Automatic, Age Consistent Reconstruction of the Corpus Callosum Guided by Coherency from In Utero Diffusion-Weighted MRI — Supplemental Material

David Hunt, Manjiri Dighe, Christopher Gatenby, and Colin Studholme, *Member, IEEE*

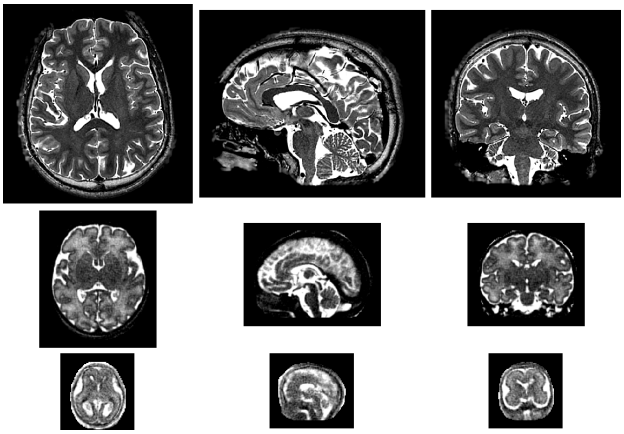

Fig. 11. Comparison between male adult (top), ≈ 36 weeks GA fetal (middle), and ≈ 20 weeks GA fetal (bottom) T_2-weighted images. Intensities have been rescaled between subjects to maximize contrast.

APPENDIX A
ANATOMICAL SCALES

Limits in resolution are particularly challenging for imaging fetal brains, where the target anatomy is much smaller than the more developed and expansive structures of adult, and even post-natal subjects. Figure 11 demonstrates this disparity in subject anatomy for three subjects, one at the youngest extreme of the present fetal study of ≈ 20 weeks GA, one at the oldest fetal extreme ≈ 36 weeks GA, and one of an adult subject. These three stages also demonstrate the rapid and drastic changes that the brain undergoes in utero, with a significant increase in landmarks during the last trimester, yet with many more still to develop before maturity.

APPENDIX B
APPLICATION TO CROSSING PHANTOM

As an intermediate example between the schematics of Fig. 1 and the fetal data in Sec. IV, we present results of a simple phantom with two distinct populations of fibers and noise. In half of the volume, the dominant direction of diffusion is along the x-axis (right-left), while the dominant direction in the other half is along the y-axis (forward-backward). The boundary between these two regions, as viewed from above in Fig. 12(a), includes a slight parabolic dependence on the y coordinate to encourage mixing between the two tracts. We introduce noise in each voxel by adding a random vector with maximum amplitude 90% of the base signal. Figure 12(b) shows the two tracts from the fibers that intersect the ROI at the center of the volume. While the dominant direction of diffusion is mostly consistent across the length of each tract, noise causes fibers to divert near the boundary or to terminate prematurely from the turning constraint, so that few fibers span the complete length of either tract (Fig. 12(c)).

In applying the proposed method to a region that is dominated by the forward-backward tract (higher z coordinates), we observe the appropriate subsets of fibers both for filtering to a coherent core and for parametrization for a full tract (Fig. 13), despite the influence of significant noise and diversions. By translating the ROI along the z-axis, we can select a set of fibers that is dominated by the left-right tract (lower z-coordinates), as shown in Fig. 14. Because of the slight parabolic feature that defines the boundary between the two diffusion regions, the left-right tract dominates for the z-centered mask. Since each tract is represented nearly evenly when the mask spans the normalized z-coordinate from $[1/4, 3/4]$, the first cycle of reconstruction does not yield a clean, single tract (Fig. 15). The tie is broken during the second cycle, which uses the first attempted reconstruction of a tract as input, whereupon the left-right tract dominates more clearly.

With some adaptation (perhaps as presented in App. F), this method could be used to explore general white matter (WM) structure without any target tract specified in order to extract the dominant patterns without the use of specific anatomical landmarks. A preliminary ROI could be selected, for example, by choosing WM regions with high fractional anisotropy.

APPENDIX C
COMPARISON TO FIBER CLUSTERING

Established methods of clustering produce sets of fibers based on spatial proximity and similarity in paths. The method we present in the present paper has the same goal of identifying coherent bundles, however our approach builds a model of

D. Hunt (e-mail: davidoh@uw.edu) and Colin Studholme are with the Department of Pediatrics, University of Washington, Seattle, WA, 98195 USA.

M. Dighe, C. Gatenby, and C. Studholme are with the Department of Radiology, University of Washington, Seattle, WA, 98195 USA.

C. Studholme is with the Department of Bioengineering, University of Washington, Seattle, WA, 98195 USA.

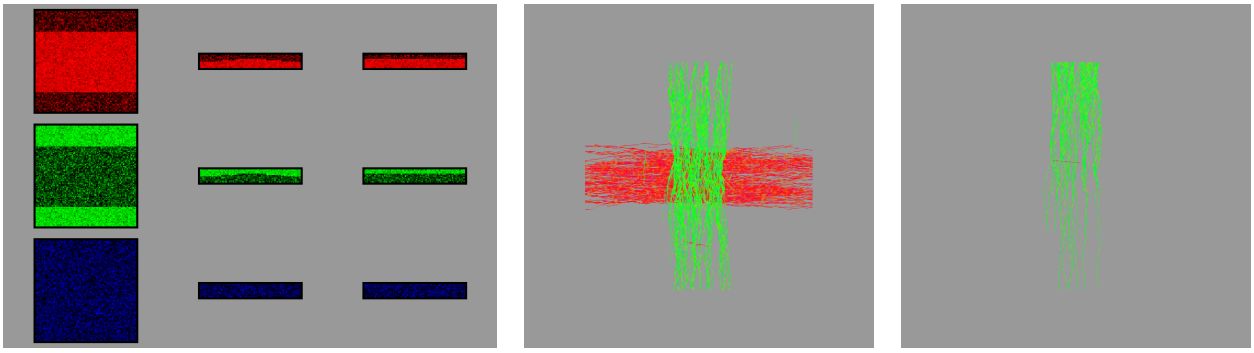

(a) (b) (c)

Fig. 12. (a) Components of simulated diffusion peaks in the x-direction (top row), y-direction (middle row), and z-direction (bottom row) from three views: axial (left column), coronal (middle column), and sagittal (right column). Bright colors indicate a large component, while black indicates a small component. (b) Axial view of the target competing tracts, one running left-right (red), the other forward-backward (green). (c) Propagation of fibers from the ROI at the boundary of the simulated volume. Fibers do not densely span the path between opposite sides.

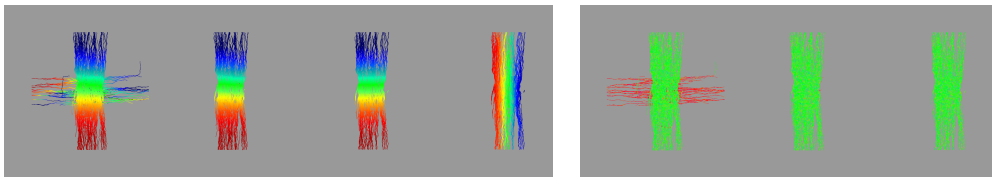

(a) (b)

Fig. 13. (a) From left to right: p_0 for an initial ROI; p_0 for the preliminary core; p_0 for the final reconstruction; p_1 for the final reconstruction. (b) The same sets of fibers as in (a), but with colors indicating the direction of diffusion at each vertex.

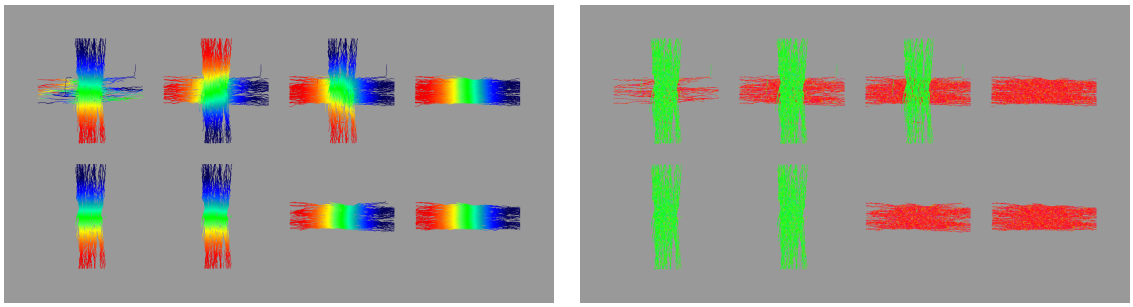

(a) (b)

Fig. 14. Initial sets of tractography (top row) and final reconstructed tracts (bottom row) for different selections of fibers. (a) Colors indicate diffusion-parallel parametrization p_0. (b) The same sets of fibers as in (a), but with colors indicating the direction of diffusion at each vertex. From left to right, the ranges of the ROI in the normalized z coordinate are: [0, 1/2]; [1/8, 5/8]; [1/4, 3/4]; [1/2, 1].

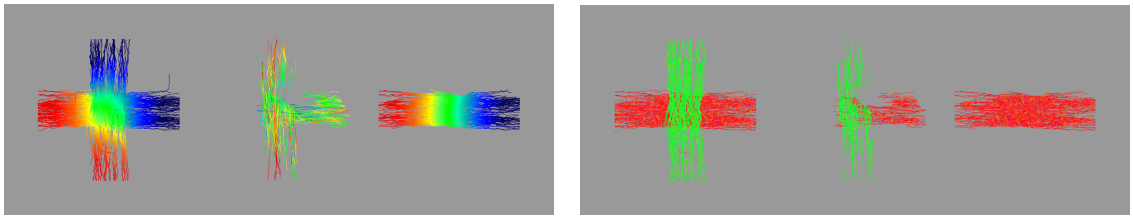

(a) (b)

Fig. 15. (a) From left to right, sets of fibers with p_0 for: initial set of tractography, first reconstruction cycle, second reconstruction cycle. Because each tract is represented nearly equally, the first cycle fails to select a single tract. (b) The same sets of fibers as in (a), but with colors indicating the direction of diffusion at each vertex.

the geometry for the entire tract and offers the advantage of identifying the spatial pattern of coherent *portions* of fibers.

As an example, we show in Fig. 16 the clustering of fibers using the method from [1], applied to clustering tractography on both the scale of the whole brain and the corpus callosum (CC) region of interest (ROI). Using the guidance of the original Nyström method proposal [2], we sample 100 fibers for each set of fibers, 4,183 for the callosal ROI, and 45,570 for whole brain tractography after an additional filter for fibers with length of at least a decimeter. The final clustering is chosen to be the realization with the maximum silhouette coefficient from 100 trials of the k-means algorithm.

As expected, the clusterings produced by this method group similar fibers together, however both ROI and whole brain clusterings exhibit fragmentation of the CC. In the case of whole brain tractography, while one cluster in Fig. 16(d) does cover much of the CC, non-callosal regions are grouped with sections of the CC. Unfortunately, robustly automating an appropriate number of clusters for the k-means algorithm is not practical over the age range and image quality for in utero DWIs. Even provided with such an automated threshold, the spatial location of the CC within the brain would still be required to identify the appropriate cluster, hence the use of a ROI (as we propose in Sec. III-C) is unavoidable.

Furthermore, because of the fragmentation of the tract of interest, a geometric model based on the dominant shape of each bundle would be required to stitch the fragments together into a single tract. Identifying which portions of fibers within a fragment are appropriate for stitching and which are not further complicates this approach. Hence for simplicity and consistency, we apply our proposed method to the entire set of fibers determined by the automatically produced ROI and utilize our geometry-based modeling to filter fibers that do not belong to the dominant tract shape. Our results show that the extraction of a coherent core is possible even with significant diversions that lead to fragmentation in fiber clustering schemes.

APPENDIX D
IMPLEMENTATION DETAILS

A. Diffusion-parallel parametrization

We apply the diffusion-parallel parametrization method to a smaller subset of fibers in the tract to reduce computational expense. To accomplish this in practice, one must ensure that the subset spans the extent of physical space sufficiently and also with adequate density so that a single connected component exists. The sparser parametrization can then be propagated to the denser tractography by choosing the nearest parametrized point.

This simple propagation scheme can lead to non-monotonic parametrizations along the length of a fiber, especially when a fiber is diverted into an adjacent tract. To correct this, we search for the largest segment of the fiber that maintains a stable linear relation between the parametrization and length along the fiber. Inspired by the Ramer-Douglas-Peucker algorithm [3], [4] for sampling curves, we sequentially add intermediate sample points and quantify the linearity of the fiber portion using the product of the coefficient of determination of the linear regression (R^2) and the magnitude of the span of the parametrization. A greedy search guides the subsequent subsections and the parametrization from the portion with the largest product is expanded to the rest of the fiber. The parametrization is normalized by reassigning negative values to 0 and values greater than 1 to 1. Portions at the tips of fibers that have constant p_0 values are useful for identifying fibers that do not match the coherent tract core (Sec. II-C).

B. Diffusion-orthogonal parametrization

Since all fibers may not span the entirety of the tract in practice, sampling the tract at multiple values of p_0 makes it more likely that we include all fibers and appropriately associate p_1 with the longer extent (the tract's width) and p_2 with the shorter extent (the tract's thickness). Because the diffusion-orthogonal parametrizations for different values of p_0 are initially calculated independently of each other, there may be a change in scale as fibers diverge or converge, or as some portions of the tract fail to propagate as far as others. To make consistent comparisons, we allow a basic linear scaling with constant offset for the p_1 and p_2 parametrizations between different p_0 selections. This allows us to assign constant diffusion-orthogonal values for the entire length of a single fiber, based on the average of those parametrizations at various values of p_0. Final diffusion-orthogonal parametrizations are shown in the results in Figs. 7(c) and (d).

If any fibers are not assigned diffusion-orthogonal parametrizations — e.g., this may arise if the p_0 sampling is not dense enough to fall within the span parametrized on the fiber — the parametrizations can be inferred from the regressed core surface of the tract. In such cases, after determining p_1 via a golden-section search along the curve of constant p_0, the value of p_2 is assigned based on the oriented (above or below $p_2 = 0.5$) distance from the core surface.

The medial line in Fig. 4 illustrates typical drifts at the tract extremes (specifically, to the left of line segment 3a and below line segment 3d). For the purposes of utilizing the parametrization to enforce simple geometric constrains, as we propose here, this drift is not problematic. For more sensitive analyses along the path of the medial line, the length corresponding to half the mean thickness can safely be ignored; alternatively, one could impose additional regularization constraints on the curve. During the construction of the medial line, highly irregular surface features, such as deep, skinny concave sections may have high ratios R but can be excluded by detecting that the straight line connecting the points is not contained within the boundary.

C. Coherent core extraction

Plots in Fig. 17 describe the performance of the core extraction for the *base* cycle, i.e. the best initial extraction attempt upon which subsequent cycles of the sequence from Fig. 1 take as a seed. We found a limit of 40 iterations to be reasonable over the age range considered, with the typical 20-week subject taking ≈ 16 iterations and the typical 35-week subject taking ≈ 35 iterations (Fig. 17(a)). Figure 17(b) shows

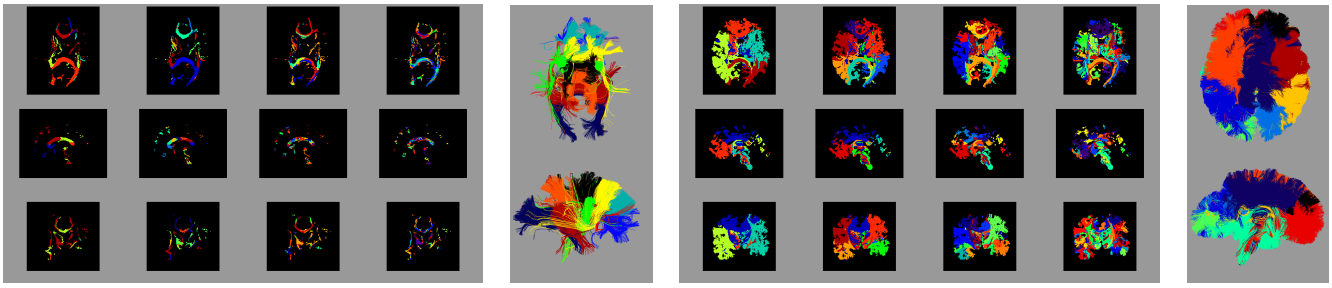

(a) (b) (c) (d)

Fig. 16. Bundles identified by fiber clustering for the same ≈ 32 weeks GA subject from Fig. 2. (a) Cross-sectional maps of clusters from the callosal ROI for number of clusters $k = 5, 10, 32, 100$ from left to right. (b) Clustered tractography for the callosal ROI corresponding to $k = 10$ clusters. (c) Cross-sectional maps of clusters from whole brain tractography for the same numbers of sets as in (a). (d) Clustered tractography (with a cut in the side view to observe the CC) from whole brain tractography corresponding to $k = 10$ clusters. Colors indicate different cluster labels.

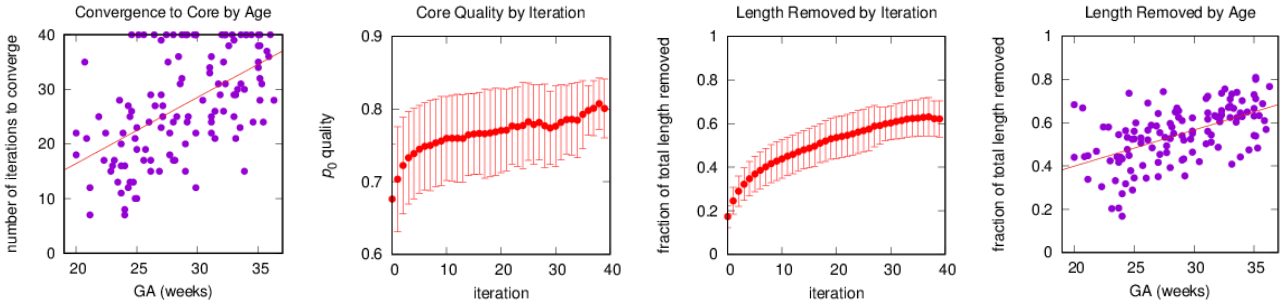

(a) (b) (c) (d)

Fig. 17. Performance of the method for identification of the base cycle's coherent core. (a) Number of iterations before convergence or termination for core extraction with line showing trend of and average of 1.2 extra iterations per week GA. (b) Average parametrization quality [Eq. (1)] across scans as a function of number of iterations of filtering for a coherent core. (c) Average fraction of total fiber length removed across scans as a function of number of iterations. (d) Fraction of total length removed for individual scans with line showing trend of removing an additional 1.6% more of the total length with each additional week GA.

the improvement in parametrization quality as the method removes more fibers, the normalized length sum of which is shown in Fig. 17(c). The nonmonotonicity in these plots comes from the variable number of iterations for convergence across scans (compare Fig. 17(a)). Figure. 17(d), similarly to (c), shows the total fraction of length of the original raw ROI tractography that is identified by the method to belong to a coherent tract core as a function of age. Not surprisingly, the fraction of diverted fiber length increases with age as the structural connectivity becomes more elaborate.

Figure 17(d) shows that a large portion, if not most of the length of the tractography associated with the initial ROI belongs to fibers that at some point do not cohere to the dominant shape of the tract. These diversions result in misleading connectivity.

Example core extractions for various thresholds are shown in Fig. 18. Thresholds are chosen by considering the application to data, based on observations from the characteristics of several trial thresholds that may vary depending on signal quality and the nature of the tract to be extracted. We show the agreement threshold of 0.3 as an alternative to the chosen threshold of 0.2 since in this example subject it also yields a reasonably coherent core. However, while the set of fibers for this example subject is likely adequate for expansion of the core for an agreement threshold of 0.3, the more restrictive

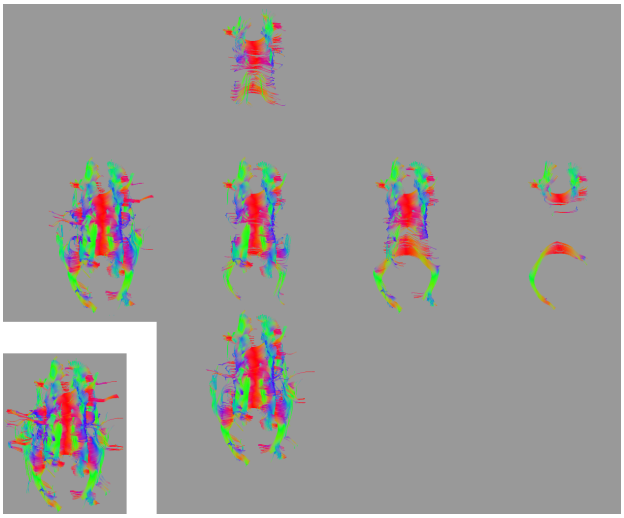

Fig. 18. Extracted cores with different filtering thresholds for a subject ≈ 26 weeks GA (different than that in Fig. 22(a) and (b)). The thresholds for agreement score [Eq. (2)] from left to right are: 0.1, 0.2, 0.3, and 0.4. The threshold for extremal parametrization values from bottom to top are: 20%, 40%, and 60%. The resulting core from the parameters chosen in this study is shown at the intersection of the row and column. For reference, the original ROI tractography is shown in the lower left.

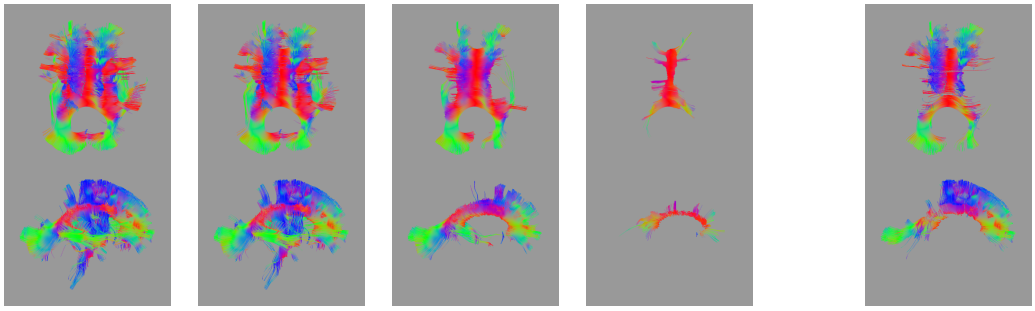

(a) (b) (c) (d) (e)

Fig. 19. Tractography from different thresholds for the permissible sharpness of turns: (a) $\pi/2$ (used in this study for initial tractography) (b) $\pi/4$, (c) $\pi/10$, and (d) $\pi/20$ (rad/mm). For comparison, the coherent core using the method from Sec. II-C is shown in (e). Colors indicate the primary direction of diffusion at each vertex.

(higher) threshold can cause problems in other subjects. Consequently, we choose the more permissive threshold of 0.2 for stability across scans.

A common alternative for filtering diverted fibers involves enforcing an appropriately chosen turning restriction so that fibers terminate upon encountering sharp turns. Examples of this filtering for various turning thresholds are shown in Fig. 19. We used the most permissive $\pi/2$ rad/mm turning constraint in Fig. 19(a) for the generation of tractography in this study. While Fig. 19(b) and (d) are too permissive and too restrictive, respectively, Fig. 19(c) does show an adequate core from a turning threshold of $\pi/10$ rad/mm. However, when compared to the core in Fig. 19(e) from the method in Sec. II-C, we see that filtering portions of fibers by turning angle alone can be too restrictive in removing portions of the superior extremes of the tract, while being simultaneously too permissive in allowing the more gradual diversions into the tapetum. In general, the length scales and shapes change with age, whereas the constant agreement score threshold [Eq. (2)] of 0.2 is adequate across all ages.

There are many possible ways to produce a coherent core from a set of fibers; while the core expansion method (Sec. II-D) is robust enough to use some other thresholds or filtering methods, utilizing the parametrization for initial filtering as proposed in Sec. II-C offers an even more complete initial set of tractography to guide the subsequent refinement and reconstruction of the tract. The precision of this proposed coherent core filtration and the stability of the iterative expansion of the coherent core allow us to use a very permissive turning threshold for the initial tractography for extracting the corpus callosum.

D. Expanding the coherent bundle

Not all initial coherent core estimates yield a complete tract, but the full width of the tract is often still recoverable. Figure 20 shows a sequence of steps for a subject with only a small initial coherent region. Despite the preliminary coherent region spanning only a fraction of the width of the CC, the iterative refinement eventually extends the coherent region to the full width of the tract to include the forceps major.

Since the target coherent bundles do not have complex topologies, we determine the corresponding point on the

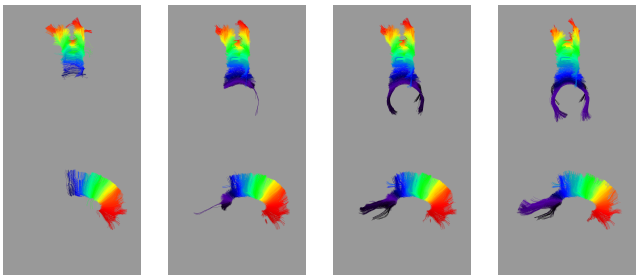

(a) (b) (c) (d)

Fig. 20. Iterations in the expansion of a partial coherent core to the full width of the CC: (a) core seed; (b) first iteration; (c) second iteration, (d) final selection. Color indicates the diffusion-orthogonal parametrization p_1 from red at one extreme to blue at the other.

regressed surface $\boldsymbol{r}_i(\boldsymbol{f}_i)$ that is nearest the fiber vertex $\boldsymbol{f}_i$ using Newton's method to find local minima in Euclidean distance metric. For a subset of vertices along a fiber, including both ends and intermittently along its path, we distribute seeds for Newton's method evenly on the full 2D extent $[0,1]\times[0,1]$ of the p_0p_1-surface. The remaining points use $\boldsymbol{r}(\boldsymbol{f}_j)$ from the nearest vertex $\boldsymbol{f}_j$ that has been mapped to the regressed surface for the single seed to find $\boldsymbol{r}(\boldsymbol{f}_i)$.

To stabilize the reconstructed tract and to prevent shifting to adjacent tracts over multiple iterations, a vertex is persistently included in the model fit for subsequent steps once it has been accepted. The vertex may still be removed from contributing to the final extraction of the tract if the vertex does not belong to the single central connected component (i.e., the contiguous set of vertices nearest $p_0 = 0.5$) or if it belongs to a portion that is not consistent with the main tract pathway (e.g., for the CC, portions that divert into the tapetum or corticospinal pathways).

To expedite parametrization of the new collection of fibers, the full parametrization $\boldsymbol{p}(\boldsymbol{f}_i)$ is inferred from the nearest point $\boldsymbol{r}(\boldsymbol{f}_i)$ on the tract core surface. To keep the properties of the parametrization consistent, only p_0 is allowed to vary along the length of the fiber, while p_1 and p_2 are assigned according to their respective averages over the fiber. We select the final tract as the one that includes the largest number of vertices, which by construction of the equally-spaced vertices for the fibers in the tractography, is the tract with the greatest sum of

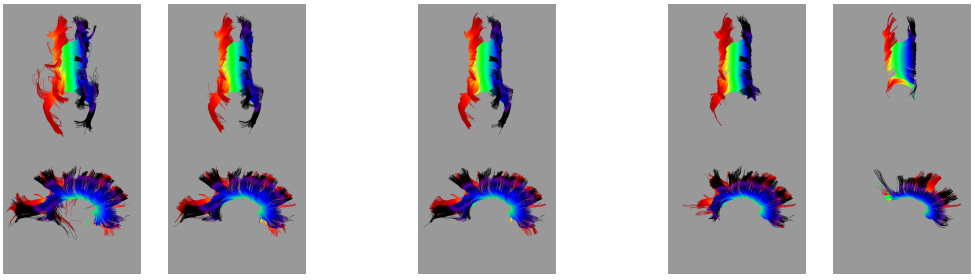

(a) (b) (c) (d) (e)

Fig. 21. Reconstructed tracts for the same ≈ 32 weeks GA subject from Fig. 2 showing varying quality for different vertex score thresholds [from Eq. (3)] of (a) 0.02, (b) 0.1, (c) 0.15 (chosen default), (d) 0.3, and (e) 0.6. Color indicates the diffusion-parallel parametrization p_0.

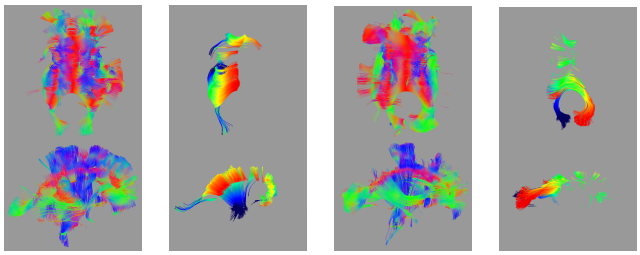

(a) (b) (c) (d)

Fig. 22. Problematic cases. (a) The majority of fibers in a scan ≈ 26 weeks GA diverge into the corticospinal tract; (b) the proposed method fails to identify a coherent half-pipe shape when it is obscured in this way. (c) A different subject ≈ 29 weeks GA with tractography dominated by diversions into the tapetum; (d) the proposed method fails to extend the coherent bundle beyond the largely diverted portion in (c). Colors in (a) and (c) correspond to local fiber direction, while colors in (b) and (d) correspond to the attempted parametrization.

fiber lengths.

Example reconstructions for different values of the score threshold are shown in Fig. 21. The default vertex-wise score threshold of 0.15, shown in Fig. 21(c) exhibits the most appropriate reconstruction. As the threshold decreases in Fig. 21(b) and (a), more inappropriate portions of fibers remain. As the threshold increases in Fig. 21(d) and (e), less of the appropriate portions of fibers remain.

In addition to the 34 typical subjects in Fig. 8, we show 2 additional challenging scans to illustrate the limitations of our approach, where diverted fibers dominate the tractography and yield inadequate reconstructions (Fig. 22). For the first case in Fig. 22(a) and (b), fibers divert away from the superior cortex toward the corticospinal tract. Because larger regions of consistent curvature belong to these diversions, our method fails to identify the target CC. For the second case in Fig. 22(c) and (d), the diversions into the tapetum compete with fibers of the forceps major. The proposed method fails to identify a clearly dominant coherent pattern that spans the width of the CC, and the final reconstruction includes only hints of more anterior regions.

APPENDIX E
EXTRACTED TRACTS WITH DIRECTIONAL COLOR CODING

Figure 23 shows the reconstructed CC with directional color coding for the same set of subjects as Fig. 8. For a more direct comparison with the improvement over the raw set of fibers, see Fig. 6. Some remnants of diversions survive on the boundary of the tracts in Fig. 23, especially for older subjects with richer structures. For example, some green is visible in some reconstructions where the corpus callosum is adjacent to the cingulum bundles. However, since these fibers are truncated at the point they begin to divert, their influence on the modelled geometry is negligible and they avoid contributing to inappropriate connectivity.

While it may appear that there are long distance diversions, these are not the result of extended lengths of fibers travelling along these paths; rather, they indicate regions where partial voluming diverts fibers, resulting in the truncation of many fibers. That the tips of these fibers do not represent the end connectivity of the tract is identified by the fact that the diffusion-parallel parametrization p_0 does not span the full range [0, 1]. The more appropriate connectivity of the tract is recovered by using the tips of fibers that do approach the extremes of the parametrization.

APPENDIX F
APPLICATION TO NONCALLOSAL TRACT

Reconstructing the complete CC is an important first step to accurately identifying other tracts and their connectivity. As an example of the utility of this reconstruction, we apply the automatically extracted CC to the identification of the set of tractography fibers for corticospinal tracts.

We begin with the subset of fibers that intersect the brain stem, identified by the same lobe segmentation mentioned in Sec. III-C. However, this set includes diversions into other tracts, including the CC, as shown in Fig. 24(a). Additionally, some fibers span both corticospinal tracts and the appropriate portions should be attributed to a single hemisphere.

We first leverage the reconstructed CC to produce a surface between points along the central $p_0 = p_2 = 0.5$ line of the regressed surface and the point in the brain stem that is farthest from the center of the CC. We then select the largest connected component of voxels that are within $\pi/3$ radians of the nearest point on the initial surface, which extends beyond the brain stem and allows us to assign a hemisphere label. Using the portions of fibers that are at least one decimeter in length contained in these regions gives two separated seeds for the corticospinal tracts. Extracting the coherent core from these seeds yields the parametrizations shown in Fig. 24(b).

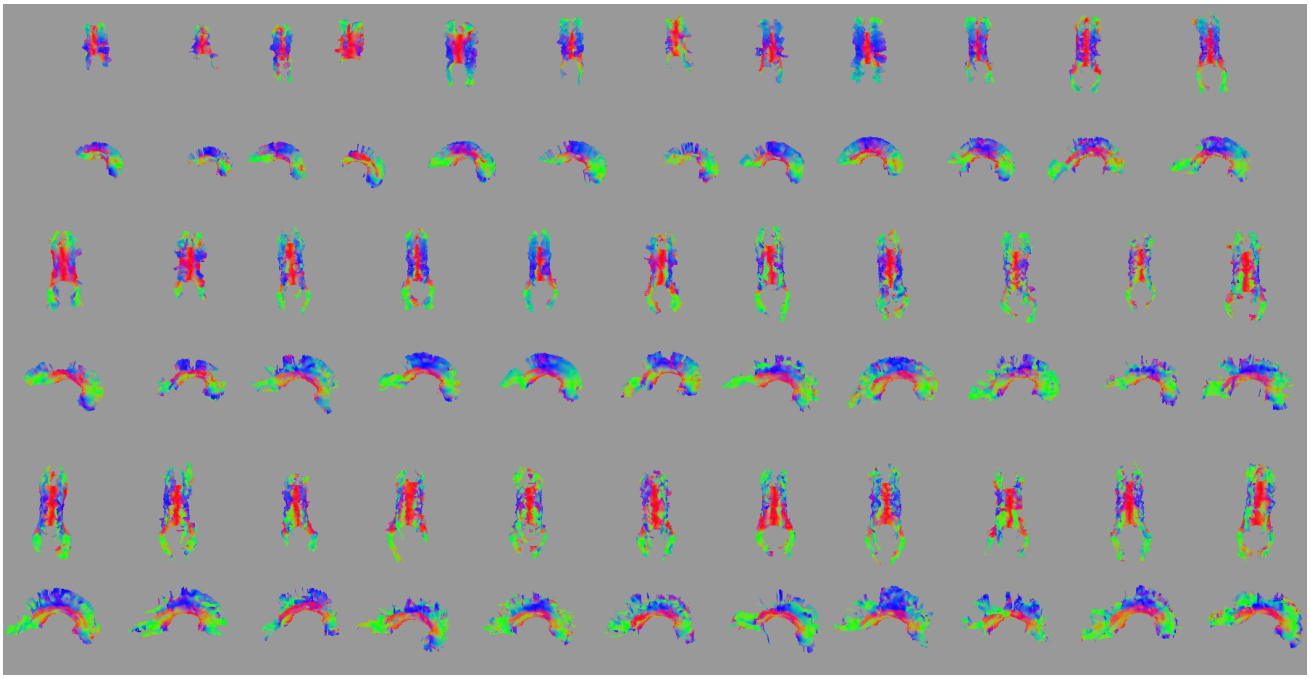

Fig. 23. Automated reconstruction of the CC for individual subjects, ordered by age (youngest of $\approx$ 20 weeks GA in the upper left to oldest of $\approx$ 36 weeks GA weeks in the lower right). Colors indicate the value of the local direction of the fiber.

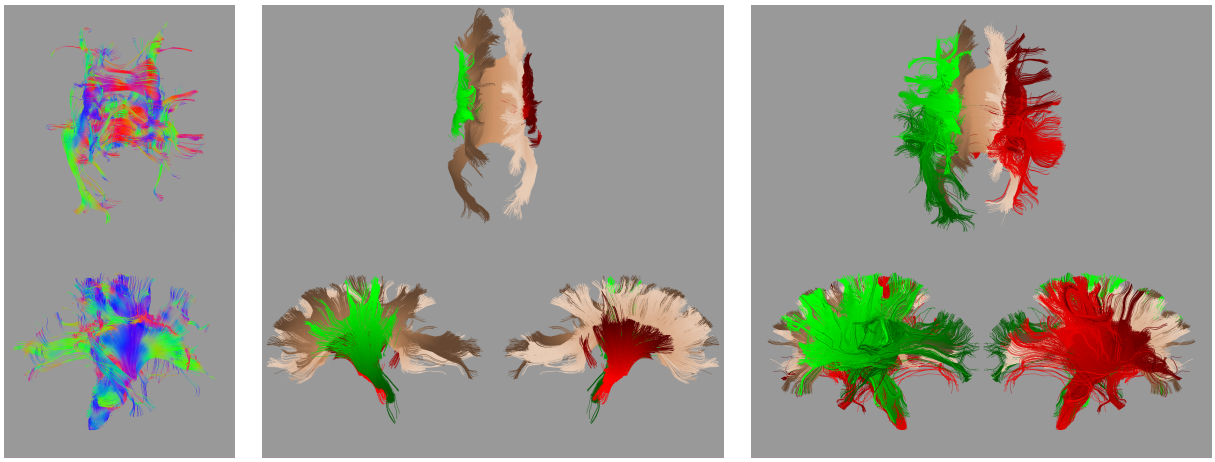

(a) (b) (c)

Fig. 24. For the same $\approx$ 32 weeks GA subject from Fig. 2: (a) Initial seed for corticospinal tracts from fibers that pass through the brain stem. Colors show fiber direction. (b) Filtered seeds for each of the corticospinal tracts. Colors indicate tract: red for right corticospinal, lime green for left corticospinal, copper for CC; shades indicate the value of the parametrization p_0. (c) Larger set of fibers for expanding the coherent core. Colors are the same as in (b).

For a more complete tract, the set of possible fibers to consider can be expanded to all those that intersect a dilation of the region defined by the conservatives seeds, shown in Fig. 24(c). This set has many diversions into other tracts, but can still be refined with a method similar to that used for the CC, perhaps benefiting from relaxing the restriction on a single half pipe for the complete, corrugated corticospinal tracts. The approach presented in this section is not optimized, but it demonstrates how the information from the reconstructed CC benefits the refinement the extraction of other tracts. Complete automated extraction of corticospinal and other tracts will be important future studies, but the focus of the current paper is the extraction of the CC. Consequently, we do not present a comprehensive extraction pipeline for other tracts.

REFERENCES

[1] L. J. O'Donnell, M. Kubicki, M. E. Shenton *et al.*, "A method for clustering white matter fiber tracts," *Am. J. Neuroradiol.*, vol. 27, no. 5, pp. 1032–1036, 2006.

[2] C. Fowlkes, S. Belongie, F. Chung *et al.*, "Spectral grouping using the Nyström method," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 2, pp. 214–225, 2004.

[3] U. Ramer, "An iterative procedure for the polygonal approximation of plane curves," *Comput. Graph. and Image Process.*, vol. 1, no. 3, pp. 244–256, 1972.

[4] D. H. Douglas and T. K. Peucker, "Algorithms for the reduction of the number of points required to represent a digitized line or its caricature," *Cartographica*, vol. 10, no. 2, pp. 112–122, 1973.