GROUP ACTION INDUCED AVERAGING FOR HARDI PROCESSING

H. Ertan Çetingül[⋆], *Bijan Afsari*[†], *Margaret J. Wright*[§], *Paul M. Thompson*[‡], *and René Vidal*[†]

[⋆]Imaging and Visualization, Siemens Corporate Research & Technology, Princeton, NJ, USA
[†]Center for Imaging Science, Johns Hopkins University, Baltimore, MD, USA
[§]QIMR and School of Psychology, University of Queensland, Brisbane, Australia
[‡]Laboratory of Neuro Imaging, UCLA School of Medicine, Los Angeles, CA, USA

ABSTRACT

We consider the problem of processing high angular resolution diffusion images described by orientation distribution functions (ODFs). Prior work showed that several processing operations, e.g., averaging, interpolation and filtering, can be reduced to averaging in the space of ODFs. However, this approach leads to anatomically erroneous results when the ODFs to be processed have very different orientations. To address this issue, we propose a group action induced distance for averaging ODFs, which leads to a novel processing framework on the spaces of orientation (the space of 3D rotations) and shape (the space of ODFs with the same orientation). Experiments demonstrate that our framework produces anatomically meaningful results.

Index Terms— biomedical image processing, information geometry, Riemannian manifolds, diffusion magnetic resonance imaging.

1. INTRODUCTION

High angular resolution diffusion imaging (HARDI) is a diffusion MRI technique that can be used to infer the tissue microstructure in vivo [1]. This requires the reconstruction of the orientation distribution function (ODF), a non-parametric probability density function (PDF) describing the anisotropy of water diffusion at a spatial location. The ODF model offers improved accuracy in resolving intra-voxel complexities over the diffusion tensor (DT) model [2], currently the de facto standard for neuroimaging.

Developing mathematical methods for processing fields of ODFs is important in many aspects. For instance, by computing the mean of a set of ODFs, one can statistically compare ODF images of several subjects. Similarly, almost every geometric transformation applied to a grayscale image requires interpolation between intensity values and this is also true in the case of ODF images. Convolution and filtering are also needed to denoise ODF images that are estimated from noisy HARDI signals. These operations should be reformulated to properly handle the mathematical structure of the space of ODFs.

Doing calculus with DTs or ODFs requires defining a metric to compare two such elements. In the case of DTI, there exist several frameworks based on the well-studied geometry of the space of second-order tensors. Typical examples employ affine-invariant [3] and Log-Euclidean metrics [4]. In the case of HARDI, one should consider the statistical manifold whose elements are PDFs. Existing frameworks for HARDI processing [5, 6] exploit the fact that ODFs are probability density functions on the 2-sphere $\mathbb{S}^2$ and that under a particular re-parametrization, the *square-root representation* of ODFs, various Riemannian operations are computable in closed-form.

H.E. Çetingül performed this work while at JHU. M.J. Wright and P.M. Thompson contributed to this work by providing the brain HARDI dataset.

In some cases, however, the methods of [5, 6] fail to provide anatomically meaningful results by producing ODF "bloating" and creating false fiber crossings. Fig. 1 shows two such cases: averaging (Fig. 1(a)) and interpolation (Fig. 1(b)). In the first case, we average ODFs with the same shape but different orientations. We expect a mean ODF with no difference in shape (red frame), but [5, 6] produce a "bloated" ODF (blue frame). In the second case, we interpolate two 1-fiber ODFs and expect to preserve their shapes (red frame), but [5, 6] produce 2-fiber (crossing) ODFs (blue frame).

To alleviate these issues, [7] suggests to separate the information from an ODF into orientation and shape components. In this work, we formalize this idea into a novel framework for HARDI processing. We propose a group action induced distance for ODFs, where the rotation group is the *orientation* space and the *shape* space is the space of ODFs with the same orientation. We compute the mean ODF by alternating between aligning the ODFs to the current estimate and updating its orientation and shape via Riemannian averaging in these two spaces. We evaluate our framework on synthetic and real data.

(a) Averaging a set of ODFs (green frame).

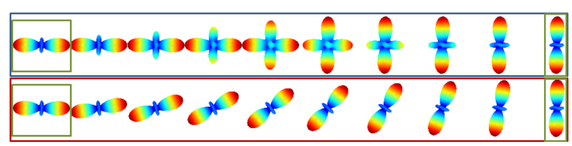

(b) Linear interpolation between pairs of ODFs (green frames)

Fig. 1. Averaging and linear interpolation of ODFs using existing frameworks (blue frame) and the proposed framework (red frame).

2. RIEMANNIAN MANIFOLD OF ODFs

On the manifold of ODFs, i.e., the space of PDFs p on $\mathbb{S}^2$, the Fisher-Rao (FR) metric [8] determines a Riemannian metric that is invariant to re-parameterizations. By using a particular re-parameterization, the *square-root representation*, the manifold of ODFs is a unit sphere in a Hilbert space with the Fisher-Rao metric being the $\mathbb{L}^2$ metric [9]. Then the space of *square-root density functions* $\psi(\mathbf{s}) \triangleq \sqrt{p(\mathbf{s})}$ is

$$\Psi = \Big\{\psi : \mathbb{S}^2 \to [0,1] \mid \forall \mathbf{s} \in \mathbb{S}^2, \psi(\mathbf{s}) \geq 0\, ;\int_{\mathbb{S}^2} \psi^2(\mathbf{s}) d\mathbf{s} = 1\Big\}. \quad (1)$$

Ψ is the non-negative orthant of a unit Hilbert sphere and the geodesic distance between $\psi_i, \psi_j \in \Psi$ is the angle between them, i.e.,

$$d_{FR}(\psi_i, \psi_j) = \cos^{-1}\langle \psi_i, \psi_j \rangle = \cos^{-1}\Big(\int_{\mathbb{S}^2} \psi_i(\mathbf{s})\psi_j(\mathbf{s}) d\mathbf{s}\Big). \quad (2)$$

Let $T_\psi\Psi$ denote the tangent space of Ψ at ψ. Under the square-root representation, several Riemannian operations such as the *exponential map* $\exp : T_\psi\Psi \to \Psi$ and *logarithm map* $\log : \Psi \to T_\psi\Psi$ can be computed in closed-form [5]. In practice the ODF p is represented in terms of its samples at M directions as the vector $\mathbf{p}\in\mathbb{R}^M$. In our implementation, we use the vector $\boldsymbol{\psi} \triangleq \sqrt{\mathbf{p}}$, the M-bin histogram of ψ, such that $\boldsymbol{\psi}\in\mathbf{\Psi}\subset\mathbb{S}^{M-1}$. For the sake of clarity, we refer to the function ψ as "ODF" in the following discussions.

Riemannian Averaging of ODFs. The Riemannian weighted average of N ODFs $\{\psi_n\}_{n=1}^N \subset \Psi$ with weights $\{w_n\}_{n=1}^N$ is the solution to $\operatorname{argmin}_\psi \sum_{n=1}^N w_n d_{FR}^2(\psi, \psi_n)$. It is computed via Riemannian gradient descent, which involves consecutive evaluations of the aforementioned logarithm and exponential maps until a convergence criterion is met. The reader is referred to [5, 11] for further details.

3. GROUP ACTION INDUCED AVERAGING FOR PROCESSING ODFs

3.1. A Group Action Induced Distance between ODFs

To explicitly represent the orientation information of an ODF, we use the action of the rotation group $\mathrm{SO}(3)$ on Ψ. The rotation group $\mathrm{SO}(3) = \{\mathbf{R} \in \mathbb{R}^{3\times3} | \mathbf{R}^\top = \mathbf{R}^{-1}, \det(\mathbf{R}) = 1\}$ is a *matrix Lie group* whose tangent space at the identity element $\mathbf{I} \in \mathbb{R}^{3\times3}$ is the *Lie algebra* $\mathrm{so}(3)$, which is the space of 3×3 skew-symmetric matrices. The geodesic distance between two elements $\mathbf{R}_i, \mathbf{R}_j \in \mathrm{SO}(3)$ is

$$d_{\mathrm{SO}(3)}(\mathbf{R}_i, \mathbf{R}_j) = (1/\sqrt{2})\|\mathrm{Log}(\mathbf{R}_i^\top \mathbf{R}_j)\|_F, \tag{3}$$

where $\|\cdot\|_F$ is the Frobenius norm and Log is the matrix logarithm (see [10, 11] for details on the Riemannian operations and optimization in $\mathrm{SO}(3)$). The action of $\mathrm{SO}(3)$ on Ψ is a rotation around the origin, i.e., if $\mathbf{R} \in \mathrm{SO}(3)$ and $\psi \in \Psi$, then the group action $\mathbf{R} \circ \psi \in \Psi$ is defined as $\mathbf{R} \circ \psi(\mathbf{s}) = \psi(\mathbf{R}\mathbf{s}), \forall \mathbf{s} \in \mathbb{S}^2$.

We define the rotation group action induced distance in Ψ as

$$d_\Psi^2(\psi_i, \psi_j) = \min_{\mathbf{R}\in\mathrm{SO}(3)} d_{FR}^2(\psi_i, \mathbf{R} \circ \psi_j) + \lambda\, d_{\mathrm{SO}(3)}^2(\mathbf{I}, \mathbf{R}). \tag{4}$$

The first term measures the dissimilarity in *shape* when the ODFs ψ_i and ψ_j are "aligned", while the second term measures the amount of alignment needed, i.e., the dissimilarity in *orientation*. The parameter λ is a trade-off parameter between the shape and orientation terms and needs to be chosen such that $\lambda > 0$. Otherwise, the ODF alignment problem in (4) may not have a unique solution due to the possible symmetries of the ODFs. Finally, notice also that d_Ψ is a rotation-invariant distance on Ψ, i.e., for all $\psi_i, \psi_j \in \Psi$ and all $\mathbf{R} \in \mathrm{SO}(3)$, we have $d_\Psi(\psi_i, \psi_j) = d_\Psi(\mathbf{R}\circ\psi_i, \mathbf{R}\circ\psi_j)$, because d_{FR} is rotation-invariant and $d_{FR}(\mathbf{R}_i \circ \psi_i, \mathbf{R}_j \circ \psi_j) = d_{FR}(\psi_i, \mathbf{R}_j\mathbf{R}_i^\top \circ \psi_j)$.

3.2. Group Action Induced Averaging of ODFs

The average of N ODFs $\{\psi_n\}_{n=1}^N$ with respect to d_Ψ is defined as

$$\tilde{\tilde{\psi}} = \operatorname*{argmin}_{\psi\in\Psi} \sum_{n=1}^N w_n (\min_{\mathbf{R}_n} d_{FR}^2(\psi, \mathbf{R}_n \circ \psi_n) + \lambda d_{\mathrm{SO}(3)}^2(\mathbf{I}, \mathbf{R}_n)), \tag{5}$$

where $\psi = \mathbf{R} \circ \phi$ is a decomposition of ψ into its orientation and shape components $\mathbf{R}$ and ϕ, respectively. We propose to compute $\tilde{\tilde{\psi}}$ via an iterative method that alternates between aligning the ODFs to the current estimate of $\tilde{\tilde{\psi}}$, updating the orientation component given the current alignments and shape component, and updating the shape component given the current alignments and orientation component.

To derive the algorithm, let $\mathbf{Q}_n = \mathbf{R}_n\mathbf{R}^\top$ and notice that

$$d_{FR}^2(\mathbf{R}\circ\phi, \mathbf{R}_n\circ\psi_n) = d_{FR}^2(\phi, \mathbf{R}_n\mathbf{R}^\top\circ\psi_n) = d_{FR}^2(\phi, \mathbf{Q}_n\circ\psi_n)$$
$$d_{\mathrm{SO}(3)}^2(\mathbf{I}, \mathbf{R}_n) = d_{\mathrm{SO}(3)}^2(\mathbf{I}, \mathbf{Q}_n\mathbf{R}) = d_{\mathrm{SO}(3)}^2(\mathbf{R}^\top, \mathbf{Q}_n). \tag{6}$$

Therefore, we can rewrite (5) as

$$\min_{\phi,\mathbf{R},\{\mathbf{Q}_n\}} \sum_{n=1}^N w_n (d_{FR}^2(\phi, \mathbf{Q}_n \circ \psi_n) + \lambda d_{\mathrm{SO}(3)}^2(\mathbf{R}^\top, \mathbf{Q}_n)). \tag{7}$$

We solve this problem using an alternating minimization strategy.

In the *alignment step*, we solve for $\mathbf{Q}_n$ assuming that we have estimates $\phi = \phi^k$ and $\mathbf{R} = \mathbf{R}^k$ at iteration k. From (7) we obtain

$$\mathbf{Q}_n^{k+1} = \operatorname*{argmin}_{\mathbf{Q}_n} d_{FR}^2(\phi^k, \mathbf{Q}_n \circ \psi_n) + \lambda d_{\mathrm{SO}(3)}^2(\mathbf{R}^{k\top}, \mathbf{Q}_n). \tag{8}$$

By setting $\mathbf{Q}_n = \mathbf{R}_n\mathbf{R}^{k\top}$ and $\psi^k = \mathbf{R}^k \circ \phi^k$, notice that this problem is equivalent to $\min_{\mathbf{R}_n} d_{FR}^2(\psi^k, \mathbf{R}_n \circ \psi_n) + \lambda d_{\mathrm{SO}(3)}^2(\mathbf{I}, \mathbf{R}_n)$, which is of the form in (4). We compute $\mathbf{Q}_n^{k+1}$ numerically using the *Nelder-Mead method* [12] extended from $\mathbb{R}^n$ to $\mathrm{SO}(3)$. The method is initialized with the *icosahedral rotation group*, a finite point subgroup of $\mathrm{SO}(3)$, and converges to a local minimum.

In the *orientation step*, we solve for $\mathbf{R}$ with $\phi = \phi^k$ and $\mathbf{Q}_n = \mathbf{Q}_n^{k+1}$. From (7) we obtain

$$\mathbf{R}^{k+1} = \operatorname*{argmin}_{\mathbf{R}} \sum_{n=1}^N w_n d_{\mathrm{SO}(3)}^2(\mathbf{R}^\top, \mathbf{Q}_n^{k+1}). \tag{9}$$

Hence, $(\mathbf{R}^{k+1})^\top$ is the Riemannian average of the current alignments, which can be computed using Riemannian gradient descent.

In the *shape step*, we solve for ϕ with $\mathbf{R} = \mathbf{R}^{k+1}$ and $\mathbf{Q}_n = \mathbf{Q}_n^{k+1}$. From (7) we obtain

$$\phi^{k+1} = \operatorname*{argmin}_{\phi} \sum_{n=1}^N w_n d_{FR}^2(\phi, \mathbf{Q}_n^{k+1} \circ \psi_n). \tag{10}$$

Thus, ϕ^{k+1} is the Riemannian average with respect to d_{FR} of the rotated ODFs $\{\mathbf{Q}_n^{k+1} \circ \psi_n\}_{n=1}^N$, which can be computed using the algorithm of [5], as described in §2. Finally, the new estimate of mean ODF $\tilde{\tilde{\psi}}$, ψ^{k+1}, is obtained by composing the orientation and shape updates as $\psi^{k+1} = \mathbf{R}^{k+1} \circ \phi^{k+1}$. The above alternating minimization algorithm converges to a local minimum of the cost function in (7). In our experiments we run a single iteration of this method starting at $\mathbf{R}^0 = \mathbf{I}$ and $\phi^0 = \psi_n$ for any $n = 1, \ldots, N$.

3.3. Group Action Induced Interpolation and Filtering of ODFs

Having presented how to perform averaging of ODFs, let us now consider the problems of interpolation and filtering. Assuming that we know N endpoints $\{\boldsymbol{x}_n\}_{n=1}^N$ in a multi-dimensional lattice, interpolation of ODFs at point $\boldsymbol{x}$ can be defined as the solution to

$$\psi(\boldsymbol{x}) = \operatorname*{argmin}_{\psi\in\Psi} \sum_{n=1}^N w_n(\boldsymbol{x}) d_\Psi^2\big(\psi, \psi(\boldsymbol{x}_n)\big), \tag{11}$$

where $\psi(\boldsymbol{x}_n)$ is the ODF at point $\boldsymbol{x}_n$ and $w_n(\boldsymbol{x})$ is the corresponding interpolation weight. Notice that $\psi(\boldsymbol{x})$ is the weighted average of $\{\psi(\boldsymbol{x}_n)\}_{n=1}^N$ and can be computed from the algorithm in §3.2.

Discrete convolution of an ODF image ψ with the filter $g \geq 0$ of spatial support $\mathcal{U}$ can be written as

$$\varphi(\boldsymbol{x}) = \operatorname*{argmin}_{\varphi\in\Psi} \sum_{\boldsymbol{u}\in\mathcal{U}} g(\boldsymbol{u}) d_\Psi^2(\varphi(\boldsymbol{u}), \psi(\boldsymbol{x} - \boldsymbol{u})). \tag{12}$$

Again, notice that the filtered ODF $\varphi(\boldsymbol{x})$ is the weighted average of $\{\psi(\boldsymbol{x} - \boldsymbol{u})\}_{\boldsymbol{u}\in\mathcal{U}}$ and can be computed from the algorithm in §3.2.

4. VALIDATION AND DISCUSSIONS

4.1. Experiments on Synthetic Data

We generate the synthetic data using the multi-tensor model, where the HARDI signal at a gradient direction $\mathbf{g}$ is a convex combination of functions $e^{-b\mathbf{g}^\top \mathbf{D}_k \mathbf{g}}$, $\mathbf{D}_k$ being the k-th tensor and $b = 3,000$ s/mm^2. We simulate the signal at 81 gradient directions, add complex Gaussian noise, and reconstruct the ODFs as described in [13].

Averaging. We generate five sets of ODFs (Fig. 2) to illustrate the difference between the averages $\bar{\psi}$ and $\bar{\tilde{\psi}}$ computed using d_{FR} and d_Ψ, respectively. The ODFs are generated by gradually rotating and/or changing the shape of the leftmost ODFs as well as adding noise to the signals they are reconstructed from. Thus, we expect to obtain mean ODFs "similar" to the fifth ODFs in the sets. In the first set, the ODFs do vary in shape but not in orientation, and hence both methods yield the same mean ODF. In the second set, the ODFs have the same shape but different orientations and averaging using d_{FR} produces a bloating effect on the resulting ODF. Our framework, on the other hand, produces an ODF with the same shape as that of the ODFs it is computed from and its orientation is the average of their orientations. For the remaining sets, averaging with d_Ψ produces more meaningful results than averaging with d_{FR}.

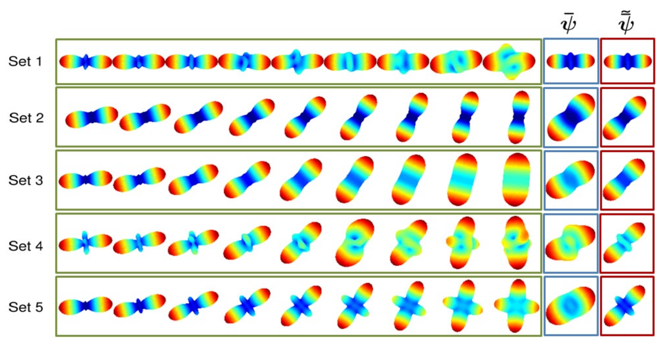

Fig. 2. Averaging ODFs (green frames): the averages $\bar{\psi}$ (blue frames) and $\bar{\tilde{\psi}}$ (red frames) are computed using d_{FR} and d_Ψ, respectively.

Interpolation. We first perform linear interpolation between three pairs of ODFs (Fig. 3(a)). In the first and second cases, interpolation using d_{FR} produces a 2-fiber ODF at the mid-point, whereas using d_Ψ interpolates the ODFs via rotation. In the third case, two 2-fiber ODFs that are different in orientation and in shape are interpolated. We observe that our framework prevents the bloating effect observed when interpolation is performed using d_{FR}. More precisely, the entropy of the interpolated ODFs is higher than that of the ODFs being interpolated when interpolation is performed using d_{FR} (Fig. 3(b)), whereas interpolation using d_Ψ results in an approximately linear relationship between the entropies. We also perform 2D interpolation between four ODFs, which have either different orientations but the same shape (Figs. 4(a)-4(b)), or different shapes including crossing configurations (Figs. 4(c)-4(d)). We observe that interpolation using d_{FR} produces ODFs with large amounts of bloating, whereas our framework yields smooth and visually appealing interpolations.

Filtering. We convolve an ODF image (Fig. 5(a)) with sharp discontinuities (in orientation) with a 5×5 Gaussian filter of standard deviation 1. Figs. 5(b)-5(c) show the resulting images after filtering using d_{FR} and d_Ψ, respectively. We observe that our framework yields a more realistic smoothed version of the input image.

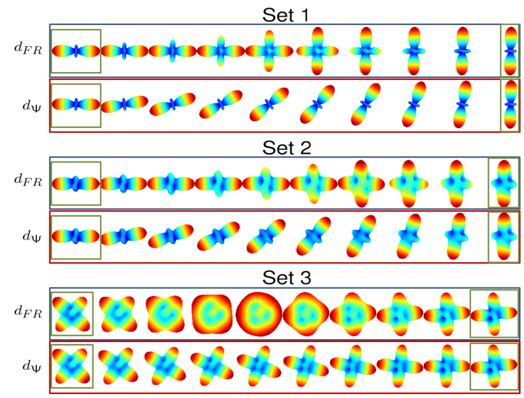

(a) Interpolation between the leftmost and rightmost ODFs

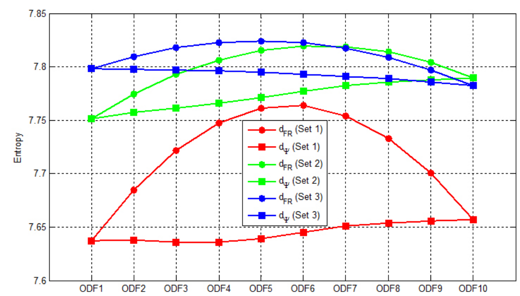

(b) Entropies of the ODFs in Sets 1, 2, and 3

Fig. 3. (a) Interpolation between pairs of ODFs using d_{FR} (blue frames) and using d_Ψ (red frames). (b) Study of bloating effects for interpolation using d_{FR} and using d_Ψ.

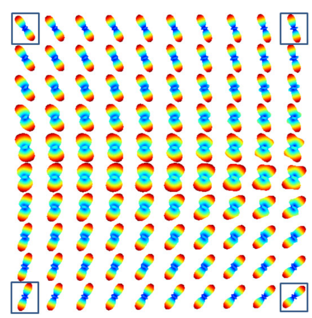

(a) Interpolation using d_{FR}

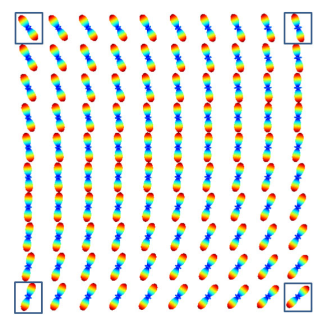

(b) Interpolation using d_Ψ

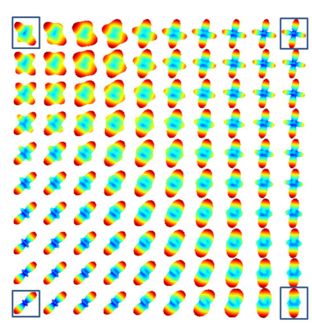

(c) Interpolation using d_{FR}

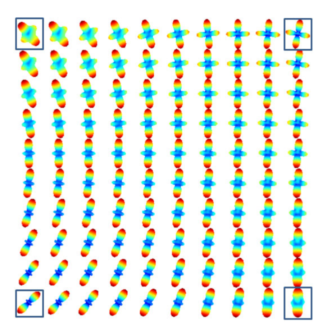

(d) Interpolation using d_Ψ

Fig. 4. 2D interpolation among four ODFs (blue frames): (a,b) Input ODFs have the same shape but different orientations; (c,d) Input ODFs have different shapes and orientations.

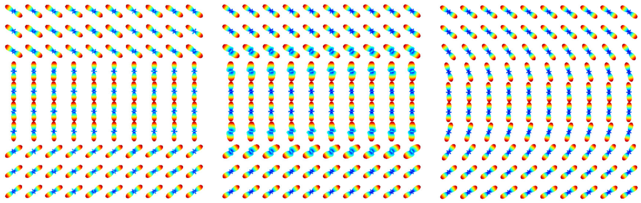

(a) (b) (c)

Fig. 5. (a) ODF image to be smoothed, (b) Image obtained after filtering using d_{FR}, (c) Image obtained after filtering using d_{Ψ}.

4.2. Experiments on Real Data

We also evaluate our framework on a human brain HARDI dataset [14] where 105 images were acquired, 11 with no diffusion weighting and 94 with diffusion weighting at $b = 1,159$ s/mm^2, by using a 128×128 acquisition matrix (1.8 mm in-plane resolution) and 55 axial slices (2 mm thick). Prior to ODF reconstruction, the diffusion weighted images of each subject are nonlinearly registered to a group-averaged template. We perform averaging of the ODFs over 20 subjects in a ROI containing parts of the corpus callosum and cingulum. These tracts do not intersect, but there exist partial volume averaging due to low image resolution, as illustrated in Figs. 6(a)-6(b). Fig. 6(d) shows the mean ODF image computed using d_{Ψ}, where both tracts are accurately delineated without ODF bloating. Our framework prevents the generation of spurious ODFs with crossing configurations by considering the orientation as a separate entity. As a result, the number of voxels with partial volume averaging (see number of the blue frames in Fig. 6(d)) is reduced compared to that in the mean ODF image computed using d_{FR} (see number of the blue frames in Fig. 6(c)). This demonstrates that by using the proposed framework, one can obtain anatomically more meaningful averages.

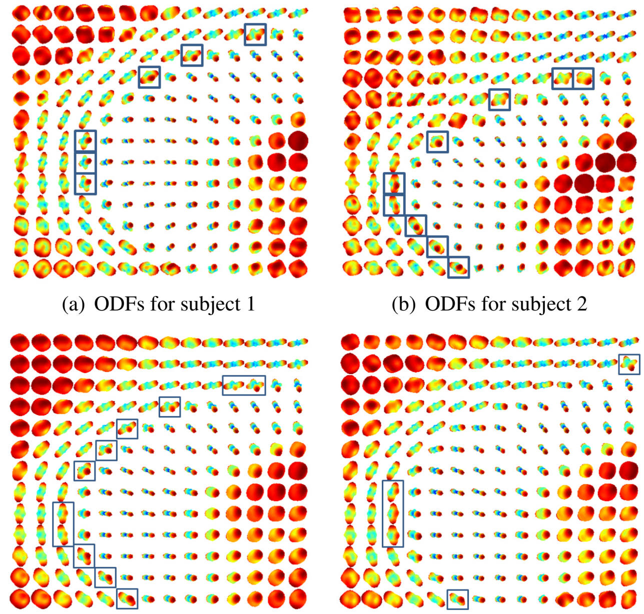

(a) ODFs for subject 1 (b) ODFs for subject 2

(c) Mean ODFs using d_{FR} (d) Mean ODFs using d_{Ψ}

Fig. 6. (a,b) ODF images of selected subjects; (c,d) The mean ODF images computed using d_{FR} and using d_{Ψ}, respectively. The ODFs in the blue frames show partial volume averaging.

5. CONCLUSIONS AND FUTURE WORK

We presented a novel framework for performing averaging, interpolation, and filtering of ODFs. We showed how these operations reduce to weighted averaging and proposed the rotation induced distance d_{Ψ}, which led to an iterative method that updates the orientation and shape of the current estimate via Riemannian averaging in the orientation and shape spaces, respectively. We demonstrated that our framework produces anatomically more meaningful results, especially in highly anisotropic regions, by reducing partial volume averaging and eliminating unrealistic changes in ODFs. As a future work, we aim to improve the quality of white matter atlases by averaging ODFs across several subjects using d_{Ψ}, and the accuracy of tractography by increasing the coherence among ODF orientations via filtering.

6. REFERENCES

[1] D.S. Tuch, "Q-ball imaging," *Magnetic Resonance in Medicine*, vol. 52, no. 6, pp. 1358–1372, 2004.

[2] P.J. Basser, J. Mattiello, and D. LeBihan, "Estimation of the effective self-diffusion tensor from the NMR spin echo," *Journal of Magnetic Resonance B*, vol. 103, no. 3, pp. 247–254, 1994.

[3] X. Pennec, P. Fillard, and N. Ayache, "A Riemannian framework for tensor computing," *International Journal of Computer Vision*, vol. 66, no. 1, pp. 41–46, 2006.

[4] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, "Log-Euclidean metrics for fast and simple calculus on diffusion tensors," *Magnetic Resonance in Medicine*, vol. 56, no. 2, pp. 411–421, 2006.

[5] A. Goh, C. Lenglet, P.M. Thompson, and R. Vidal, "A nonparametric Riemannian framework for processing high angular resolution diffusion images and its applications to ODF-based morphometry," *NeuroImage*, vol. 56, no. 1, pp. 1181–1201, 2011.

[6] J. Cheng, A. Ghosh, T. Jiang, and R. Deriche, "A Riemannian framework for orientation distribution function computing," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2009*, 2009, vol. LNCS 5761, pp. 911–918.

[7] S. Ncube and A. Srivastava, "A novel Riemannian metric for analyzing HARDI data," in *Proceedings of the SPIE*, 2011, vol. 7962.

[8] C.R. Rao, "Information and accuracy attainable in the estimation of statistical parameters," *Bulletin of the Calcutta Mathematical Society*, vol. 37, pp. 81–89, 1945.

[9] A. Srivastava, I. Jermyn, and S. Joshi, "Riemannian analysis of probability density functions with applications in vision," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1–8.

[10] M. Moakher, "Means and averaging in the group of rotations," *SIAM Journal on Matrix Analysis and Applications*, vol. 24, no. 1, pp. 1–16, 2002.

[11] B. Afsari, R. Tron, and R. Vidal, "On the convergence of gradient descent for finding the Riemannian center of mass," *arXiv:1201.0925v1 [math.DG]*.

[12] J.A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308–313, 1965.

[13] A. Goh, C. Lenglet, P.M. Thompson, and R. Vidal, "Estimating orientation distribution functions with probability density constraints and spatial regularity," in *Medical Image Computing and Computer-Assisted Intervention*, 2009, vol. LNCS 5761, pp. 877–885.

[14] M.-C. Chiang, M. Barysheva, A.D. Lee, S.K. Madsen, A.D. Klunder, A.W. Toga, K. McMahon, G.I. de Zubicaray, M. Meredith, M.J. Wright, A. Srivastava, N. Balov, and P.M. Thompson, "Brain fiber architecture, genetics, and intelligence: A high angular resolution diffusion imaging (HARDI) study," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2008*, 2008, vol. LNCS 5241, pp. 1060–1067.