

White Matter Tract Clustering and Correspondence in Populations

Lauren O'Donnell^{1,2} and Carl-Fredrik Westin^{1,3}

¹ MIT Computer Science and Artificial Intelligence Lab

² Harvard-MIT Division of Health Sciences and Technology, Cambridge MA, USA
lauren@csail.mit.edu

³ Laboratory of Mathematics in Imaging,
Brigham and Women's Hospital, Harvard Medical School, Boston MA, USA
westin@bwh.harvard.edu

Abstract. We present a novel method for finding white matter fiber correspondences and clusters across a population of brains. Our input is a collection of paths from tractography in every brain. Using spectral methods we embed each path as a vector in a high dimensional space. We create the embedding space so that it is common across all brains, consequently similar paths in all brains will map to points near each other in the space. By performing clustering in this space we are able to find matching fiber tract clusters in all brains. In addition, we automatically obtain correspondence of tractographic paths across brains: by selecting one or several paths of interest in one brain, the most similar paths in all brains are obtained as the nearest points in the high-dimensional space.

1 Introduction

Diffusion MRI measures water diffusion in tissue. Where many cells have similar orientations, such as in muscle or nervous tissue, the MRI measurements describe that local cellular pattern. Starting from these measurements of diffusion (represented as diffusion tensors) in the brains of several subjects, we produce an approximation of the neural fiber tract anatomy of the population.

Due to the complexity of the data, manual exploration of diffusion tensor images for the purposes of segmentation, registration, diagnosis, surgical planning, etc. is currently very difficult. One may choose to visualize tensors directly using small glyphs, or attempt to reconstruct anatomical structure by interpolating paths that follow the principal direction of diffusion. Detailed three-dimensional neuroanatomical knowledge is necessary to pinpoint the expected location of a white matter fiber, so it is not easy to choose a region of interest for display of the glyphs or paths. By creating an automatic segmentation of all possible paths, one could reduce this data interaction problem to a simple choice of which group of paths to display. Not only would this aid in visualization of the data, but also in the automatic quantification of properties of interest such as anisotropy measures, and in finding cross-patient anatomical correspondences.

Early work in grouping of tractographic paths by Brun et. al has used color to enhance visual perception of connectivity using information from spectral

embeddings [2, 1]. Related work in segmentation of tractographic paths has employed clustering approaches which divide the problem into two parts: the decision of how to quantify distance (or similarity) between paths, and the choice of clustering method. Several distance measures have been proposed in the literature. In one of the earliest approaches, Ding et al. calculate the mean distance separating paths using pointwise correspondences between path segments [4]. Their method is specific to paths which have been seeded in one image slice, but in later approaches more general distances have been defined. Brun et al. introduced a 9-D vector for tract shape approximation which they define as the mean and lower triangular part of the covariance matrix of the points on a path [1]. They compute the distance between paths as the Euclidean distance between the corresponding 9-D vectors. Gerig et al. and Corouge et al. propose distances that do pointwise comparison of tract shapes: they define three measures related to the Hausdorff distance [6, 3]. In more recent work by Jonasson et al. (who use paths through high angular resolution diffusion data) a path similarity measure is calculated based on the number of times two paths share the same voxel [7]. Two general types of clustering methods have been employed in the literature, hierarchical clustering [3, 6, 10], and the spectral clustering approach [1, 7].

To our knowledge, there is little prior work on the topic of matching tractographic paths across patients. In work by Zhang et al., a two-step process first performs clustering on each patient, then describes the clusters with 9-D vectors (the average start point, end point, and “middle point” of all paths in the cluster). These feature vectors are used to match clusters across patients [10].

2 Methods

Our method has several steps: tractography, estimation of distances/similarities between paths, spectral embedding, and finally clustering and/or anatomical region of interest (ROI) selection in the population. Figure 1 illustrates this process using data from the population of brains described further in Section 3. In the rest of this section we explain each step in our algorithm and its parameters.




Fig. 1. Visual overview of the method using data from the population: two paths, a random sample of 500 embedding vectors, and population tract clusters in one brain

2.1 Tractography

We perform tractography using open-source software (www.slicer.org), which does second-order Runge-Kutta integration. We produce paths starting at all

points in a region of interest that covers most of the brain. Paths stop when the FA (fractional anisotropy) or curvature become too low or high, respectively. For the experiments in this paper, the region of interest was the area with FA greater than 0.3, and the settings used for the thresholds were FA of 0.1 and curvature of 1.1. In addition, to limit the number of paths, paths shorter than a length cutoff of approximately 30mm were discarded. To give an idea of the numbers of paths involved, there are about $256 \times 256 \times 50$, or 3 million, voxels in a dataset, of which we select approximately 300,000 for path seeding. After pruning the number of paths is between 5,000 and 25,000 for one brain.

2.2 Similarity Measure for Tract Paths

We use a symmetrized Hausdorff distance, as suggested in [6], and we compare paths in all brains. The Hausdorff distance H_{ij} is defined as the maximum of pointwise minimum distances between a pair of curves i and j . We symmetrize the distance by averaging H_{ij} with H_{ji} . In the common case where a short and long path are aligned, the average distance is a more forgiving measure than the maximum distance, but it preserves some information about the length mismatch, unlike the minimum distance. We convert the distance measure to a similarity measure W_{ij} using a Gaussian kernel where σ controls the distance over which paths can influence each other. For the experiments in this paper, σ was set to 10.

$$W_{ij} = e^{-H_{ij}/\sigma^2} \quad (1)$$

2.3 Spectral Clustering

Spectral clustering methods group data using eigenvectors of a data affinity matrix. The method we employ is described in [8] and is identical to the k-way normalized cuts procedure from [5], except we do not need to postprocess the clusters to correct for oversegmentation. First the symmetric tract similarity matrix W is scaled using a diagonal matrix D whose entries are the row (or column) sums of W . This produces L :

$$L = D^{-\frac{1}{2}} W D^{-\frac{1}{2}} \quad (2)$$

Embedding vectors are then calculated from the eigensystem of L ($L = U\Lambda U^T$) by applying a scaling to the rows of U . The scaling converts the eigenvectors to the solution one would get by solving the generalized eigensystem $(D - W)y = D\lambda y$, which minimizes the Normalized Cut [9]. From [5], the embedding vector for the j th tract is

$$E_j = \frac{1}{\sqrt{D_{jj}}}(U_{j,2}, U_{j,3}, \dots, U_{j,n}) \quad (3)$$

where n is the number of eigenvectors chosen for the embedding, and the column index of U starts at 2 in order to skip the first (major and smoothest) eigenvector. Care is needed in choosing n because the eigenvectors become noisy as the

eigenvalues decrease. For the experiments in this paper we have set n to 20. In spectral clustering, the next step is to cluster the embedding vectors, usually by k-means.

2.4 The Nystrom Method

In our application it is important to avoid computing the full W matrix, which for 5 brains would be at least $25,000 \times 25,000$. The Nystrom method is a way to interpolate values of a function, using known values and interpolation weights. We use this method to approximate the eigenvectors of W , as described in [5].

Instead of computing W directly, a random subset of paths is selected for comparison, and the similarity measure is calculated between each path and that subset. The matrix A contains similarities between paths in the subset, and the matrix B contains similarities between other paths and the selected subset. From [5], the approximate eigenvectors of W are calculated as

$$\bar{U} = \begin{bmatrix} U \\ B^T U \Lambda^{-1} \end{bmatrix} \quad (4)$$

where $U = A U \Lambda^{-1}$ and the ordering of the rows of \bar{U} is such that those corresponding to rows of A are first, followed by those corresponding to rows from B^T . To get an intuition for why this works, imagine that you have performed clustering using the information in A and you would like to add some new data to the problem without starting over. Since $A = A^T$ we can rewrite equation 4 as

$$\bar{U} = \begin{bmatrix} A^T U \Lambda^{-1} \\ B^T U \Lambda^{-1} \end{bmatrix} \quad (5)$$

So this process can be thought of as estimating the rows of \bar{U} by projecting the new data from B into the same space where the old data in A were projected.

There are two more details when using this approximation method. First of all, the normalization of A and B using row sums of W must be performed somehow before calculating \bar{U} . In [5] this is shown to be possible without computing W itself. The second detail is that the estimated eigenvectors of W (the columns of \bar{U}) should be orthonormal. In [5] the authors give two methods for orthogonalizing the columns. However we have found this to be unnecessary because in practice, after the columns of \bar{U} are normalized to length one, we have always found $\bar{U}^T \bar{U}$ to be close to the identity matrix.

3 Results

We employed our method to cluster tractographic paths in a population of 5 brains. For each, the input to tractography was a DTI scan with voxel size $0.86 \times 0.86 \times 5$ mm. The total number of tract paths from all brains was 36,003. We randomly selected 1000 tract paths from the population, and compared all other paths to these in order to generate A , B , and \bar{U} . We then performed spectral




Fig. 2. Anterior, right, and superior views of the five brains in the population. The colors show correspondence via spectral embedding.

clustering to generate 100 clusters, a number that we find in practice is large enough to avoid grouping very dissimilar tracts. Interesting anatomical clusters exist at many size scales so choosing the number of clusters is not easy.

The output clustered/embedded brains are shown in Figure 2. The correspondence of the colors directly shows the correspondence in embedding space.




Fig. 3. Anatomical correspondences: selected clusters, displayed in all 5 brains. The two leftmost images show the corpus callosum viewed superiorly and from the right. Of the 100 clusters found, 10 were manually chosen as belonging to the corpus callosum. The third images from the left show a single cluster containing the cingulum bundles, viewed superiorly. Finally, the rightmost images show the two clusters that contain the left and right uncinate fasciculi, viewed anteriorly and from the right.

To generate this visualization, the paths in each cluster were colored based on the centroid of that cluster in the embedding space. Mapping the centroid vectors to colors was performed as described in [1], where the first three components of the embedding vector are used to define red, green, and blue.

In Figure 3 we show more detailed anatomical correspondences. Using the same 100 clusters, one of the brains in the population was visualized in order to identify clusters of interest. The cluster indices aren't in themselves significant, but clusters 3, 8, 12, 14, 19, 33, 54, 66, 77, and 81 were found to correspond to the corpus callosum, cluster 31 happened to correspond to both cingulum bundles, and clusters 29 and 42 represented the uncinate fasciculi. After selecting these clusters in one brain, the same clusters were then displayed in all brains to demonstrate correspondence.

4 Discussion

The first issue with the presented method is that it depends on paths from tractography, a process prone to errors due to noise and partial voluming, with limited validity in regions of low tensor anisotropy. The alternative approach of working directly in voxel space by defining a suitable voxel-to-voxel similarity measure is possible, however capturing long range connectivity may be more difficult. Also in principle our method could be applied to paths produced by another algorithm or paths through data with another diffusion model.

Another issue of interest is that because the scan protocol was chosen to acquire the images in a consistent way, in this work we explored performing clustering without rigid alignment. This can be seen especially in the fourth brain from the top in Figure 2, where the brain stem is angled very differently from the rest of the brains. This may be the reason for the less-structured uncinate fasciculus clusters for that brain in Figure 3. The robustness of the method is shown by the overall success of the correspondence, however we believe correspondence in the population will improve with rigid and potentially nonrigid alignment and perhaps normalization for brain size.

A question that might be posed about our choice of method is whether we could instead use a sparse W and sparse eigensolvers. Our W matrix is not particularly sparse, unlike the case of a W matrix that comes from similarities between pixels, because neighborhood relationships between paths are not limited by a rectangular grid. In addition, in [5] it is shown that the Nystrom approximation has comparable performance and much better running time when compared to a sparse solver.

It also makes sense to ask how this method will scale to a large population of brains. Our current implementation would likely scale up to a small clinical study, but for population sizes in the 100's a more memory-careful implementation would be necessary. One difficulty is holding all embedding vectors in memory to perform k-means clustering. But if a subset of the embedding vectors is chosen for clustering, and the subset is representative enough of the population, possibly the cluster centroids would be close to those obtained by doing the full clustering.

Another potential extension is the incorporation of additional information into the creation of the affinity matrix. For example if there already exist segmentations of the brains in question, the affinity should increase for two paths that pass through similar segmented anatomy. Another important point about the affinity matrix is that the choice of Hausdorff distance is somewhat arbitrary (though reasonable) and we hope to be able to compare various distance measures now that we have a working framework for clustering studies.

5 Conclusion

We have presented a novel method for obtaining anatomical clusters and correspondences across brains using paths from tractography through DTI data. To our knowledge, this is the first method for obtaining dense correspondences (path-to-path correspondences), as well as cluster-to-cluster correspondences across a population of brains. We believe there are many interesting future applications for this method, including studies of anatomy in populations and DTI visualization using automatic segmentation and labeling of tract clusters.

Acknowledgments

We would like to thank the HST Neuroimaging Training Grant, NIH P41-RR13218 (NAC), NIH 1-R01-NS051826-01, and NIH NCCR Morphometry Biomedical Informatics Research Network (BIRN) grant U24 RR021382. This work is part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149.

References

1. A. Brun, H. Knutsson, H.-J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber traces using normalized cuts. In *MICCAI*, pages 368–375, 2004.
2. A. Brun, H.-J. Park, H. Knutsson, and C.-F. Westin. Coloring of DT-MRI fiber traces using Laplacian eigenmaps. In *EUROCAST*, pages 564–572, 2003.
3. I. Corouge, S. Gouttard, and G. Gerig. Towards a shape model of white matter fiber bundles using diffusion tensor MRI. In *ISBI*, pages 344–347, 2004.
4. Z. Ding, J. C. Gore, and A. W. Anderson. Classification and quantification of neuronal fiber pathways using diffusion tensor MRI. *MRM*, 49:716–721, 2003.
5. C. Fowlkes, S. Belongie, F. Chung, and J. Malik. Spectral grouping using the Nystrom method. *PAMI*, 26(2):214–225, 2004.
6. G. Gerig, S. Gouttard, and I. Corouge. Analysis of brain white matter via fiber tract modeling. In *EMBS*, page 426, 2004.
7. L. Jonasson, P. Hagmann, J.-P. Thiran, and V. J. Wedeen. Fiber tracts of high angular resolution diffusion MRI are easily segmented with spectral clustering. In *ISMRM*, 2005.
8. A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *NIPS*, volume 14, 2001.
9. J. Shi and J. Malik. Normalized cuts and image segmentation. *PAMI*, 22(8):888–905, 2000.
10. S. Zhang and D. H. Laidlaw. DTI fiber clustering and cross-subject cluster analysis. In *ISMRM*, 2005.