Self-supervised segmentation and characterization of fiber bundles in anatomic tracing data

Vaanathi Sundaresan[a,*], Julia F. Lehman[b], Chiara Maffei[c], Suzanne N. Haber[b,d], Anastasia Yendiki[c]

[a]*Department of Computational and Data Sciences, Indian Institute of Science, Bengaluru, Karnataka 560012, India*
[b]*Department of Pharmacology and Physiology, University of Rochester School of Medicine, Rochester, NY, United States*
[c]*Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, Charlestown, MA, United States*
[d]*McLean Hospital, Belmont, MA, United States*

Abstract

Anatomic tracing is the gold standard tool for delineating brain connections and for validating more recently developed imaging approaches such as diffusion MRI tractography. A key step in the analysis of data from tracer experiments is the careful, manual charting of fiber trajectories on histological sections. This is a very time-consuming process, which limits the amount of annotated tracer data that are available for validation studies. Thus, there is a need to accelerate this process by developing a method for computer-assisted segmentation. Such a method must be robust to the common artifacts in tracer data, including variations in the intensity of stained axons and background, as well as spatial distortions introduced by sectioning and mounting the tissue. The method should also achieve satisfactory performance using limited manually charted data for training. Here we propose the first deep-learning method, with a self-supervised loss function, for segmentation of fiber bundles on histological sections from macaque brains that have received tracer injections. We address the limited availability of manual labels with a semi-supervised training technique that takes advantage of unlabeled data to improve performance. We also introduce anatomic and across-section continuity constraints to improve accuracy. We show that our method can be trained on manually charted sections from a single case and segment unseen sections from different cases, with a true positive rate of

~0.80. We further demonstrate the utility of our method by quantifying the density of fiber bundles as they travel through different white-matter pathways. We show that fiber bundles originating in the same injection site have different levels of density when they travel through different pathways, a finding that can have implications for microstructure-informed tractography methods. The code for our method is available at `https://github.com/v-sundaresan/fiberbundle_seg_tracing`.

Keywords: Anatomic tracing, fiber bundle detection, self-supervised, contrastive loss, temporal ensembling, fiber density

1. Introduction

Higher cortical function emerges from a combination of functional specialization at each cortical location and connectivity between locations, which, together, comprise complex anatomic networks (Haber et al., 2022; Geschwind, 1965). Understanding those network connections is crucial for detecting abnormalities in disease. Anatomic tracing methods allow us to visualize brain connections by identifying the trajectories of individual axons from their origin to their termination. This includes the routes that axons follow to reach each of the major white matter bundles, their position as they travel within these bundles, and their exit points from the bundles to their terminal fields. An example is shown in Figure 1, which illustrates trajectories of fiber bundles from four different cortical injection sites, as they travel to and through the internal capsule (IC). Combining data from multiple injections, as in the figure, is invaluable for investigating the topographic organization of fibers from different cortical areas within large white-matter pathways such as the IC.

As these methods are not applicable to human subjects, we typically rely on tracer studies in nonhuman primates (NHPs) for accurate identification of cortical connections (Haber et al., 2022; Lehman et al., 2011; Öngür & Price, 2000). These studies provide the foundation for understanding the organization of white-matter

*Corresponding author
  Email address: `vaanathi@iisc.ac.in` (Vaanathi Sundaresan)
  URL: `http://cds.iisc.ac.in/faculty/vaanathi/` (Vaanathi Sundaresan)

Figure 1: **Organization of fibers in the internal capsule, as revealed by tracer injection studies.** Sagittal view of nonhuman primate brain illustrates the trajectories of fibers from four cortical injection sites, as they travel to and through the internal capsule. Orange: anterior cingulate fibers; yellow: dorsal prefrontal fibers; green: premotor fibers; blue: motor fibers.

pathways and for assessing the accuracy of pathways reconstructed by non-invasive neuroimaging in humans (Haber et al., 2023; Safadi et al., 2018; Jbabdi et al., 2013; Haber et al., 2022). In particular, the comparison of tracer injections to diffusion MRI (dMRI) tractography in the same NHP brain has generated important insights, e.g., on the fiber configurations that confound dMRI and on how dMRI data should be acquired and analyzed to maximize the accuracy of pathways reconstructed by tractography (Grisot et al., 2021; Maffei et al., 2022; Schilling et al., 2019; Yendiki et al., 2022).

Public databases of tracer injection data (e.g., Stephan et al. (2001); Kötter (2004); Bakker et al. (2012)) provide information on which cortical or subcortical regions are connected to each other (i.e., a "connectivity matrix"), but not on the trajectories that axon bundles follow to get from one region to the other. The full trajectories are needed, e.g., to map the topographic organization of white-matter bundles as in Figure 1, or to determine the exact locations in white matter where errors of tractography occur. The main challenge in building a database that contains the full trajectories of axon bundles from tracer experiments is that this would require extensive manual annotation. Figure 2 shows examples of manual chartings of fiber bundles from a tracer injection in the frontopolar cortex of a macaque monkey. This manual annotation is labor intensive and time consuming. The development of a

Figure 2: **Manually annotated fiber bundles from a tracer injection study.** Photomicrographs show coronal sections (1-3) from a macaque brain that received a tracer injection in the frontopolar cortex, with terminal fields at different cortical locations (1a, 2a, 3a). In the rostrocaudal direction, a fiber bundle stalk (1b) branches into two fiber bundles in prefrontal white matter (2b) and travels laterally in the external capsule (3b) and medially in the corpus callosum (3c). Manual chartings of dense and moderately dense bundles are shown as green and orange outlines, respectively.

computer-assisted tool for segmenting fiber bundles in tracer data is thus critical for accelerating this process and facilitating the work of anatomists.

There are two types of challenges in the development of a computer-assisted method for segmenting fiber bundles in anatomic tracer data: artifacts in the histological sections, and limitations in the available ground-truth annotations of the bundles. First, the staining and digitization of histological sections introduces substantial variation in the intensity of both the stained axons and the background, between different sections and cases. Second, spatial distortions introduced by staining and mounting make it difficult to ensure the consistency of segmentation labels between consecutive sections. Third, the fiber bundles that we aim to segment may have similar texture characteristics to, e.g., terminal fields or background staining, leading to false positives (FPs) in the segmentation. Finally, the manually drawn outlines of fiber bundles that can be used for training are only available on a limited number of sections, and may not include all fiber bundles in a section.

We seek to overcome the above challenges with the use of semi-/self-supervised segmentation approaches that can extract maximal information from limited training data, while being sufficiently generalizable in the presence of typical variation across

datasets from different brains and injection sites. Semi-/self-supervised methods have been shown to work well on generic noisy data and limited labels with uncertainties (Dinsdale et al., 2022; Chen et al., 2020; Feyjie et al., 2020; Perone et al., 2019; Sundaresan et al., 2022; Fischer et al., 2023; Du et al., 2023). In particular, contrastive learning, which aims to learn image features that are similar or different between segmentation classes (Chen et al., 2020; Zhao et al., 2023), has been used to segment histopathological images (Wu et al., 2022; Lai et al., 2021). Similarly, perturbation-based self-ensembling and temporal ensembling, where average predictions from prior epochs are used as pseudo-labels for training the current epoch (Li et al., 2020; Perone et al., 2019), have been shown to perform well in segmentation tasks with minimal manual annotations for training.

Prior work on segmentation of axons in microscopy data has been focused mainly on segmenting individual axons in nm-scale images with typical fields of view in the order of 1 mm or less (Zaimi et al., 2016; Mesbah et al., 2016; Naito et al., 2017; Zaimi et al., 2018; Wei et al., 2021). These methods are not directly applicable to our task. The relevant prior work on segmenting fiber bundles in whole-brain, μm-scale, histological sections from tracer experiments is quite scarce and has only been applied to marmoset brains (Skibbe et al., 2019; Woodward et al., 2020). These methods used the U-Net model (Ronneberger et al., 2015), showing the reliability and robustness of this architecture in tracer data segmentation. However, the U-Net model in these methods was trained in a fully supervised manner, which would be suboptimal for our case due to limited manual chartings. Hence, our goal is to use the U-Net architecture as a backbone within a more flexible, multi-tasking framework, trained in a semi-supervised manner to address the variability in the data and the shortage of manual chartings.

We propose the first deep learning-based method for computer-assisted fiber bundle detection on anatomic tracer data from macaque brains, using only a few manually labeled sections. We use an anatomy-constrained, self-supervised loss for contrastive learning within a multi-tasking model with a U-Net backbone, and a semi-supervised

temporal ensembling training technique for efficient improvement of predictive performance. In particular, we use contrastive learning to learn the contextual features of manually charted fiber bundles, given that these chartings consist of fiber-dense areas on a relatively homogeneous background. We use temporal ensembling to further enhance this contextual learning and improve robustness to noise via the averaging of predictions. We also reduce FPs with the use of continuity priors across predictions from consecutive sections. In addition to segmenting fiber bundles, our tool estimates the density of fibers within each bundle. We evaluate our method on sections from different brains and tracer injection sites, and we quantify the density of fiber bundles in various white-matter pathways. The tool is publicly available and can be deployed, e.g., for quantitative analyses of tracer data or for validation of tractography.

2. Method

2.1. Data

Ethics statement: All tracer experiments were performed in accordance with the Institute of Laboratory Animal Resources Guide for the Care and Use of Laboratory Animals and approved by the University of Rochester Committee on Animal Resources. See Lehman et al. (2011); Haynes & Haber (2013); Haber et al. (2006) for more details on tracer injection, immunocytochemistry and histological processing.

We use digitized, coronal histological sections from 13 macaques (M1 – M13), with a slice thickness of 50µm and in-plane resolution of 0.4µm. Every 24[th] section was previously processed to visualize the tracer, resulting in a slice gap of 400μm. Manual charting of fiber bundles had been done previously by an expert neuroanatomist under dark-field illumination with a 4.0 or 6.4x objective, using Neurolucida software (MBF Bioscience). Examples are shown in Figure 2. Fibers traveling closely were outlined as a bundle. Fiber bundle orientations were marked by charting some of the individual fibers within the bundles. These orientations were used as visual markers for identifying fiber bundles in the consecutive sections. The 2D outlines were combined across slices using IMOD software (Boulder Laboratory; Kremer et al. (1996)) to create 3D renderings of pathways (examples shown in Figure 1). These

renderings were used to further refine bundle contours and ensure spatial consistency across sections. The density of fibers within each bundle were assessed visually and the bundle was categorized as dense or moderate (shown in green and orange outlines, respectively, in Figure 2). For more information on manual annotation, see Grisot et al. (2021).

After the histological samples were digitized, manually charted region masks were realigned to them using an affine transform with 6 degress of freedom. Manual chartings were available for 88 sections from 3 macaques (M1 – M3).

1. **Dataset 1 (DS1)** was used for training and consists of 465 sections, including 25 charted sections (out of 88) from one macaque (M1), which has an injection in ventro-lateral prefrontal cortex (vlPFC), and 440 unlabeled sections from 10 macaques.

2. **Dataset 2 (DS2)** was used for testing and consists of 63 charted sections from the other 2 annotated macaques (M2 & M3). The macaques M2 and M3 have injections in the vlPFC and frontal pole, respectively. Hence, macaques M2 & M3 allow us to evaluate our method in a case with a similar injection site as the training case (vlPFC, slightly ventral to the injection site in M1) and a case with a different injection site (frontal pole).

All datasets were downsampled in-plane by a factor of 4 for training and testing.

2.2. Computer-assisted segmentation and characterization of fiber bundles

The aim of this work is to provide an end-to-end solution for segmentation and characterization of fiber bundles similar to the manually annotated examples shown in Figure 2.

The workflow of the proposed method is illustrated in Figure 3. The first step is the detection of bundles using an anatomy-constrained, self-supervised learning technique. The second step is the characterization of fibers within individual bundles by estimating their fiber density (FD) as they travel through different white-matter pathways.

Step 1: Automated detection of fiber bundles

A sample region with fiber bundles (magnified)

Photomicrograph of a coronal section

Step 2: Automated characterization of fibers within bundles in various regions in different pathways

A sample fiber bundle (magnified)

Grayscale

Segmented fibers within the bundle

Fiber density (**FD**)

Figure 3: **Proposed method workflow.** The first step involves automated detection of fiber bundles, and the second step the estimation of the fiber densities within bundles. The ROIs on photomicrographs of coronal sections are magnified at each step to show the individual bundles and fibers within them.

2.2.1. Step 1: Automated detection of fiber bundles

We train our detection method by using (1) an encoder-decoder architecture for segmenting fiber bundles, while simultaneously discriminating fiber bundles from background with a self-supervised contrastive loss, and (2) a temporal ensembling framework to efficiently use sections without manual charting from different brains.

Anatomy-constrained, self-supervised learning: We build a multi-tasking model as shown in Figure 4 by using a 2D U-Net (Ronneberger et al., 2015), which is one of the most successful architectures for medical image segmentation tasks (Panayides et al., 2020). The multi-tasking model consists of a U-Net backbone (F_{Seg}) for segmenting the fiber bundles and an auxiliary classification arm (F_{Class}) for discriminating fiber patches from background patches. We provide randomly sampled RGB patches of size $\times 256 \times 3$ as input. F_{Class} is connected to the bottleneck of the encoder of F_{Seg}, where the feature maps are passed through a downsampling module followed by two fully connected layers ($fc1024$, $fc256$) and an output layer with two nodes (fiber bundle vs background). The downsampling module consists of two max-pooling layers, each followed by two 3×3 convolution layers to extract high-level global features in the patches. We used focal loss (eqn. 1) for training F_{Seg},

because it handles class imbalance well (Lin et al., 2017). The focal loss is given by:

$$FL(p_t) = -\alpha_t(1 - p_t)^\gamma log(p_t), \quad p_t = \begin{cases} p, & \text{if } y = 1 \\ (1-p), & \text{otherwise} \end{cases} \quad (1)$$

where α and γ are weighing and focusing parameters, respectively, and $p \in [0,1]$ is the predicted probability for the fiber bundle class.

Figure 4: **Fiber bundle detection.** Network architecture used for segmentation, and the use of continuity priors from previous sections for false positive reduction.

The sparse nature of the anatomic tracing data makes it hard to draw exact boundaries. As described earlier, the main goal of the manual charting is to circumscribe areas that contain fibers traveling close to each other, which is a much harder task than annotating a contiguous structure with clear boundaries (e.g., caudate nucleus). Hence, the manual labels may sometimes not include all fiber regions. In addition, there may be significant texture variations and noise in the background. Therefore, to learn intrinsic texture/intensity variations in addition to the fiber features from the manual charting alone, we use a self-supervised technique for training F_{Class}. Specifically, we use a contrastive loss function based on SimCLR (Chen et al., 2020), where augmented data from each sample constitute the positive example to the sample while the rest are treated as negatives for the loss calculation. In SimCLR, augmentation

by random cropping and color distortions of the image patches were shown to perform well. In our case, we adapt this approach by choosing augmentations better suited to our problem: (i) random cropping of patches closer to the input patch ($< 20\mu$m), anatomically constrained within the white matter (by iterative sampling of patches until a mean intensity criterion is satisfied), and (ii) noise injection followed by Gaussian blurring (with a randomly chosen $\sigma \in [0.05, 0.3]$). The self-supervised loss with the above augmentations has two advantages: (1) effective separation between fiber and non-fiber background patches, and (2) identification of fiber patches even in the presence of artifacts, aided by the shared weights in the encoder of F_{Seg}. We used the contrastive loss (Chen et al., 2020) (eqn. 2) between positive pairs of patches (i, j) of F_{Class}, given by:

$$CL(i, j) = \frac{exp(sim(f_i, f_j)/\tau)}{\sum_{k=1}^{2N} I_{k \neq i} \, exp(sim(f_i, f_k)/\tau)}, \quad sim(x, y) = \frac{x^T y}{||x|| \, ||y||} \qquad (2)$$

where f is the output of F_{Class}, $sim(.)$ is the cosine similarity function, I is an indicator function such that $I_{k \neq i} = 1$ if $k \neq i$, else 0, and τ is the temperature parameter.

Temporal ensembling (TE) training: Data comes from 13 macaque brains, M1– M13, where only $\sim$6% of sections were manually charted. This would be insufficient for this challenging detection problem. Hence, after pretraining the model for Np epochs using the manually charted sections alone, we use the additional unlabeled sections for training both F_{Seg} and F_{Class}, with temporal ensembling (Perone et al., 2019). In this technique, predictions from the previous r epochs ($[P_{N-r}, ..., P_{N-1}]$) are averaged and thresholded to obtain the target label for the current epoch N (we empirically set $r = 3$). We use focal loss for pretraining the encoder-decoder of F_{Seg}, because contrastive loss is determined in a self-supervised manner in F_{Class}, and mainly used for learning inherent texture variations. For the first 3 epochs after pretraining, predictions from the pretrained model P_{N_p} are used for label generation. Averaging predictions reduces segmentation noise and aid in adapting the model to data from different brains.

Inference on test brain sections: We obtain predictions of fiber bundle labels by applying the segmentation part F_{Seg} of the model on the whole coronal sections (or patches of size 1024 × 1024 in the case of sections with dimensions larger than 1024 voxels).

Continuity prior for false positive removal: We take advantage of the spatial continuity of fiber bundles across consecutive sections to remove obvious FPs. Hence, we use the bundle segmentation from adjacent sections to inform the segmentation in the current section as follows: We downsample the sections by a factor of 10 and align the sections approximately along the center of ventricles (or along the lateral edges of the brain for sections without ventricles) to roughly form 3D histological volumes. We apply a triplanar U-Net architecture used in Sundaresan et al. (2021) to obtain a 3D low-res segmentation of main dense fiber bundles, which we then upsample to the original dimensions. For each section, we compute the average of the segmented fiber bundle masks from the two nearest neighboring sections (e.g., in rostral and/or caudal directions, if available) from the 3D segmentation. We remove any detected fiber bundle region in the current section if its distance from the averaged bundles of neighboring sections is >0.5mm.

In an additional post-processing step to reduce FPs due to noise, we reject predicted regions with area <2mm^2 and those near the brain outline (<0.5mm).

Implementation details: For training, we used the Adam optimizer (Kingma & Ba, 2014) ($\epsilon = 10^{-3}$), batch size = 8, pretraining epochs (N_p) = 100 and train with TE for 100 epochs. The convergence occurred at ∼90 epochs with early stopping using a patience value of 25 epochs. For focal loss, we use $\alpha = 0.25$; $\gamma = 2$. For contrastive loss, we used $\tau = 0.5$. The hyperparameters are chosen empirically. For F_{Seg}, we augment data using translation (offset ∈ [-50, 50] voxels), rotation ($\theta \in$ [-20^o, 20^o]), horizontal/vertical flipping and scaling ($s \in$ [0.9, 1.2]). The model is implemented using PyTorch 1.10.0 on Nvidia GeForce RTX 3090, taking ∼10 mins/epoch for ∼22,000 samples, with training:validation = 90:10.

2.2.2. Step 2: Automated characterization of fibers within bundles

We further characterize the segmented fiber bundles by estimating the density of fibers in each bundle. We binarize the image intensities within the boundaries of each segmented fiber bundle by enhancing the contrast with contrast-limited adaptive histogram equalization (Zuiderveld, 1994) and thresholding at the 95th percentile of intensity values. Example binary fiber maps are shown in Figure 3. We then calculate the fiber density (FD) as the number of voxels above the threshold over the total number of voxels in the bundle area.

2.3. Experiments

2.3.1. Experimental setup

We perform 5-fold cross-validation on 465 sections (440 unlabeled + 25 labeled) from DS1 with a training-validation-testing split ratio of 80-13-5 sections (for each fold, only manually charted sections are used for testing). We then train the model on DS1 and test it on the unseen dataset DS2 (sections from macaques different from the training one).

We also perform an ablation study of the method on DS2. This allows us to show the impact of different components of our architecture on bundle detection performance: (i) F_{Seg} with cross-entropy loss (CE loss), (ii) F_{Seg} with focal loss, (iii) F_{Seg} with addition of F_{Class} with contrastive loss (focal loss + ss_con loss), (iv) F_{Seg} and F_{Class} with TE (focal loss + ss_con loss + TE). We use the same postprocessing for all cases (i-iv), to isolate the effect of the above components.

The fiber bundles from each injection site reach their destinations by travelling through the large white-matter pathways such as internal capsule (IC), corpus callosum (CC) and uncinate fasciculus (UF). In certain pathways, fibers from the same injection site travel closely bundled with each other. In other pathways, fibers from different injection sites are intermingled. Thus, how tightly packed fibers from the same injection site remain as they travel through the white matter depends more on the pathway that they are traveling through than the injection site that they came from. We quantify this empirical observation by determining the density of fibers

within each of our segmented fiber bundle areas and comparing that density among areas that lie in three different WM pathways: IC, CC, UF.

For this analysis, the IC, CC, and UF are identified by an expert neuroanatomist. The predicted fiber bundle areas that overlap with one of these pathways by more than 40% are determined and fiber densities are obtained for these bundles using the method described in section 2.2.2.

2.3.2. Performance evaluation metrics

We define true positive (TP) bundles as the manually charted fiber bundles that are correctly predicted by the proposed method, and false positive (FP) bundles as the regions that are predicted as fiber bundles by our method but are not included in the manual chartings. We evaluate the performance of our method based on the following metrics:

1. True positive rate (TPR): number of TP bundles / total number of true bundles charted manually.

2. Average number of FPs (FP_{avg}): number of FP bundles / number of test sections.

3. Difference between the estimated fiber density of the manually charted and automatically detected bundles (Δ_{FD}).

3. Results

3.1. Cross-validation on DS1

Figure 5(a) shows free-response receiver operating characteristic (FROC) curves for fiber bundle detection using 5-fold cross-validation on DS1 (consisting of sections from the manually labelled brain M1 and 10 unlabeled brains M4 – M13). We obtain a TPR of 0.85 at 3.7 FPs/section at the elbow point (shown in dotted lines) for a threshold value of 0.4.

Figure 5: **Cross-validation on DS1.** (a) FROC curves for fiber bundle detection. (b) Boxplots of TPR and Δ_{FD} at $FP_{avg}{=}2$ FPs/section, after postprocessing.

Figure 5(b) shows the boxplots of TPR and Δ_{FD} values after postprocessing, with corresponding performance values reported in Table 1. We observe a significant reduction of FPavg (p $<$ 0.05) after postprocessing, mainly due to continuity constraints, for much lower changes in TPR values. Typically, FD ranged between $\sim$2-20%, and we obtained mean Δ_{FD} = -1.9%. The negative difference indicates slightly higher FD in the automatically segmented bundles than the manually charted ones, potentially due to a tighter fit of the area boundaries around the fibers in the former.

Table 1: **Cross-validation on DS1 and ablation study on DS2.** Mean and standard error values are reported. CE loss - Cross-entropy loss, Focal loss - F_{Seg} with focal loss, ss_con loss - F_{Class} with self-supervised contrastive loss, TE - temporal ensembling. (*) indicates significant improvements in the results compared to the previous row, determined using paired two-tailed T-tests. The best performance in the ablation study is highlighted in bold. $\uparrow$/$\downarrow$ indicate that higher/lower values lead to better results.

	TPR $\uparrow$		$\|\Delta_{FD}\|$ **(%)** $\downarrow$		**FP$_{avg}$** $\downarrow$	
	M1(injection site: vlPFC)		**M1(injection site: vlPFC)**		**M1(injection site: vlPFC)**	
5-fold cross-validation	0.83 $\pm$ 0.05		1.9 $\pm$ 0.09		2.0	
Ablation study						
	M2(injection site: vlPFC)	**M3(injection site: frontal pole)**	**M2(injection site: vlPFC)**	**M3(injection site: frontal pole)**	**M2(injection site: vlPFC)**	**M3(injection site: frontal pole)**
CE loss	0.68$\pm$0.04	0.69$\pm$0.05	4.8$\pm$0.12	5.2$\pm$0.18	5.5	7.5
Focal loss	*0.75$\pm$0.05	*0.76$\pm$0.04	*3.5$\pm$0.11	*4.0$\pm$0.17	*4.5	*4.0
Focal loss + ss_con loss	0.79$\pm$0.04	*0.81$\pm$0.03	*2.8$\pm$0.09	*3.3$\pm$0.14	*3.9	3.5
Focal loss + ss_con loss + TE	*0.81$\pm$0.04	0.82$\pm$0.03	*2.2$\pm$0.09	*3.0$\pm$0.13	*3.5	*2.5

3.2. Ablation study results on DS2

We train the method on dataset DS1 for ablation study cases (i-iv), test on DS2 (consisting of brains M2 and M3, different from those used for training) and use a threshold of 0.4 to obtain binary maps. We use the same postprocessing for all cases (i-iv) of the study.

Table 1 reports numeric results and Figures 6 and 7 show example images from the ablation study. Figure 6 shows fiber bundles in the CC and cingulum (a) and in the prefrontal white matter (b). The bundles in Figure 6(a) and 6(b) had been annotated, respectively, as dense and moderately dense. Figure 7 shows bundles in the prefrontal white matter (a) and in the IC (b).

Figure 6: **Ablation study on brain M2 (injection site in vlPFC).** (a, b) Sections with ROIs enlarged (white dotted box), show examples of bundles that had been manually annotated as dense (a) and moderately dense (b); (i – iv) Ablation study results on the ROIs with true positive, false positive and false negative bundles shown in yellow, red and blue outlines respectively (the proposed method highlighted in green box (iv)). Further enlarged ROIs (orange dotted box) containing fibers in the original RGB, grayscale and fiber binary maps.

As observed in the table, we obtain consistent performance trends in the ablation study between the M2 and M3 brains, which had different injection sites and hence different fiber trajectories. In both cases, CE loss (i) shows the worst performance.

Figure 7: **Ablation study on brain M3 (injection site in frontal pole).** (a, b) Sections with ROIs enlarged (white dotted box); (i – iv) Ablation study results on the ROIs with true positive, false positive and false negative bundles shown in yellow, red and blue outlines respectively (the proposed method highlighted in green box (iv)). Further enlarged ROIs (orange dotted box) containing fibers in the original RGB, grayscale and fiber binary maps.

From both Figures 6 and 7, among all the methods, experiments using focal loss (ii - iv) yield significantly better performance than CE loss (i), suggesting that focal loss is better at handling the heavy class imbalance. The self-supervised contrastive loss (ss_con loss) significantly increases TPR for M3 (injection site in frontal pole) and reduces FP_{avg} in M2 (injection site in vlPFC) due to the better discrimination between subtle variations in the background intensity and texture. We also observe a significant reduction in Δ_{FD} for focal loss + ss_con loss (iii) in both M2 and M3 brains, indicating more refined, tighter boundaries of fiber bundles. Hence, the contrastive loss not only reduces FPs, but also improves the segmentation of predicted regions. Using TE (iv) further improves detection, especially increasing the TPR of dense bundles and reducing FP_{avg}. The value of r (number of prior epochs to predict the target labels) in TE plays a crucial role in the reduction of prediction noise. We set r = 3 because it significantly reduces FP_{avg} over $r = 1$ (p < 0.01) but provides FP_{avg} values that are not significantly different from those with higher $r = 5$ ($p = 0.52$).

Corpus Callosum (CC)

M2: injection at vlPFC   M3: injection at Frontal Pole

Uncinate Fasiculus (UF)

M2: injection at vlPFC   M3: injection at Frontal Pole

Internal Capsule (IC)

M2: injection at vlPFC   M3: injection at Frontal Pole

Predicted bundles

- M2 (injection in vlPFC)
- M3 (injection in frontal pole)

Manually charted bundles

CC – Corpus callosum, UF – Uncinate fasiculus, IC – Internal capsule

Figure 8: **Fiber bundle characterization. Top panel:** examples of predicted fiber bundles from brain M2 (injection in vlPFC) and brain M3 (injection in frontal pole), in three different pathways: corpus callosum (CC), internal capsule (IC) and uncinate fasciculus (UF). **Middle and bottom panel:** Boxplots of fiber density (FD) for predicted and manually charted bundles in the IC, CC and UF, for brains M2 (injection in vlPFC – blue) and M3 (injection in frontal pole – orange).

3.3. Fiber densities in different WM pathways

Figure 8 shows examples of predicted fiber bundles in the CC, IC, and UF, for brains M2 (injection in vlPFC) and M3 (injection in frontal pole), along with boxplots of FD for the predicted and manually labelled bundles. Table 2 reports the mean and standard error of FD in these pathways for M2 and M3. The comparison of FD

between three white matter pathways (IC, CC, UF) shows that fibers from both the vlPFC injection (M2) and frontal pole injection (M3) are more densely packed in IC than in CC and UF. Also, from the boxplots in Figure 8, the interquartile range of FD is greater in CC and UF when compared to IC. We observe that FD is higher for predicted than manually labelled bundles, in almost all pathways. This indicates that predicted bundle areas have a tighter boundary around fibers than manually charted areas. However, the trend in FD across pathways is consistent between the predicted and manually charted areas, suggesting that they circumscribe a consistent amount of fibers. A one-way ANOVA, with FD of predicted bundles as the dependent variable and white-matter pathway (IC, CC, UF) as the independent variable, shows that the differences in FD are significant between pathways ($F = 26.1$, $p < 0.001$).

Table 2: **Fiber density (FD) by pathway.** Mean and standard errors of FD in corpus callosum (CC), internal capsule (IC) and uncinate fasciculus (UF) for sections from brain M2 (injection in vlPFC) and brain M3 (injection in frontal pole).

Pathways	FD (%)	
	M2	**M3**
CC	11.3 ± 0.15	11.1 ± 0.15
UF	11.1 ± 0.13	11.6 ± 0.17
IC	13.1 ± 0.12	13.1 ± 0.07

4. Discussion

We propose a method for computer-assisted segmentation and characterization of fiber bundles in histological sections from macaque brains that have received tracer injections. Our method does not require a large number of manually labeled sections ($<$ 10% of training data). We use a self-supervised, contrastive-learning technique with temporal ensembling that enables our model to leverage information from unlabeled sections during training, and to overcome intensity variations in histological sections and inconsistency in manually labeled boundaries. Our method achieves TPR > 0.80 on test sections from different macaque brains, one with a similar injection site as the manually labeled case used for training (vlPFC) and one with a different injection site (frontal pole) and hence different fiber trajectories. Given that we expect the

segmentation to always be inspected by an expert anatomist as a final step, a TPR of 0.8 represents an excellent starting point that will reduce the amount of manual intervention needed and thus accelerate the work of anatomists substantially. As more labeled cases become available with the use of our computer-assisted method, it will be possible to retrain our model and further improve its performance.

This is the first work on segmentation of fiber bundles in tracer data from macaque brains. The performance of our method compares favorably to prior work in marmoset brains, which reported a voxel-wise TPR of 0.7 (Woodward et al. 2020). The main sources of FPs in our method are terminal fields (shown in Figure 2) and artifacts such as glare or dust particles. In addition, FPs may occur along the white/gray matter interface, due to intensity variations. Figures 6-8 illustrate the typical variability in the intensity and contrast characteristics across sections from different brains. The use of continuity priors and ss_con loss was highly useful in reducing FPs and making our method more robust to this variability. The inclusion of the continuity priors in the training framework was not possible due to the lack of a sufficient number of manual chartings from consecutive sections. Hence, a future direction of this work, as more labeled cases become available, could explore integrating the priors within the training framework for further reduction of FPs. This may also reduce performance variation (indicated by standard deviation). Furthermore, with the availability of more labeled cases, it may be possible to train a model to detect fiber bundles and terminal fields as separate classes.

Typically, the detection improves and encloses more fibers, even without ss_con loss and TE, for densely packed fiber bundles as shown in Figure 6, because greater fiber density leads to greater texture differences between the fiber bundle and the background. These variations in the individual fiber bundle densities may impact performance differently between cases. For instance, in the case of M3 (tracer injected in the frontal pole), the fiber bundles in IC have relatively higher densities and less variations when compared to M2 (tracer injected at vlPFC) as shown in Figures 6 and 7. Sections in M2 especially show heavy contrast and texture variations in fibers

in the IC. Therefore, in our ablation study, we observe a greater improvement in M2 compared to M3 (Table 1, showing significant improvement in TPR in M2 using focal loss, ss_con loss and TE). This is because the initial predictive performance of the method just with focal loss is better for M3 compared to M2, due to the higher fiber density of the former. The proposed method with contrastive loss and TE equalizes performance across cases.

Our method lays the groundwork for accelerating the annotation of tracer data and building databases that contain not just the end points but the full trajectory of axon bundles. This will enable larger-scale studies of the topographic organization of axons within white-matter pathways (like the example of Figure 1), as well as more comprehensive validation of pathways reconstructed by dMRI or other, novel imaging modalities. It will also enable quantitative studies on the geometric properties of axon bundles (e.g., density, curvature, orientation dispersion) and how these properties vary among brain regions.

As an example of such a quantification, we compared the density of fibers projecting from two different injection sites (vlPFC, frontal pole) as they traveled through three white-matter pathways (IC, CC, UF). This allowed us to provide quantitative evidence for the empirical observation that fibers from the same injection site stay more tightly bundled when they travel through some pathways than others. In our results, fibers from the same injection site stayed close to each other as they traveled through the IC, but were more spread out (and presumably intermingled with fibers from other cortical areas) as they traveled through the CC and UF. The anterior limb of the IC is a narrow structure where fibers from the prefrontal cortex (PFC) are topologically organized based on their cortical origin (e.g., ventromedial, dorsomedial, or dorsolateral PFC) (Safadi et al., 2018; Lehman et al., 2011). The UF contains intertwined fibers running between the vlPFC and several destinations; some of these fibers follow the UF all the way to the temporal lobe, and others use the UF as a conduit to reach other white-matter pathways (Lehman et al., 2011). Our finding is particularly relevant for microstructure-informed tractography methods (Daducci

et al., 2016). Such methods assume that differences in microstructural properties between white-matter bundles can help disentangle the long-range trajectories of these bundles. This, however, may be less effective in areas where fibers from multiple origins are intermingled, rather than neatly organized in spatially separable bundles. Extending the quantitative analyses that we performed here to more injection sites and pathways will be important for shedding light on this issue.

5. Conclusion

We have developed a method for segmentation of fiber bundles that can greatly reduce the time needed for anatomists to annotate histological sections from anatomic tracing experiments. Facilitating the annotation of more cases in a semi-automated fashion will allow us to generate larger training datasets that can be used to further improve the performance of our method in the future. It will also enable larger-scale studies to validate tractography algorithms, or to extract quantitative information from tracing data and analyze the precise route and geometric properties of axon bundles across multiple seed regions. The code for our method is available at `https://github.com/v-sundaresan/fiberbundle_seg_tracing`.

Data and code availability

The python implementation of our proposed method is available at `https://github.com/v-sundaresan/fiberbundle_seg_tracing`.

Author contributions

V.S. contributed to methodology, software, validation, formal analysis, investigation and wrote the original draft, with inputs from all authors. S.H. and J.L. contributed to conceptualization, tracer data acquisition, histology data processing and manual annotations. C.M. contributed to investigation and analysis. A.Y. contributed to conceptualization, methodology, funding acquisition and project supervision. All authors contributed to manuscript revision, read, and approved the submitted version.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This work was supported by the National Institute of Mental Health to SNH (R01-MH045573, P50-MH106435), the National Institute of Neurological Disorders and Stroke to AY (R01-NS119911), and the National Institute of Biomedical Imaging and Bioengineering to AY (R01-EB021265).

References

Bakker, R., Wachtler, T., & Diesmann, M. (2012). Cocomac 2.0 and the future of tract-tracing databases. *Frontiers in neuroinformatics*, *6*, 30.

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597–1607). PMLR.

Daducci, A., Dal Palú, A., Descoteaux, M., & Thiran, J.-P. (2016). Microstructure informed tractography: pitfalls and open challenges. *Frontiers in neuroscience*, *10*, 247.

Dinsdale, N. K., Bluemke, E., Sundaresan, V., Jenkinson, M., Smith, S. M., & Namburete, A. I. (2022). Challenges for machine learning in clinical translation of big data imaging studies. *Neuron*, *110*, 3866–3881.

Du, J., Zhang, X., Liu, P., & Wang, T. (2023). Coarse-refined consistency learning using pixel-level features for semi-supervised medical image segmentation. *IEEE Journal of Biomedical and Health Informatics*, .

Feyjie, A. R., Azad, R., Pedersoli, M., Kauffman, C., Ayed, I. B., & Dolz, J. (2020). Semi-supervised few-shot learning for medical image segmentation. *arXiv preprint arXiv:2003.08462*, .

Fischer, M., Hepp, T., Gatidis, S., & Yang, B. (2023). Self-supervised contrastive learning with random walks for medical image segmentation with limited annotations. *Computerized Medical Imaging and Graphics*, *104*, 102174.

Geschwind, N. (1965). Disconnexion syndromes in animals and man. *Brain*, *88*, 585–585.

Grisot, G., Haber, S. N., & Yendiki, A. (2021). Diffusion mri and anatomic tracing in the same brain reveal common failure modes of tractography. *Neuroimage*, *239*, 118300.

Haber, S. N., Kim, K.-S., Mailly, P., & Calzavara, R. (2006). Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. *Journal of Neuroscience*, *26*, 8368–8376.

Haber, S. N., Lehman, J., Maffei, C., & Yendiki, A. (2023). The rostral zona incerta: A subcortical integrative hub and potential deep brain stimulation target for obsessive-compulsive disorder. *Biological Psychiatry*, *93*, 1010–1022.

Haber, S. N., Liu, H., Seidlitz, J., & Bullmore, E. (2022). Prefrontal connectomics: from anatomy to human imaging. *Neuropsychopharmacology*, *47*, 20–40.

Haynes, W. I., & Haber, S. N. (2013). The organization of prefrontal-subthalamic inputs in primates provides an anatomical substrate for both functional specificity and integration: implications for basal ganglia models and deep brain stimulation. *Journal of Neuroscience*, *33*, 4804–4814.

Jbabdi, S., Lehman, J. F., Haber, S. N., & Behrens, T. E. (2013). Human and monkey ventral prefrontal fibers use the same organizational principles to reach their targets: tracing versus tractography. *Journal of Neuroscience*, *33*, 3190–3201.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, .

Kötter, R. (2004). Online retrieval, processing, and visualization of primate connectivity data from the cocomac database. *Neuroinformatics*, *2*, 127–144.

Kremer, J. R., Mastronarde, D. N., & McIntosh, J. R. (1996). Computer visualization of three-dimensional image data using imod. *Journal of structural biology*, *116*, 71–76.

Lai, Z., Wang, C., Hu, Z., Dugger, B. N., Cheung, S.-C., & Chuah, C.-N. (2021). A semi-supervised learning for segmentation of gigapixel histopathology images from brain tissues. In *2021 43rd annual international conference of the IEEE engineering in Medicine & Biology Society (EMBC)* (pp. 1920–1923). IEEE.

Lehman, J. F., Greenberg, B. D., McIntyre, C. C., Rasmussen, S. A., & Haber, S. N. (2011). Rules ventral prefrontal cortical axons use to reach their targets: implications for diffusion tensor imaging tractography and deep brain stimulation for psychiatric illness. *Journal of Neuroscience*, *31*, 10392–10402.

Li, X., Yu, L., Chen, H., Fu, C.-W., Xing, L., & Heng, P.-A. (2020). Transformation-consistent self-ensembling model for semisupervised medical image segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, *32*, 523–534.

Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision* (pp. 2980–2988).

Maffei, C., Girard, G., Schilling, K. G., Aydogan, D. B., Adluru, N., Zhylka, A., Wu, Y., Mancini, M., Hamamci, A., Sarica, A. et al. (2022). Insights from the irontract challenge: Optimal methods for mapping brain pathways from multi-shell diffusion mri. *NeuroImage*, *257*, 119327.

Mesbah, R., McCane, B., & Mills, S. (2016). Deep convolutional encoder-decoder for myelin and axon segmentation. In *2016 International Conference on Image and Vision Computing New Zealand (IVCNZ)* (pp. 1–6). IEEE.

Naito, T., Nagashima, Y., Taira, K., Uchio, N., Tsuji, S., & Shimizu, J. (2017). Identification and segmentation of myelinated nerve fibers in a cross-sectional optical microscopic image using a deep learning model. *Journal of neuroscience methods*, *291*, 141–149.

Öngür, D., & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral cortex*, *10*, 206–219.

Panayides, A. S., Amini, A., Filipovic, N. D., Sharma, A., Tsaftaris, S. A., Young, A., Foran, D., Do, N., Golemati, S., Kurc, T. et al. (2020). Ai in medical imaging informatics: current challenges and future directions. *IEEE journal of biomedical and health informatics*, *24*, 1837–1857.

Perone, C. S., Ballester, P., Barros, R. C., & Cohen-Adad, J. (2019). Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. *NeuroImage*, *194*, 1–11.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18* (pp. 234–241). Springer.

Safadi, Z., Grisot, G., Jbabdi, S., Behrens, T. E., Heilbronner, S. R., McLaughlin, N. C., Mandeville, J., Versace, A., Phillips, M. L., Lehman, J. F. et al. (2018). Functional segmentation of the anterior limb of the internal capsule: linking white matter abnormalities to specific connections. *Journal of Neuroscience*, *38*, 2106–2117.

Schilling, K. G., Gao, Y., Stepniewska, I., Janve, V., Landman, B. A., & Anderson, A. W. (2019). Anatomical accuracy of standard-practice tractography algorithms in the motor system-a histological validation in the squirrel monkey brain. *Magnetic resonance imaging*, *55*, 7–25.

Skibbe, H., Watakabe, A., Nakae, K., Gutierrez, C. E., Tsukada, H., Hata, J., Kawase, T., Gong, R., Woodward, A., Doya, K. et al. (2019). Marmonet: a pipeline for automated projection mapping of the common marmoset brain from whole-brain serial two-photon tomography. *arXiv preprint arXiv:1908.00876*, .

Stephan, K. E., Kamper, L., Bozkurt, A., Burns, G. A., Young, M. P., & Kötter, R. (2001). Advanced database methodology for the collation of connectivity data on the macaque brain (cocomac). *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *356*, 1159–1186.

Sundaresan, V., Dinsdale, N. K., Jenkinson, M., & Griffanti, L. (2022). Omni-supervised domain adversarial training for white matter hyperintensity segmentation in the uk biobank. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)* (pp. 1–4). IEEE.

Sundaresan, V., Zamboni, G., Rothwell, P. M., Jenkinson, M., & Griffanti, L. (2021). Triplanar

ensemble u-net model for white matter hyperintensities segmentation on mr images. *Medical image analysis*, *73*, 102184.

Wei, D., Lee, K., Li, H., Lu, R., Bae, J. A., Liu, Z., Zhang, L., dos Santos, M., Lin, Z., Uram, T. et al. (2021). Axonem dataset: 3d axon instance segmentation of brain cortical regions. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24* (pp. 175–185). Springer.

Woodward, A., Gong, R., Abe, H., Nakae, K., Hata, J., Skibbe, H., Yamaguchi, Y., Ishii, S., Okano, H., Yamamori, T. et al. (2020). The nanozoomer artificial intelligence connectomics pipeline for tracer injection studies of the marmoset brain. *Brain Structure and Function*, *225*, 1225–1243.

Wu, H., Wang, Z., Song, Y., Yang, L., & Qin, J. (2022). Cross-patch dense contrastive learning for semi-supervised segmentation of cellular nuclei in histopathologic images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 11666–11675).

Yendiki, A., Aggarwal, M., Axer, M., Howard, A. F., van Walsum, A.-M. v. C., & Haber, S. N. (2022). Post mortem mapping of connectional anatomy for the validation of diffusion mri. *NeuroImage*, *256*, 119146.

Zaimi, A., Duval, T., Gasecka, A., Côté, D., Stikov, N., & Cohen-Adad, J. (2016). Axonseg: open source software for axon and myelin segmentation and morphometric analysis. *Frontiers in neuroinformatics*, *10*, 37.

Zaimi, A., Wabartha, M., Herman, V., Antonsanti, P.-L., Perone, C. S., & Cohen-Adad, J. (2018). Axondeepseg: automatic axon and myelin segmentation from microscopy data using convolutional neural networks. *Scientific reports*, *8*, 3816.

Zhao, X., Qi, Z., Wang, S., Wang, Q., Wu, X., Mao, Y., & Zhang, L. (2023). Rcps: Rectified contrastive pseudo supervision for semi-supervised medical image segmentation. *arXiv preprint arXiv:2301.05500*, .

Zuiderveld, K. (1994). Contrast limited adaptive histogram equalization. *Graphics gems*, (pp. 474–485).