

Shape Modelling for Tract Selection

Jonathan D. Clayden, Martin D. King, and Chris A. Clark

Institute of Child Health, University College London, UK
E-mail: j.clayden@ucl.ac.uk

Abstract. Probabilistic tractography provides estimates of the probability of a structural connection between points or regions in a brain volume, based on information from diffusion MRI. The ability to estimate the uncertainty associated with reconstructed pathways is valuable, but noise in the image data leads to premature termination or erroneous trajectories in sampled streamlines. In this work we describe automated methods, based on a probabilistic model of tract shape variability between individuals, which can be applied to select seed points in order to maximise consistency in tract segmentation; and to discard streamlines which are unlikely to belong to the tract of interest. Our method is shown to ameliorate false positives and remove the widely observed falloff in connection probability with distance from the seed region due to noise, two important problems in the tractography literature. Moreover, the need to apply an arbitrary threshold to connection probability maps is entirely obviated by our approach, thus removing a significant user-specified parameter from the tractography pipeline.

1 Introduction

Probabilistic tractography uses diffusion MRI (dMRI) data to provide estimates of the probability of a connection existing between a seed point, or seed region, and all other points within a brain volume. When the seed region is placed within a white matter tract, areas of high probability are typically found within other sections of the same tract. The first step towards estimating these probabilities of connection is to derive an orientation distribution function (ODF) for each voxel in the brain, which characterises the orientations of local structure. Several alternative methods for calculating such an ODF have been described [1], some of which are based on a specific model of diffusion, while others take a model-free approach. Probabilistic streamlines are then generated by alternating between sampling from these ODFs and stepping along the sampled direction. The probability of connection between the seed region and any other voxel is then estimated as the proportion of these streamlines that visit the target voxel.

Unfortunately, the probabilities of connection estimated by this Monte Carlo method are strongly affected by nuisance characteristics of the basic data, particularly noise, as well as limitations of the applied diffusion model. Streamlines may be deflected away from the tract of interest or prematurely truncated due to the nearby presence of a disparate tract, or due to ambiguity in the estimated

ODFs, or because of noise—and the estimated probability of connection at a given voxel may be affected in turn by any or all of these. The most common method of compensation for these effects is to threshold the visitation map to avoid including pathways which are unlikely to belong to the tract of interest. But this approach is very sensitive to an arbitrary user-specified parameter, the threshold level; and relies on the flawed assumption that false positive pathways are nonrepeatable and spatially dispersed. Moreover, it cannot correct for effects which lead to *underestimation* of the probability of connection, such as the premature termination of streamlines. A method to compensate for the latter has been proposed by Morris *et al.* [2], which uses a “null connection map” to differentiate true connections from chance events, but a threshold is still required, and the technique cannot compensate for the effects of neighbouring pathways.

Seed points or regions may be placed by an observer, or transferred from a reference tract or atlas by registering dMRI data to a standard brain template. In either case, seed regions typically have no special anatomical significance, but are instead located to maximise the chance of reconstructing the tract of interest as fully as possible. Unfortunately, direct transfer of a seed region from an atlas space to diffusion space is generally not a reliable basis for consistent tract reconstruction, although recent work by Clayden *et al.* [3] described how a probabilistic model of tract shape variability can be used to select one or more suitable seed points from within such a region. A related approach was applied to the clustering of deterministic streamlines by Maddah *et al.* [4], whereby a tract trajectory model was used to infer cluster membership.

In this work we describe how the shape modelling approach can be applied not just to the choice of seed points, but also to the selection of streamlines which accurately represent a tract of interest. Using a reference tract for prior information, but also allowing for the topological variability of a given tract between individuals, we demonstrate dramatic improvements in patterns of estimated connectivity, without the need for a user-defined threshold to be applied.

2 Methods

The tract shape model used for this work is based on that described in [3]. Streamlines are represented by uniform cubic B-splines. The knot separation distance is fixed for each tract of interest, but is invariably larger than the typical width of an image voxel, so that small scale directional perturbations are of less importance than the large scale topology of the tract.

Given a set of seed points for a dMRI data set, indexed by i , each of which generates a set of streamlines, indexed by j , a single B-spline is initially fitted to the spatial median of the streamline set. The data, \mathbf{m}^i , which are relevant to the model then consist of the lengths of this B-spline either side of the seed point— L_1^i and L_2^i —and the angles, ϕ_u^i , between the straight lines connecting successive knot points, and the corresponding lines in the reference tract. The B-spline is transformed into MNI standard space for the purpose of calculating

these lengths and angles only. The location index u is, by convention, negative on one side of the seed point and positive on the other side.

Given the observed data, \mathbf{m}^i , for tract i , the model likelihood is given by

$$P(\mathbf{m}^i | \Theta) = P(L_1^i | L_1^*, \mathbf{L}_1) P(L_2^i | L_2^*, \mathbf{L}_2) \prod_{u=1}^{\tilde{L}_1^i} P(\phi_{-u}^i | \alpha_u) \prod_{u=1}^{\tilde{L}_2^i} P(\phi_u^i | \alpha_u) \quad (1)$$

where L_1^* and L_2^* are the lengths of the reference tract corresponding to L_1^i and L_2^i respectively; $\tilde{L}_1^i = \min\{L_1^i, L_1^*\}$ and equivalently for \tilde{L}_2^i ; and $\Theta = \{\mathbf{L}_1, \mathbf{L}_2, (\alpha_u)\}$ is a set of model parameters. The distributions over each variable are given by

$$\begin{aligned} L_1^i | L_1^* &\sim \text{Multinomial}(\mathbf{L}_1) \\ L_2^i | L_2^* &\sim \text{Multinomial}(\mathbf{L}_2) \\ \frac{\cos \phi_u^i + 1}{2} &\sim \text{Beta}(\alpha_u, 1) . \end{aligned} \quad (2)$$

The model parameters are fitted using an Expectation–Maximisation (EM) algorithm, the E-step of which calculates a posterior probability of each tract representing the best match to the reference tract [5]. All tracts are assumed to be *a priori* equiprobable matches. We use the implementation of this algorithm provided by the TractoR software package (<http://code.google.com/p/tractor/>). For the M-step we apply the hyperprior $\alpha_u - 1 \sim \text{Exponential}(\lambda)$, thereby constraining each α_u to ensure that smaller deviations from the reference tract are always more likely (i.e. $\alpha_u \geq 1$), and simultaneously avoiding model overfitting for small data sets. We take $\lambda = 1$ throughout this work.

The fitted model and posterior matching probabilities enable us to select one or more seed points which produce sets of probabilistic streamlines whose medians are accurate representations of the tract of interest for that subject. However, some individual streamlines may not resemble the median, and therefore may not accurately represent the tract of interest. To establish this, we additionally apply the modelling framework described above to streamline selection.

In this streamline-level selection phase, we begin by fitting a B-spline to each streamline, j , individually, and recovering a data vector, \mathbf{d}^{ij} , which describes it. This data vector is analogous to \mathbf{m}^i for the median. Treating (1) as a function of the data, with the model parameters fixed to those estimated by the EM algorithm, denoted $\hat{\Theta}$, we calculate the probability of each streamline under the model, which in turn allows us to derive the value

$$\pi^{ij} = \frac{P(\mathbf{d}^{ij} | \hat{\Theta})}{P(\mathbf{m}^i | \hat{\Theta})} . \quad (3)$$

We then retain streamlines probabilistically, such that

$$\text{Pr}(\text{keep streamline } j) = \min\{\pi^{ij}, 1\} . \quad (4)$$

Hence, streamline j will be retained with certainty if it has higher probability under the model than the median line itself; otherwise it may be kept if it has not

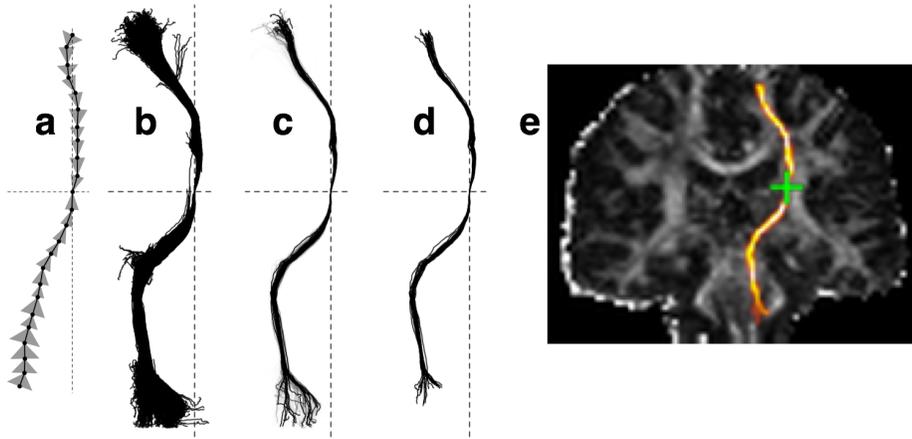


Fig. 1. Stages of the pruning process, shown in coronal projection. The knots of the reference tract are shown with 95% confidence intervals on the orientations of each tract segment (a). We also show a full set of 5000 probabilistic streamlines at full opacity (b), and with the alpha level for streamline j given by π^{ij} (c). The rejection algorithm is applied to the set, and remaining streamlines are then truncated to the length of the reference tract (d). A visitation map is finally calculated from this subset (e).

much lower probability. Since heavily truncated streamlines—and those following paths that differ substantially from the reference tract—will be associated with much lower values of π^{ij} , the contributions of such paths to estimates of the probability of connection will be annulled. The final step of our algorithm is to truncate all remaining streamlines to the length of the reference tract in the portions distal to the seed point. This is necessary for consistent results because the reference tract provides no orientational information in these regions, and so inappropriate trajectories have no effect on the value of (3).

This process of streamline pruning is illustrated by Fig. 1 for the left pyramidal tract. The model fitted from the median streamlines embodies the variability in tract topology across the whole data set. The amount of deviation “allowed” by the model over each segment of the reference tract is shown in Fig. 1a, in terms of the 95% confidence intervals on the angular deviation from the reference tract, which is controlled by the α_u parameters in the model. It can be seen that these confidence intervals tend to be wider towards the ends of the tract, particularly at the inferior extreme, due to greater uncertainty or variability in this region of the structure. Mapping the level of transparency to the value of π^{ij} for each streamline in the visualisation makes the effect of the method clear (Fig. 1c): some spread in the trajectories can be observed at the inferior extreme, in line with the greater local uncertainty in the model, but other branching structures are no longer visible. The probabilistic streamline retention algorithm is applied, the streamlines are truncated to the length of the reference, and a visitation map is produced (Fig. 1e).

3 Experiments and Results

Eight young, healthy right-handed volunteers (four male; mean age 31.9 ± 5.3 yr) underwent a dMRI protocol on three separate occasions. Scans were performed on a GE Signa LX 1.5 T clinical scanner using 64 noncollinear diffusion directions at a b value of 1000 s mm^{-2} , and 7 $b = 0$ images. Reconstructed voxel dimensions were $2 \times 2 \times 2 \text{ mm}$. ODFs were calculated using the Markov chain Monte Carlo method of Behrens *et al.*, and all tractography was performed using the “ProbTrack” algorithm described by the same authors [6]. Reference tracts were created using a published white matter atlas [7], as described in [8].

For each dMRI data set, initial seed points for each tract of interest were placed by transferring reference seeds from MNI standard space to diffusion space, using the FLIRT linear registration algorithm [9]. A neighbourhood of $7 \times 7 \times 7$ voxels (volume 2744 mm^3) centred at this point was then used as a source of seed points for the modelling process. However, seed voxels with a fractional anisotropy (FA) of less than 0.2 were excluded to save time, since such voxels are very likely to be outside white matter. Throughout our experiments, a single seed point from this neighbourhood was retained by the seed-level selection phase for simplicity—although our approach generalises to multiple seed points without modification.

In order to investigate the effects of the streamline-level selection which is the main novelty in this work, we begin by examining the lengths of streamlines retained. Histograms of the streamline lengths in Fig. 1, before and after pruning, are shown in Fig. 2. It is immediately evident from this figure that there is far greater homogeneity in streamline length after the pruning algorithm has been applied. Inferior to the seed point, in particular, the bimodal distribution seen before pruning—due to a short and erroneous pathway followed by a plurality of probabilistic streamlines—is completely absent after pruning.

A significant effect of discarding prematurely truncated streamlines is the removal of the usual dependence of visitation count on distance from the seed point. Fig. 3 shows that while applying a 1% threshold to visitation maps can remove most—though, in this case, not all—false positive pathways, visitation counts are conspicuously reduced at the superior and inferior extremes of the tract. After applying the pruning algorithm this issue disappears (Fig. 3c).

The effect on diffusion tensor parameters of applying each of the three treatments is shown graphically in Fig. 4. It is apparent that the three different treatments produce substantially different patterns of FA and MD across the data set, with greatest dispersion for the untreated case corresponding to Fig. 3a. To further quantify the effects on these widely used parameters, we used a simple random effects model to estimate their group means and variance components under each treatment. Treating equivalent tracts in the two hemispheres as repeated measurements, indexed by n , we model the measurement of FA or mean diffusivity (MD) in the m th scan of the l th subject with $f_{lmn} = \mu + \Delta_l + \delta_{lm} + \varepsilon_{lmn}$, where

$$\Delta_l \sim N(0, \sigma_b^2) \quad \delta_{lm} \sim N(0, \sigma_w^2) \quad \varepsilon_{lmn} \sim N(0, \sigma_e^2). \quad (5)$$

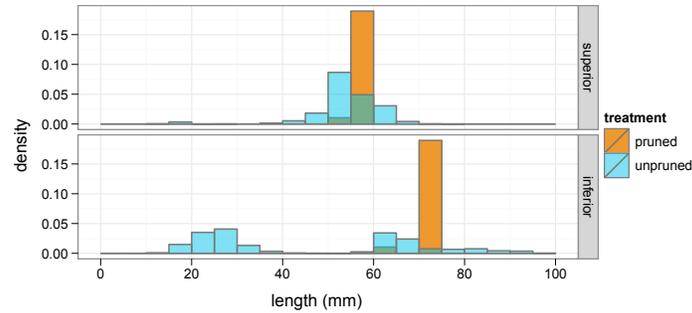


Fig. 2. Histograms showing the lengths of the pruned and unpruned streamline sets from Fig. 1, on the superior (top) and inferior (bottom) sides of the seed point.

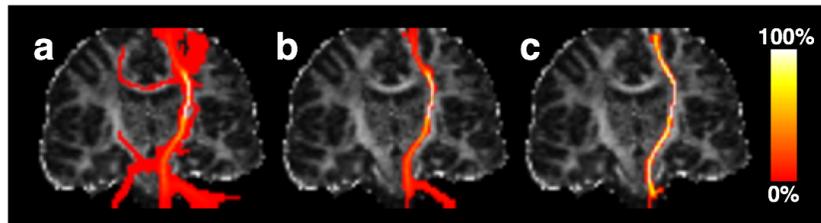


Fig. 3. An untreated visitation map for a left pyramidal tract in the data set (a). Equivalent visitation maps after thresholding at 1% of initiated streamlines (b), and after pruning (c) are also shown. No threshold is applied in the latter case. The colour scale indicates the proportion of streamlines passing through each voxel, with red indicating fewest and yellow most. The underlying images are FA maps.

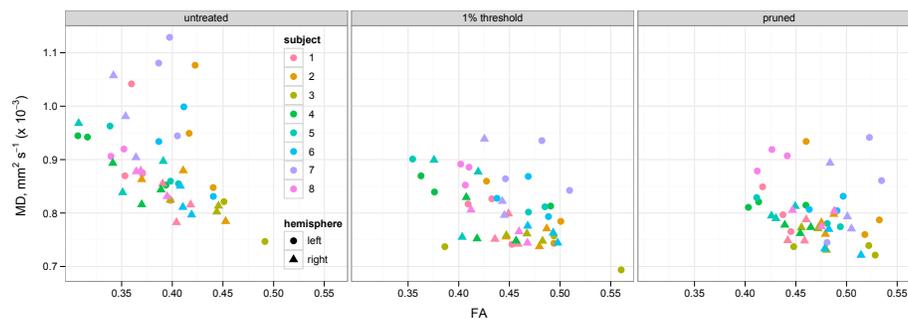


Fig. 4. Scatter plot of FA against MD based on binarised visitation maps for bilateral pyramidal tracts, using different thresholding and pruning strategies on the data.

Table 1. Estimated group mean and error standard deviations for two major tracts.

	1% threshold		5% threshold		pruned		untreated	
	μ	σ_e	μ	σ_e	μ	σ_e	μ	σ_e
pyramidal tracts, FA	0.448	0.038	0.481	0.038	0.469	0.024	0.389	0.030
pyramidal tracts, MD [†]	0.804	0.044	0.772	0.038	0.798	0.042	0.891	0.067
cingulum bundles, FA	0.374	0.033	0.438	0.045	0.386	0.044	0.285	0.021
cingulum bundles, MD [†]	0.782	0.033	0.740	0.035	0.763	0.030	0.907	0.035

† mm² s⁻¹ ($\times 10^{-3}$)

Between-subject (σ_b^2), within-subject (σ_w^2) and error (σ_e^2) variances over the data set are thereby distinguished from one another. This model was fitted using the “nlme” package for the R statistical environment (<http://stat.bell-labs.com/NLME/>), using the restricted maximum likelihood method [10]. The group mean and error variances for pyramidal tracts and cingulum bundles are summarised in Table 1, using visitation thresholds of 1% and 5% as well as untreated and pruned data. Mean FA and MD vary substantially depending on the threshold level applied, with the untreated results differing noticeably from the rest. The pruning algorithm produces values between those corresponding to thresholds of 1% and 5%, and on average the smallest error variance.

4 Discussion

We have demonstrated in this work a process by which a model of tract topology, combined with a predefined reference tract, can be used to select seed points for optimal tract segmentation, and also to retain or reject individual streamlines based on their probabilities under the model. The latter “pruning” method is a substantial improvement over standard thresholding approaches.

The absence of any user-specified parameters is a major advantage of the technique. It is rarely advisable to calculate parameters of interest over a region segmented using untreated tractography output (e.g. Fig. 3a), but we observe from Fig. 4 and Table 1 that the absolute recovered values and variances of such parameters are strongly dependent on the chosen threshold level. Moreover, the tacit assumption that erroneous pathways are nonrepeatable is false—as shown by the remaining false positive in Fig. 3b—and so finding a single threshold level which works well for different tracts, or even different parts of a single tract, is essentially impossible. Unlike a simple threshold, our model is sensitive to the meaning of the streamline data, and flexible enough to allow appropriate deviation from the reference tract whilst rejecting streamlines which do not follow its whole length, or branch off it. Since the streamline rejection criterion, (4), is specific to each subject, variation in tract shape from individual to individual is implicitly accounted for. Compared to region-of-interest approaches to streamline selection, our approach is very much less labour-intensive. Although the technique described in [2] also attempts to remove irrelevant tractography output, it is not tract specific and continues to rely upon a user-specified threshold.

In the present study we have limited application of our technique to probabilistic tractography using single seed points. Whilst the method could be directly applied to probabilistic or deterministic fibre tracking output derived from a neighbourhood of seed voxels, it could not, in its present form, be applied to whole-brain deterministic tractography. In addition, our focus has been on segmenting very specific pathways in groups of subjects with very high consistency, rather than covering the entire extent of the complex tracts of interest. These decisions impose some limitations on the immediate scope of this work, but the general approach has broad applicability. All tractography techniques raise the question of validation, but an increasing number of studies are vindicating fibre tracking, and our approach does not make validation any more difficult.

Finally, we have shown that the method ameliorates the usual falloff in connection probability with distance from the seed region. By retaining only sampled streamlines which accurately represent the tract of interest, estimates of connection probability are more robust, and reflect only genuine uncertainty in the tract location, rather than the effects of noise.

Acknowledgment: This work was supported by EPSRC grant EP/C536851/1.

References

1. Jones, D.: Studying connections in the living human brain with diffusion MRI. *Cortex* **44** (2008) 936–952
2. Morris, D., Embleton, K., Parker, G.: Probabilistic fibre tracking: differentiation of connections from chance events. *NeuroImage* **42** (2008) 1329–1339
3. Clayden, J., Storkey, A., Bastin, M.: A probabilistic model-based approach to consistent white matter tract segmentation. *IEEE Trans Med Imag* **26** (2007) 1555–1561
4. Maddah, M., Zöllei, L., Grimson, W., Westin, C.F., Wells, W.: A mathematical framework for incorporating anatomical knowledge in DT-MRI analysis. In: *Proc ISBI*. (2008) 105–108
5. Clayden, J., Storkey, A., Muñoz Maniega, S., Bastin, M.: Reproducibility of tract segmentation between sessions using an unsupervised modelling-based approach. *NeuroImage* **45** (2009) 377–385
6. Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., Clare, S., Matthews, P., Brady, J., Smith, S.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med* **50** (2003) 1077–1088
7. Hua, K., Zhang, J., Wakana, S., Jiang, H., Li, X., Reich, D., Calabresi, P., Pekar, J., van Zijl, P., Mori, S.: Tract probability maps in stereotaxic spaces: Analyses of white matter anatomy and tract-specific quantification. *NeuroImage* **39** (2008) 336–347
8. Muñoz Maniega, S., Bastin, M., McIntosh, A., Lawrie, S., Clayden, J.: Atlas-based reference tracts improve automatic white matter segmentation with neighbourhood tractography. In: *Proc ISMRM*. (2008) 3318
9. Jenkinson, M., Smith, S.: A global optimisation method for robust affine registration of brain images. *Med Image Anal* **5** (2001) 143–156
10. Bates, D., Pinheiro, J.: Computational methods for multilevel modelling. Technical report, Bell Laboratories (1998)