

SPARSE RIEMANNIAN MANIFOLD CLUSTERING FOR HARDI SEGMENTATION

H. Ertan Çetingül and René Vidal

Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21218, USA

ABSTRACT

We address the problem of segmenting high angular resolution diffusion images of the brain into cerebral regions corresponding to distinct white matter fiber bundles. We cast this problem as a manifold clustering problem in which distinct fiber bundles correspond to different submanifolds of the space of orientation distribution functions (ODFs). Our approach integrates tools from sparse representation theory into a graph theoretic segmentation framework. By exploiting the Riemannian properties of the space of ODFs, we learn a sparse representation for the ODF at each voxel and infer the segmentation by applying spectral clustering to a similarity matrix built from these representations. We evaluate the performance of our method via experiments on synthetic, phantom and real data.

Index Terms— image segmentation, compressed sensing, graph theory, harmonic analysis, diffusion magnetic resonance imaging.

1. INTRODUCTION

White matter fiber clustering provides significant insights into anatomical structures of critical importance in atlas generation and statistical analysis of several neurodegenerative diseases. This task can particularly benefit from advanced diffusion MR imaging (dMRI) techniques such as high angular resolution diffusion imaging (HARDI), which has become a remarkable tool to study the microstructure of biological tissues. HARDI provides adequate information to estimate the orientation distribution function (ODF) [1], which characterizes the probability distribution of water diffusion along several directions. The ODF offers improved accuracy over the diffusion tensor (DT) model in quantifying the anisotropy of the medium and resolving complexities such as crossing or fanning fibers. However, in order to make this particular representation beneficial for diagnostic and clinical applications, novel segmentation schemes need to be developed so as to properly handle the mathematical structure of the data.

In contrast to the plethora of works on DT image segmentation proposed in the last decade –see [2] for a concise summary–, there are only a handful of segmentation methods for HARDI data. In [3, 4], the segmentation is performed on a non-Euclidean 5-D feature space, named as the position orientation space. The algorithm is implemented with a hidden Markov random field framework in [3] and with a level set framework in [4]. However, it yields binary, i.e., object vs. background, segmentation of relatively small cerebral regions due to its heavy computational load. In [5], the ODF is modeled as a mixture of von Mises-Fisher distributions and the parametric segmentation is obtained by using hidden Markov measure field models. The scheme in [6] uses the spherical harmonic (SH) expansion of the HARDI signal at every voxel to construct a feature vector of real SH coefficients. It subsequently uses the ℓ_2 -norm as a similarity measure and applies spectral embedding and clustering of selected fiber bundles via diffusion maps. A similar SH-based

This work was supported by startup funds from the JHU WSE.

representation for the ODFs is employed in [2] and the segmentation of specific white matter regions, such as the corpus callosum and the corticospinal tract, is performed via statistical surface evolution.

In this work, we present an algorithm to automatically segment HARDI data into multiple cerebral regions corresponding to different white matter fiber bundles. Our approach integrates tools from sparse representation theory into a graph theoretic segmentation framework that generalizes the sparse subspace clustering algorithm [7] from Euclidean data to ODFs. Specifically, we exploit the Riemannian properties of the space of ODFs to reformulate the problem of computing the sparse representation of an ODF with respect to its neighboring ODFs. These sparse representations are then used to build a block-sparse similarity matrix to which we apply spectral clustering for identifying the regions with distinct diffusion properties.

2. ODF ESTIMATION WITH SPHERICAL HARMONICS

In order to estimate the ODFs, we employ the reconstruction method proposed in [8]. This method follows the basic principle in Q-ball imaging (QBI) [1], i.e., estimating the ODFs directly from HARDI signals via the Funk-Radon transform. In particular, it approximates the signal at each gradient direction by a linear combination of the modified SH basis and reconstructs sharp ODFs while enforcing spatial regularity and nonnegativity of the ODF field. Section 2.1 reviews the basics of the spherical harmonic representation and Section 2.2 summarizes the ODF reconstruction method.

2.1. Modified Spherical Harmonic Basis

Due to the physics of the diffusion phenomenon, the HARDI signal $S : S^2 \mapsto \mathbb{C}$ is often assumed to be real and antipodally symmetric. Therefore, using the fact that any complex-valued function defined on the 2-sphere S^2 can be expressed as a combination of spherical harmonics, the SH representation of the signal uses a modified SH basis that is also real and symmetric. Specifically, let us consider the spherical harmonic function $Y_l^m : S^2 \mapsto \mathbb{C}$ of degree l and order m ,

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

where P_l^m is the associated Legendre polynomial of the same degree and order, and $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$ are the polar and azimuthal angles, respectively, representing the direction of interest (θ, ϕ) in spherical coordinates. Then, following [9], the modified basis of degree L contains $R = \frac{(L+1)(L+2)}{2}$ spherical harmonics of the form

$$Y_r = \begin{cases} \sqrt{2} \operatorname{Re}\{Y_l^{|m|}\} & \text{if } -l \leq m < 0, \\ Y_l^m & \text{if } m = 0, \\ \sqrt{2}(-1)^{m+1} \operatorname{Im}\{Y_l^m\} & \text{if } 0 < m \leq l, \end{cases} \quad (2)$$

with indices $r(l, m) = \frac{l^2+l+2}{2} + m$, where $l = 0, 2, 4, \dots, L$, $|m| \leq l$, and $\operatorname{Re}\{\cdot\}$ and $\operatorname{Im}\{\cdot\}$ are the real and imaginary parts, respectively.

2.2. ODF Reconstruction

Let S_0 denote the signal with no diffusion weighting, i.e., the baseline signal, and let $S(\theta, \phi)$ be the HARDI signal in the gradient direction (θ, ϕ) . In single shell QBI with solid angle consideration [10], the ODF in the spatial direction (ϑ, φ) is given by

$$p(\vartheta, \varphi)|_{\vartheta=\theta, \varphi=\phi} = \frac{1}{4\pi} + \frac{1}{16\pi^2} \text{FRT} \left\{ \nabla_b^2 \ln \left(-\ln \left(\frac{S(\theta, \phi)}{S_0} \right) \right) \right\}, \quad (3)$$

where FRT is the Funk-Radon transform and ∇_b^2 is the Laplace-Beltrami operator independent of the radial component.

The ODF, p , is estimated from G samples of the signal $S(\theta_n, \phi_n)$, $n = 1, 2, \dots, G$, by approximating the signal vector

$$\mathbf{s} := \left[\ln \left(-\ln \left(\frac{S(\theta_1, \phi_1)}{S_0} \right) \right), \dots, \ln \left(-\ln \left(\frac{S(\theta_G, \phi_G)}{S_0} \right) \right) \right]^\top \quad (4)$$

as $\mathbf{s} \approx \mathbf{B}\mathbf{c}$, where \mathbf{B} is the $G \times R$ SH basis matrix whose n -th row is $\mathbf{B}_n = [Y_1(\theta_n, \phi_n), \dots, Y_R(\theta_n, \phi_n)]$ and $\mathbf{c} \in \mathbb{R}^R$ is the unknown vector of SH coefficients parameterizing the signal. In the estimation, we consider the SH basis of degree $L = 4 \Rightarrow R = 15$. In [10], \mathbf{c} is estimated as the least-squares solution of the approximation, whereas in [8], the problem is reformulated to ensure the nonnegativity of the ODFs and to enforce spatial regularity between neighboring ODFs. In particular, assuming a tessellation scheme with $M = 162$ gradient directions $\{(\vartheta_n, \varphi_n)\}_{n=1}^M$, the ODF reconstructed at voxel i is represented in terms of its samples at these M directions as the following vector $\mathbf{p}_i \in \mathbb{R}^M$

$$\mathbf{p}_i \sim \frac{1}{4\pi} \mathbf{1} + \frac{1}{16\pi^2} \mathbf{CLP}\mathbf{c}_i \quad \text{s.t.} \quad \|\mathbf{p}_i\|_1 = 1, \quad (5)$$

where $\mathbf{1}$ is the $M \times 1$ vector of 1's, \mathbf{C} is the $M \times R$ SH basis matrix whose n -th row is $\mathbf{C}_n = [Y_1(\vartheta_n, \varphi_n), \dots, Y_R(\vartheta_n, \varphi_n)]$, \mathbf{L} is the $R \times R$ diagonal matrix of Laplace-Beltrami eigenvalues, and \mathbf{P} is the $R \times R$ diagonal Funk-Radon transform matrix. In [8], by using (5) and $p(\vartheta_n, \varphi_n) \geq 0, \forall n$, the SH coefficient vectors at N voxels, $\{\mathbf{c}_i\}_{i=1}^N$, are found simultaneously as the solution to the problem

$$\min_{\{\mathbf{c}_i\}} \sum_{i=1}^N \|\mathbf{s}_i - \mathbf{B}\mathbf{c}_i\|_2^2 + \sum_{j \in \mathcal{N}_i} \nu_{ij} \|\mathbf{c}_i - \mathbf{c}_j\|_2^2 \quad \text{s.t.} \quad \mathbf{CLP}\mathbf{c}_i \geq -4\pi\mathbf{1}, \quad (6)$$

where \mathcal{N}_i is a spatial neighborhood of the voxel i and ν_{ij} is a non-negative similarity factor between the data at voxels i and j . Once $\{\mathbf{c}_i\}$ are estimated, the ODF vectors $\{\mathbf{p}_i\}$ are computed from (5).

In this work, we will use two distinct feature vectors for segmenting the N ODFs: 1) the square-root parametrization [11] $\mathbf{p}^{\text{sq}} := \sqrt{\mathbf{p}}$ of the ODF vector and 2) its SH coefficient vector [12] $\mathbf{c}^{\text{sq}} := \mathbf{C}^\dagger \mathbf{p}^{\text{sq}}$, where \mathbf{C}^\dagger denotes the pseudoinverse of \mathbf{C} . Notice that both vectors lie on hyperspheres of different dimensions, i.e., $\mathbf{p}^{\text{sq}} \in \mathbb{S}^{M-1}$ and $\mathbf{c}^{\text{sq}} \in \mathbb{S}^{R-1}$. For simplicity, in what follows we will refer to both of them as ODFs and denote them by $\mathbf{q} \in \mathbb{S}^d$ with $d \in \{M-1, R-1\}$.

3. SEGMENTATION ON THE SPACE OF ODFs

In order to segment ODF fields into regions of distinct diffusion properties, the algorithm needs to handle multiple manifolds. Clustering intersecting manifolds is particularly important for analyzing crossing fiber bundles. In those cases, methods such as Riemannian Manifold Clustering [13] may not be directly applicable as the similarity between two voxels is defined by using nearest neighbors and two neighbors may lie in different submanifolds. The sparse subspace

clustering (SSC) method in [7], on the other hand, uses sparse representation to cluster data in multiple intersecting linear or affine subspaces. Therefore, we propose to extend the SSC from Euclidean data in multiple subspaces to ODFs in multiple submanifolds.

3.1. Subspace Clustering via Sparse Representation

Sparse representation (SR) theory is based on the idea that ‘‘many signals can have a concise representation when expressed in a proper basis’’. Specifically, it aims to find whether a vector $\mathbf{z} \in \mathbb{R}^D$ is K -sparse with respect to a dictionary $\{\psi_i\}_{i=1}^N$, i.e., if $\mathbf{z} = \sum_{i=1}^N w_i \psi_i = \Psi \mathbf{w}$ with $K \ll N$ nonzero coefficients $\{w_i\}$. It is shown in [14, 15] that under some conditions, one can recover the sparse representation \mathbf{w} of \mathbf{z} by solving the problem $\min \|\mathbf{w}\|_1$ subject to $\mathbf{z} = \Psi \mathbf{w}$.

The SSC algorithm [7] uses SR to cluster points drawn from a union of k linear or affine independent subspaces. In particular, each data point can be written as a linear or affine combination of all other points. By searching for the sparsest combination, the point $\mathbf{z}_i \in \mathbb{R}^D$ is represented in terms of points lying in the same subspace that it is drawn from. The SR denoted by $\mathbf{w}_i \in \mathbb{R}^{N-1}$ with entries $\{w_{ij}\}_{j \neq i}$ can be computed by solving the optimization problem

$$\min_{\{w_{ij}\}} \sum_{j \neq i} |w_{ij}| + \mu \left\| \mathbf{z}_i - \sum_{j \neq i} w_{ij} \mathbf{z}_j \right\|_2 \quad \text{s.t.} \quad \sum_{j \neq i} w_{ij} = 1, \quad (7)$$

where $\mu > 0$ is a constant trade-off parameter between the sparsity of the solution and the reconstruction error. The SRs of all data points are used to form the matrix of coefficients $\mathbf{W} = [w_{ij}] = [\bar{\mathbf{w}}_1, \dots, \bar{\mathbf{w}}_N] \in \mathbb{R}^{N \times N}$, where $\bar{\mathbf{w}}_i$ is the vector with a zero inserted at the i -th entry of the sparsest solution \mathbf{w}_i . Following the construction of a similarity matrix $\mathbf{A} = [a_{ij}]$, where $a_{ij} = |w_{ij}| + |w_{ji}|$, the segmentation of the data is obtained via spectral clustering [16], i.e., by applying k-means to the eigenvectors of the Laplacian matrix (computed from \mathbf{A}) corresponding to its k smallest eigenvalues.

3.2. Extending SSC to Multiple Submanifolds

Recall that SSC aims to represent a point \mathbf{z}_i as the sparsest *linear* or *affine* combination of the other data points. From a data reconstruction perspective, this representation can be considered as the *sparsest* solution of $\min \|\mathbf{z}_i - \tilde{\mathbf{z}}_i\|_2$, where $\tilde{\mathbf{z}}_i = \sum_{j \neq i} w_{ij} \mathbf{z}_j$. However, in the case of the ODF feature vectors $\mathbf{q} \in \{\mathbf{p}^{\text{sq}}, \mathbf{c}^{\text{sq}}\}$, one has to use the Riemannian properties of the manifold for minimizing the geodesic distance between an ODF \mathbf{q}_i and its reconstruction $\tilde{\mathbf{q}}_i$ from all other ODFs. We thereby use Riemannian interpolation to compute $\tilde{\mathbf{q}}_i$ as [13]

$$\tilde{\mathbf{q}}_i = \exp_{\mathbf{q}_i} \left(\sum_{j \neq i} w_{ij} \log_{\mathbf{q}_i}(\mathbf{q}_j) \right), \quad (8)$$

where $\log : \mathcal{M} \mapsto T_{\mathbf{q}}\mathcal{M}$ is the logarithm map from the manifold \mathcal{M} , i.e., the hypersphere \mathbb{S}^d , to its tangent space $T_{\mathbf{q}}\mathcal{M}$ at \mathbf{q} , and $\exp : T_{\mathbf{q}}\mathcal{M} \mapsto \mathcal{M}$ is the exponential map. Now, since the geodesic distance between \mathbf{q}_i and $\tilde{\mathbf{q}}_i$ is given by $\|\sum_{j \neq i} w_{ij} \log_{\mathbf{q}_i}(\mathbf{q}_j)\|_2$, the problem of writing the ODF \mathbf{q}_i as a sparse combination of ODFs $\{\mathbf{q}_j\}_{j \neq i}$ can be posed as the following convex optimization problem

$$\min_{\{w_{ij}\}} \sum_{j \neq i} |w_{ij}| + \mu \left\| \sum_{j \neq i} w_{ij} \log_{\mathbf{q}_i}(\mathbf{q}_j) \right\|_2 \quad \text{s.t.} \quad \sum_{j \neq i} w_{ij} = 1. \quad (9)$$

Here, it is worth noting that the logarithm map of \mathbf{q}_j at \mathbf{q}_i (from the hypersphere \mathbb{S}^d to its tangent space $T_{\mathbf{q}_i}\mathbb{S}^d$ at \mathbf{q}_i) has the following closed form expression

$$\log_{\mathbf{q}_i}(\mathbf{q}_j) = \frac{\mathbf{q}_j - (\mathbf{q}_i^\top \mathbf{q}_j) \mathbf{q}_i}{\|\mathbf{q}_j - (\mathbf{q}_i^\top \mathbf{q}_j) \mathbf{q}_i\|_2} \arccos(\mathbf{q}_i^\top \mathbf{q}_j). \quad (10)$$

More importantly, the problem in (9) is the direct extension of (7) from the Euclidean space to the Riemannian manifold.

One potential disadvantage of the method proposed so far is that it requires solving N optimization problems in $N-1$ variables. Therefore, solving the problem in (9) may be computationally expensive when the number of voxels is large. To address this issue, we find the sparse representation of an ODF in a neighborhood of the voxel i , denoted by Ω_i . The cardinality of Ω_i is critical in achieving the desired sparsity: A smaller value for $|\Omega_i|$ decreases computation time at the expense of a possible loss of sparsity, whereas a larger value favors sparsity, yet increases computation time. In our experiments, we choose $|\Omega_i|$ to be between 5 to 20 times the length of the feature vector \mathbf{q} . This results in a modified version of the optimization problem at each voxel, which is given by

$$\min_{\{w_{ij}\}} \sum_{j \in \Omega_i} |w_{ij}| + \mu \left\| \sum_{j \in \Omega_i} w_{ij} \log_{\mathbf{q}_i}(\mathbf{q}_j) \right\|_2 \quad \text{s.t.} \quad \sum_{j \in \Omega_i} w_{ij} = 1. \quad (11)$$

In our experiments, we set $\mu = 0.01$, as suggested in [7]. By solving the problem in (11) for all the ODFs in the field, we obtain the matrix of weights \mathbf{W} , which is further used for constructing the similarity matrix \mathbf{A} , as described in Section 3.1. The segmentation of the data is then found by applying spectral clustering to \mathbf{A} .

4. VALIDATION AND DISCUSSIONS

We compare our Sparse Riemannian Manifold Clustering (SRMC) algorithm to Normalized Cut (Ncut) [17] and to k-means (kM) on ODFs using both feature vectors $\{\mathbf{p}^{\text{sq}}, \mathbf{c}^{\text{sq}}\}$. For Ncut, we construct the affinity matrix with entries $w_{ij} = \exp(-\rho(\arccos(\mathbf{q}_i^\top \mathbf{q}_j)))$, where ρ normalizes the angular distances to $[0, 1]$. For k-means, we compute the dissimilarity between \mathbf{q}_i and \mathbf{q}_j as $1 - \mathbf{q}_i^\top \mathbf{q}_j$. We compare these algorithms on the segmentation of synthetic ODF fields corrupted with varying levels of noise. We also test our algorithm on the segmentation of fiber bundles in the diffusion MR data of a phantom and regions such as the corpus callosum and the cingulum in the brain.

Experiments on Synthetic Data. We first evaluate the performances of SRMC, Ncut and k-means at varying levels of noise. For this purpose, we generate synthetic data using the multi-tensor model in [9], where the HARDI signals $\{S(\theta_n, \phi_n)\}$ at 162 gradient directions, with $S_0 = 1$ and $b = 3000$ s/mm², are simulated to represent an isotropic background and ODFs of 1, 2, or 3 fibers. Noisy HARDI signals are generated by adding complex Gaussian noise with zero mean and standard deviation $\sigma = \frac{S_0}{\zeta}$, where ζ is the signal-to-noise ratio (SNR). The ODFs are reconstructed as described in Section 2.2 and three synthetic ODF fields at different levels of complexity, i.e., intersecting linear fibers (Fig. 1(a)), a circular fiber (Fig. 1(b)), and intersecting curved fibers (Fig. 1(c)), are generated. We set the number of clusters to $k = (5, 2, 4)$ for these data, respectively, and obtain the segmentation. Since the ground truth segmentation is known, we quantify the performances of the methods by using the ratio of the number of correctly segmented voxels to the total number of voxels.

Fig. 1(d) shows the average segmentation accuracy (%) over these ODF fields. We observe that, for SRMC and Ncut, the square-root representation \mathbf{p}^{sq} slightly outperforms the SH coefficient vector \mathbf{c}^{sq} when $\text{SNR} \geq 5$. In addition, SRMC gives more accurate segmentations than Ncut for all SNRs and outperforms k-means for all SNRs except when $\text{SNR} = 5$. However, all three methods fail to provide acceptable accuracy rates at $\text{SNR} = 3$. Notice also that SRMC with \mathbf{p}^{sq} features achieves accuracies of around 94% and 84% at SNR levels of 10 and 5, respectively. This is a very promising result because 1) these levels of SNR are in the range of image noise often observed

in real dMRI data and 2) the final segmentation obtained by SRMC does not enforce spatial regularization.

Experiments on Phantom Data. We conduct a second set of experiments on the dMRI data of the phantom¹ shown at the top left portion of Fig. 2(a). The HARDI signals are acquired with a 64×64 image matrix, an isotropic spatial resolution of 3 mm, and a diffusion sensitization at $b = 1500$ s/mm² applied along a set of 64 gradient directions. The ODFs are reconstructed as described in Section 2.2 and two regions of interest (ROIs) shown in Fig. 2(a) are selected for segmentation. Following the quantitative results in the synthetic experiments, we select the square-root parametrization \mathbf{p}^{sq} as a feature and set the number of clusters to $k = (3, 4)$ for the two ROIs, respectively. As shown in Fig. 2(b), where different clusters are represented with distinct colors, SRMC accurately segments crossing and fanning fiber bundles. In addition, in the first ROI, the region of intersection could be identified as a separate cluster if the number of clusters is $k \geq 4$. Recall also that since SRMC does not enforce spatial regularity, the output might look “noisy” as in the case of the second ROI. Our method is, however, automatic and its output could be used to initialize energy minimization-type segmentation schemes.

Experiments on Real Data. We finally test our algorithm in the segmentation of the corpus callosum (CC) and the cingulum (CG) from a brain dMRI dataset. The CC is the major communication conduit linking the two hemispheres of the brain and is responsible for distinct cognitive processes, whereas the CG bundle serves to connect the cingulate cortex with other regions and its integrity plays a critical role in the maintenance of the processing of cognitive functions. Reduced fractional anisotropy in these cerebral regions is often linked with atrophy and neurodegeneration.

For this experiment, we use a HARDI human brain dataset of 40 subjects [8] where 105 images were acquired, 11 with no diffusion weighting and 94 with diffusion weighting at $b = 1159$ s/mm², by using a 128×128 acquisition matrix (1.8 mm in-plane resolution) and 55 axial slices (2 mm thick). Specifically, we use the mean ODFs computed as described in [11] and focus on segmenting the aforementioned regions from two consecutive sagittal slices corresponding to ROIs on the left hemisphere and the central sagittal plane (Fig. 2(c)). Notice that in the first ROI, both the CC and the CG are visible, whereas in the second ROI, only the CC can be identified. We choose the square-root parametrization \mathbf{p}^{sq} as features and set the number of clusters to $k = (4, 3)$ for the two ROIs, respectively. Fig. 2(c) also shows the results of the proposed SRMC method. In the first slice, we observe that the corpus callosum and the cingulum are correctly segmented from the background. In fact, we oversegmented the data by setting $k = 4$ to reduce the effect of the fornix on the segmentation of the cingulum. In the second slice, the corpus callosum is successfully segmented as well in the absence of the cingulum.

5. CONCLUSIONS

We have presented a method called Sparse Riemannian Manifold Clustering (SRMC) to segment HARDI data described by ODFs into multiple clusters. SRMC integrates tools from sparse representation theory into spectral clustering and exploits the Riemannian properties of the space of ODFs to reformulate the problem of computing sparse representations for ODFs. We observed that SRMC gives more accurate segmentations than the natural extensions of normalized cut and k-means to ODFs at varying levels of noise. SRMC also produces accurate boundaries between cerebral regions with different diffusion characteristics. These boundaries could be further improved

¹<http://www.lnao.fr/spip.php?article107>

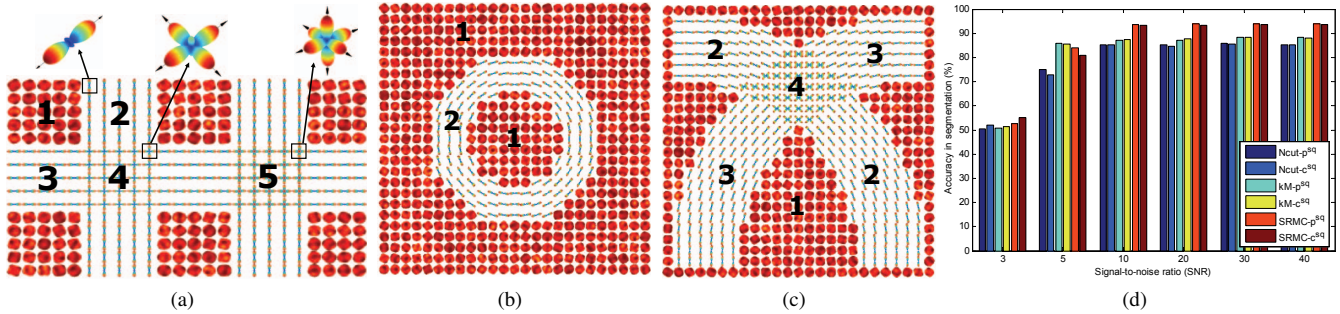


Fig. 1. Synthetic ODF fields with (a) three in-plane and one out-of-plane fibers, generating five distinct diffusion profiles, i.e., five clusters, (b) a circular fiber (two clusters), (c) two intersecting curved fibers (four clusters). (d) Segmentation accuracy (%) at varying levels of noise.




Fig. 2. (a) dMRI data of the phantom and its ODFs with two ROIs, (b) Segmentation of distinct fiber bundles in those ROIs, (c) ROIs in two consecutive sagittal slices of a brain ODF data and the segmentation of the corpus callosum (cc) and the cingulum (cg).

by adding spatial regularization. By restricting fiber extraction and tracking to be within these boundaries, we aim to improve the quality of white matter atlases of the brain. However, our future work will initially focus on further reducing the computational load in solving the problem in (11) by using homotopy techniques [18]. We will also allow the user to interact with the method by marking scribbles with known class labels so as to correct erroneous segmentations.

6. REFERENCES

- [1] D.S. Tuch, “Q-ball imaging,” *Magnetic Resonance in Medicine*, vol. 52, no. 6, pp. 1358–1372, 2004.
- [2] M. Descoteaux and R. Deriche, “High angular resolution diffusion MRI segmentation using region-based statistical surface evolution,” *Jour. of Mathematical Imaging and Vision*, vol. 33, no. 2, pp. 239–252, 2009.
- [3] P. Hagmann, L. Jonasson, T. Deffieux, R. Meuli, J.-P. Thiran, and V.J. Wedeen, “Fibertract segmentation in position orientation space from high angular resolution diffusion MRI,” *NeuroImage*, vol. 32, no. 2, pp. 665–675, 2006.
- [4] L. Jonasson, X. Bresson, J.-P. Thiran, V.J. Wedeen, and P. Hagmann, “Representing diffusion MRI in 5-D simplifies regularization and segmentation of white matter tracts,” *IEEE Trans. on Medical Imaging*, vol. 26, no. 11, pp. 1547–1554, 2007.
- [5] T. McGraw, B. Vemuri, R. Yeziarski, and T. Mareci, “Segmentation of high angular resolution diffusion MRI modeled as a field of von Mises-Fisher mixtures,” in *ECCV*, 2006, pp. 463–475.
- [6] D. Wassermann, M. Descoteaux, and R. Deriche, “Diffusion maps clustering for magnetic resonance Q-ball imaging segmentation,” *Int. Jour. of Biomedical Imaging*, 2008.
- [7] E. Elhamifar and R. Vidal, “Sparse subspace clustering,” in *IEEE Conf. on CVPR*, 2009, pp. 2790–2797.
- [8] A. Goh, C. Lenglet, P.M. Thompson, and R. Vidal, “Estimating orientation distribution functions with probability density constraints and spatial regularity,” in *MICCAI*, 2009, pp. 877–885.
- [9] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, “Regularized, fast and robust analytical Q-ball imaging,” *Magnetic Resonance in Medicine*, vol. 58, pp. 497–510, 2007.
- [10] I. Aganj, C. Lenglet, and G. Sapiro, “ODF reconstruction in Q-ball imaging with solid angle consideration,” in *IEEE Int. Symp. on Biomedical Imaging*, 2009, pp. 1398–1401.
- [11] A. Goh, C. Lenglet, P.M. Thompson, and R. Vidal, “A nonparametric Riemannian framework for processing High Angular Resolution Diffusion Images (HARDI),” in *IEEE Conf. on CVPR*, 2009, pp. 2496–2503.
- [12] J. Cheng, A. Ghosh, T. Jiang, and R. Deriche, “A Riemannian framework for orientation distribution function computing,” in *MICCAI*, 2009, pp. 911–918.
- [13] A. Goh and R. Vidal, “Clustering and dimensionality reduction on Riemannian manifolds,” in *IEEE Conf. on CVPR*, 2008, pp. 1–7.
- [14] D.L. Donoho, “For most large underdetermined systems of linear equations the minimal ℓ^1 -norm solution is also the sparsest solution,” *Comm. on Pure and Applied Math.*, vol. 59, no. 6, pp. 797–829, Jun. 2006.
- [15] E. Candès and M. Wakin, “An introduction to compressive sampling,” *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, Mar. 2008.
- [16] U. von Luxburg, “A tutorial on spectral clustering,” *Statistics and Computing*, vol. 17, 2007.
- [17] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.
- [18] M.S. Asif and J. Romberg, “Dynamic updating for ℓ_1 minimization,” *IEEE Jour. of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 421–434, 2010.