c) HKS for two less concentrated fibers crossing at 90 degrees; (e) HKS for two unequally weighted fibers crossing at 90 degrees; (f) HKS for three fibers, each with an equal weighting. *Red colors* indicate high HKS values and *blue colors* indicate low HKS values. A high HKS value indicates a high Gaussian curvature point, i.e., a potential ODF local maximum.

2.3 Peak matching using Möbius Transformation and Möbius Voting

As any ODF is defined on a spherical domain, we can use stereographic projection to project the 3D spherical surface into the 2D complex plane. In the Cartesian coordinates (x, y, z) on the sphere and (X, Y) on the plane, the projection and its inverse are given by the formulae:

$$(X, Y) = (\frac{x}{1-z}, \frac{y}{1-z})$$

$$(x, y, z) = (\frac{2X}{1+X^2+Y^2}, \frac{2Y}{1+X^2+Y^2}, \frac{-1+X^2+Y^2}{1+X^2+Y^2}) \tag{5}$$

In geometry, a Möbius transformation of the plane is a rational function of the form:

$$f(z) = \frac{az+b}{cz+d} \tag{6}$$

- involving one complex variable z; here the *coefficients a, b, c, d* are complex numbers satisfying $ad - bc \neq 0$. In order to compare two ODFs, it is desirable to adjust the conformal mapping to match the geometric features on the ODF as well as possible. So we use the least-squares approach to compute an optimal Möbius transformation [18].

$$\{a, b, c, d\} = \arg\min \sum_{i=1}^{n} g(z_i) \left| \frac{az_i + b}{cz_i + d} - \tau_i \right|^2 \tag{7}$$

where z_i is the stereographic projection of p_i, τ_i is the projection of q_i. g is the conformal factor from the plane to the sphere, which may be simplified as:

$$g(z) = \frac{4}{1+z\bar{z}} \tag{8}$$

with $\bar{z}$ is the complex conjugate of complex number z. To ensure the peaks are matched in the right order, we use the "Möbius voting" concept [19]. Prior work [19] reported an algorithm to automatically discover surface point correspondences by optimizing Möbius transformation parameters. Isometries are a subset of the Möbius group. As the Möbius group is low-dimensional (6 DOF for the topological sphere), one can compute a closed-form Möbius transformation after selecting three matching

points. By permuting matching points, one can obtain "votes" for predicted corre-spondences by measuring the mismatch energy values. We borrow this idea here to find the best diffeomorphism to match different ODFs. The Möbius group is much bigger than congruence, e.g., voxel rotation, which has been typically used for HARDI registration before [5]. Möbius transformation still only has limited numbers of parameters, making subsequent optimization work simpler. In this study, we per-mute all possible combinations of matched features to minimize the L^2 norm between the transformed versions of ODF1 and ODF2 (if we assume we are mapping ODF1 to ODF2).

2.4   Overall framework

The current framework is summarized as follows:

(a) Use a preliminary whole-brain registration method (affine or non-linear) to reg-ister the entire diffusion image from one subject to a target subject. This en-sures all relevant brain tissues are roughly matched;

(b) Calculate the ODF at each voxel for both the source image (to be registered) and the target image;

(c) Compute the HKS for each ODF in both source and target and determine the peak number and locations for each ODF;

(d) For corresponding ODFs between source and target, use Möbius transformation (and voting) to achieve finer-scale alignment of peaks.

3   Experimental Results

3.1   ODF Rotation in a Synthetic Example

We studied different registration situations to match ODFs with different numbers of peaks ranging from 1 to 3 peaks. In a prior work [1], it was found that the number of detectable peaks in a voxel generally lies between 1 and 3, with more than one peak in more than 40% of voxels. A voxel with more than 3 fibers crossing is less plausible for neurobiological reasons, unless the voxels are so large that many tracts are present and in such case a very high angular sampling would be needed to resolve more peaks. To illustrate the simplest case, multiple fibers were created, crossing at 90 degrees with equal volume fractions. Here we chose $\lambda_1/\lambda_2=5:1$ as the ratio of ei-genvalues for each cylinder-shape tensor (FA= 0.77, typical for white matter in the brain) to generate simulated ODFs using discrete mixtures of Gaussian distributions (Eq. 9). The basic idea for ODF registration is to match the dominant directions. Our registrations included matching pairs of ODFs with different numbers of peaks, with results shown in Fig. 2.

$$S(q) = S_0 \cdot \sum_{i=1}^{n} w_i \cdot e^{-b \cdot q^T \cdot D_i \cdot q} \qquad \text{where} \quad \sum_{i=1}^{n} w_i = 1 \qquad (9)$$
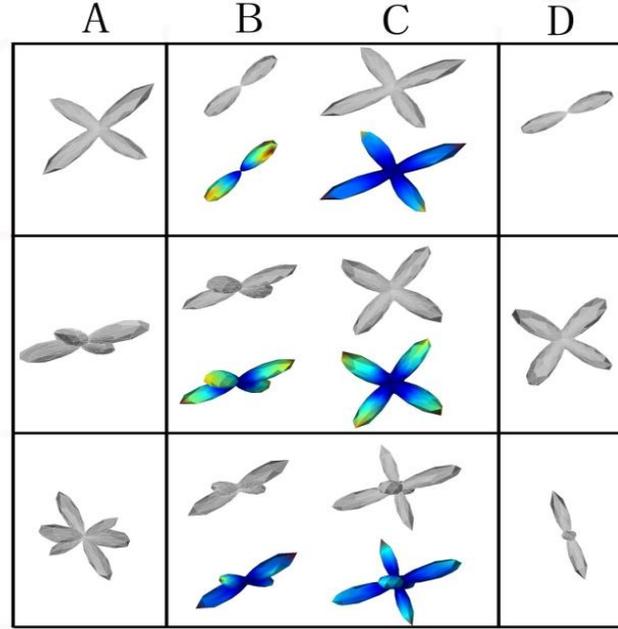
Fig. 2. Illustration of ODF rotation in the synthetic dataset. In each row, two ODFs and their corresponding HKS are shown in Columns B and C; column A shows the ODF computed from registering C to B, and column D shows the ODF computed from registering B to C. The first row shows the registration between 1 and 2, the second row for 2 vs. 2 and the last row is for 2 vs. 3.

The symmetric Kullback–Leibler distance (sKL) between the average registered ODFs and the template ODF was used to quantitatively evaluate the proposed method, according to Eq. **10**:

$$sKL = \frac{1}{2} \cdot \sum_{i=1}^{642} \left[p_i * \log\left(\frac{p_i}{q_i}\right) + q_i * \log\left(\frac{q_i}{p_i}\right) \right] \tag{10}$$

In this equation, p_i and q_i (i=1, 2…642) are the sampled points on the two ODF surfaces, as described in Section **2.1**.

In the simulation experiment (**Fig. 3**), we generated two fibers crossing at 90 degrees, and diffusion-weighted measurements were sampled from 94 directions (same as the protocol used for the human brain experiments in Section **3.2**). Rician noise (SNR=5) in was added to each of the diffusion-weighted measurements and the ODF was computed using TDF theory (Eq. **1**). We repeated this process 100 times, and we then calculated the averaged ODF with and without the adjustment of our proposed method. The sKL between the original ODF and averaged ODF - without adjustment - is 2.22 times higher than the sKL between original ODF and averaged ODF with the adjustment of our proposed method. This suggests how noise may affect the results, and how our proposed method should improve the statistical results. With the added noise, the averaged ODF – once it has been adjusted with our proposed method – offers a significant advantage in recognizing the dominant directions, compared to

using the averaged ODF without adjustment. In the next section, we evaluate our methods on human brain data.
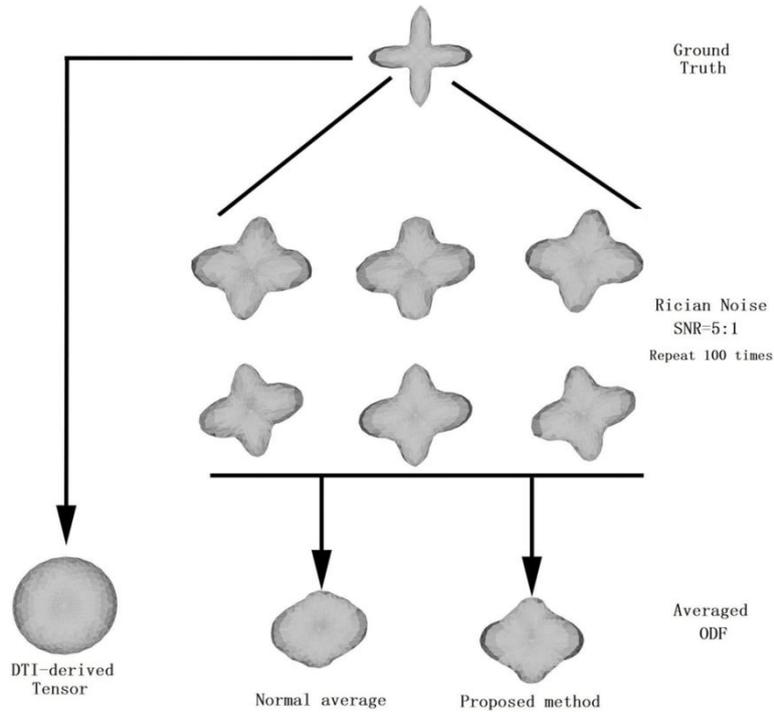

Fig. 3**. Validation of our proposed method on a synthetic dataset**. Here we define the ground truth as two fibers, crossing at 90 degrees. For this simulation, diffusion-weighed measurements were generated in each of 94 directions – the same gradient vector table as is used for our human datasets – and Rician noise (SNR=5) was added. Then the ODF was reconstructed using TDF theory. This process was repeated 100 times and all ODFs were averaged with and without adjustment by the proposed method. Without the alignment of peaks, the average ODF is very close to a sphere, and it is hard to see any evidence of peaks in the resulting ODF at all. The ODF that is created after peak matching is still not sharp – as it results from averaging many very noisy datasets, but it has more readily identifiable peaks. There is an evident advantage in retaining the dominant directions in the averaged ODF.

3.2 Real Brain ODF Transformation

Diffusion-weighted scans were acquired from 10 healthy adult participants (mean age: 23.8±2.4 years, 5 male) on a 4T MRI scanner. 105 gradient images were collected including 11 baseline (b0) images with no diffusion sensitization and 94 diffusion-weighted images (b-value: 1159 s/mm^2) (see [20] for details). Firstly, a non-linear registration method [21] was used to match the overall brain shape. After this pre-registration, TDF-derived ODFs for each subject were computed.

Before describing the adjustment of ODFs using our proposed methods, we first evaluate the necessity of performing ODF adjustment when computing ODF-based statistics. Here, we investigated the variation in the ODFs' dominant directions, among ten registered subjects [21]. The dominant direction of the ODF was defined as the direction with the maximum ODF surface value. **Fig. 4** shows the maximum dominant direction discrepancy (in units of radians) across ten subjects' corresponding voxels, for one brain slice. To further quantify the result, we calculated the variability in the dominant direction for the ROI-based mean ODF. In the corpus callosum ROI, the mean variation was 0.1 radian (4.6^0) while it was as high as 0.6 radian (33.8^0) in the fiber crossing region of the superior longitudinal fasciculus for ten registered subjects. This enormous degree of variance suggests that the ODF peak adjustment would be beneficial for ODF-based statistics, or else comparable regions would not be averaged together.
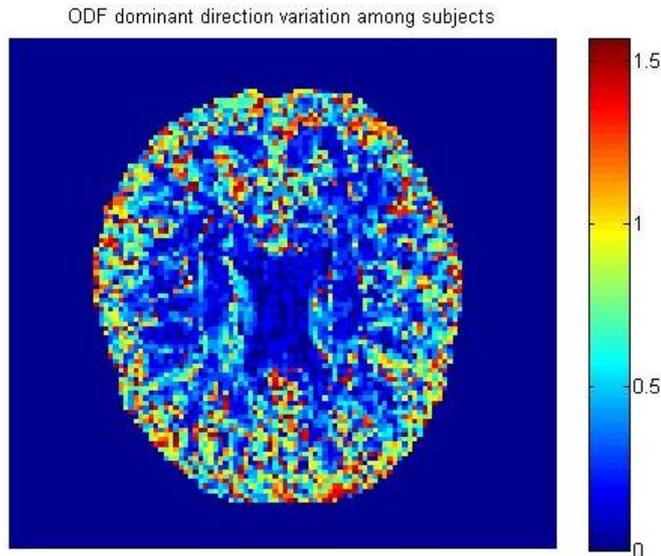

Fig. 4**. ODF dominant direction variation across subjects**. This result represents the ODF dominant direction discrepancy among ten registered subjects, in units of radians. The higher the value is, the greater is the disagreement in the ODF dominant directions among registered subjects. If fiber direction vary this much even in registered data, it makes sense to investigate further alignment of ODF peaks, before comparing ODFs across subjects.

Then we chose one subject as the target, we used our method to more accurately register all other subjects' ODF to the target. (Note that more complex methods could be used to select a target, e.g. the subject with least summed distance to the others, in some appropriate metric). **Fig. 5a** is the template subject's ODF; **Fig. 5b** shows the directly averaged ODFs across 10 subjects, using standard pointwise averaging of the ODF data across subjects, in spherical coordinates. **Fig. 5c** shows averaged registered ODFs across 10 subjects, after using the methods proposed in this paper. Compared to **Fig. 5a**, **Fig. 5b**'s main fiber tracts are blurred away due to the variance among the

subjects, and due to the lack of a method to align and constructively reinforce them. In **Fig. 5c**, the main fiber tracts have been retained and even reinforced after averaging ODF data across subjects.
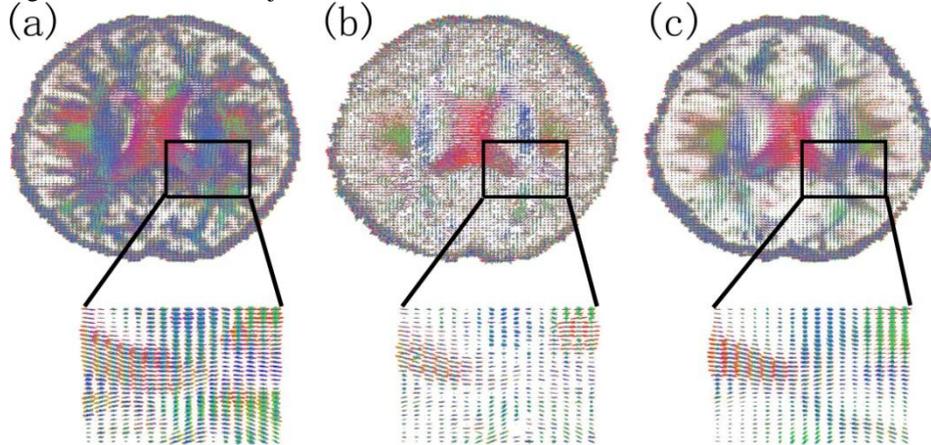

Fig. 5. ODF plots from individuals, and computation of population averages. **(a)** A typical subject's ODF field - this subject's scan was used as the target for registration, and other subjects' data was aligned to it; **(b)** directly averaged ODF across 10 subjects, without any ODF alignment - note the blurring of the features; **(c)** averaged registered ODF across 10 subjects, using the methods proposed in this paper. The colors represent the dominant fiber direction, mapped as a 3D vector to the RGB space: *red* for left-right, *blue* for superior-inferior, and *green* for anterior-posterior.

4   Conclusion and Future Work

Here we proposed a general method - based on the HKS and Möbius transformation - to register fields of spherical functions, such those that arise in analyses of diffusion imaging data from a population. The novelty of our work is twofold. First, we use HKS, which is proportional to Gaussian curvature, to objectively locate peaks; second, we optimize the matching energy between ODF by Möbius voting, to find the best transformation parameters.

Our paper has some limitations. It may not always be meaningful to align ODFs with different numbers of peaks; in this case, going further and explicitly extracting global paths for the fibers may help to resolve ambiguities. This may also help to compute a spatial coherent matching field across voxels. In future, we plan further numerical analysis of the accuracy of the ODF re-orientation at an individual level and of the accuracy when averaging entire ODF fields. Comparative studies with other ODF registration methods (e.g., with fidelity metrics based on information theory) would also be of interest.

Perhaps surprisingly, diffusion imaging has exquisite angular resolution but this is generally thrown away and reduced to a single scalar measure per voxel, prior to cross subject comparisons and population analyses. Here we show that by matching peaks

on ODFs, we can average data in way that constructively reinforces the available features. This effort should be useful in ongoing projects to define characteristic patterns of fiber integrity and geometry in disease, as well as changes with disease progression and over the human lifespan. In addition, Möbius transformations can do more than just rotate the full ODF, and could be part of future Möbius-driven global nonlinear registration, using a cost function derived from the local Möbius transformation to drive a global 3D image flow and then apply the image flow back to the ODF. Here we show HKS and Möbius transformation can reliably used for local maxima detection and matching in ODF registration.

References

1. Anonymous
2. Leow, A.D., Zhan, L., Zhu, S.W., Hageman, N,S., Chiang, M.C., Barysheva, M., Toga, A.W., McMahon, K., de Zubicaray, G.I., Wright, M.J., Thompson, P.M.: White matter integrity measured by fractional anisotropy correlates poorly with actual individual fiber anisotropy. In Proceedings of the 6th IEEE international Symposium on Biomedical Imaging: From Nano to Macro, ISBI2009, 622-625 (2009).
3. Lee, A.D., Lepore, N., Brun, C., Chou, Y.Y., Barysheva, M., Chiang, M.C., Madsen, S.K., de Zubicaray, G.I., McMahon, K.L., Wright, M.J., Toga, A.W., Thompson, P.M.: Tensor-based analysis of genetic influences on brain integrity using DTI in 100 twins. Med Image Comput Comput Assist Interv, 12:967-974 (2009).
4. Jin, Y., Shi, Y.G., Jahanshad, N., Aganj, I., Sapiro, G., Toga, A.W., Thompson, P.M.: 3D elastic registration improves HARDI-derived fiber alignment and automated tract clustering. In Proceedings of the 8th IEEE international Symposium on Biomedical Imaging: From Nano to Macro, ISBI2011, 822-826 (2011).
5. Barmpoutis, A., Hwang, M.S., Howland, D., Forder, J.R., Vemuri, B.C.: Regularized positive-definite fourth order tensor field estimation from DW-MRI. Neuroimage. 45: 153-162 (2009).
6. Geng, X., Ross, T.J., Zhan, W., Gu, H., Chao, Y.P., Lin, C.P., Christensen, G.E., Schuff, N., Yang, Y.: Diffusion MRI registration using orientation distribution functions. Inf Process Med Imaging, 21:626–637 (2009).
7. Cheng, G., Vemuri, B.C., Carney, P.R., Mareci, T.H.: Non-rigid Registration of High Angular Resolution Diffusion Images Represented by Gaussian Mixture Fields. Med Image Comput Comput Assist Interv, 5761: 190-197 (2009).
8. Hong, X., Arlinghaus, L.R., Anderson, A.W.: Spatial normalization of the fiber orientation distribution based on high angular resolution diffusion imaging data. Magn Reson Med. 61:1520–1527 (2009).
9. Bloy, L., Verma, R.: Demons registration of high angular resolution diffusion images. ISBI 1013–1016 (2010).
10. Yap, P.T., Chen, Y., An, H., Yang, Y., Gilmore, J.H., Lin, W., Shen, D.: SPHERE: SPherical Harmonic Elastic REgistration of HARDI data. Neuroimage. 55:545–556 (2011).
11. Goh, A., Lenglet, C., Thompson, P.M., Vidal, R.: A non-parametric Riemannian framework for processing high angular resolution diffusion images and its applications to ODF-based morphometry. Neuroimage. 56:1181–1201 (2011).
12. Chiang, M.C., Leow, A.D., Klunder, A.D., Dutton, R.A., Barysheva, M., Rose, S.E., McMahon, K.L., de Zubicaray, G.I., Toga, A.W., Thompson, P.M.: Fluid registration of

diffusion tensor images using information theory. IEEE Trans Med Imaging, 27:442–456 (2008).

13. 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. Magn Reson Med. 61:205–214. (2009).

14. Jian, B., Vemuri, B.C.: A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. IEEE Trans Med Imaging 26:1464–1471. (2007).

15. Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., Harel, N.: Reconstruction of the orientation distribution func-tion in single- and multiple-shell q-ball imaging within constant solid angle. Magn Reson Med. 64:554–566. (2010).

16. Sun, J.,Ovsjanikov, M., Guibas, L.: A concise and prov-ably informative multi-scale signature based on heat diffusion. Computer Graphics Forum 28(5):1383–1392.(2009).

17. Hsu, E.P. Stochastic Analysis on Manifolds. American Mathematical Society (2002).

18. Gu, X.F., Wang, Y.L., Chan, T.F., Thompson, P.M., Yau S.T.: Genus zero surface conformal mapping and its application to brain surface mapping. IEEE Trans. Med. Imag., 23(8):949–958 (2004).

19. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. SIGGRAPH, 72:1–12. (2009).

20. Anonymous

21. Jahanshad, N., Lee, A.D., Barysheva, M., McMahon, K.L., de Zubicaray, G.I., Martin, N.G., Wright, M.J., Toga, A.W., Thompson, P.M.: Genetics Influences on Brain Asymmetry: A DTI Study of 374 Twins and Siblings. Neuroimage 15; 52(2):455-69 (2010).