



Mapping mesoscale axonal projections in the mouse brain using a 3D convolutional network

Drew Friedmann^{a,b,1}, Albert Pun^{a,b,1}, Eliza L. Adams^{a,c}, Jan H. Lui^{a,b}, Justus M. Kebschull^{a,b}, Sophie M. Grutzner^{a,b}, Caitlin Castagnola^a, Marc Tessier-Lavigne^a, and Liqun Luo^{a,b,2}

^aDepartment of Biology, Stanford University, Stanford, CA 94305; ^bHoward Hughes Medical Institute, Stanford University, Stanford, CA 94305; and ^cNeurosciences Graduate Program, Stanford University, Stanford, CA 94305

Contributed by Liqun Luo, March 23, 2020 (sent for review October 22, 2019; reviewed by Gregory S. X. E. Jefferis and Hanchuan Peng)

The projection targets of a neuronal population are a key feature of its anatomical characteristics. Historically, tissue sectioning, confocal microscopy, and manual scoring of specific regions of interest have been used to generate coarse summaries of meso-scale projectomes. We present here TrailMap, a three-dimensional (3D) convolutional network for extracting axonal projections from intact cleared mouse brains imaged by light-sheet microscopy. TrailMap allows region-based quantification of total axon content in large and complex 3D structures after registration to a standard reference atlas. The identification of axonal structures as thin as one voxel benefits from data augmentation but also requires a loss function that tolerates errors in annotation. A network trained with volumes of serotonergic axons in all major brain regions can be generalized to map and quantify axons from thalamocortical, deep cerebellar, and cortical projection neurons, validating transfer learning as a tool to adapt the model to novel categories of axonal morphology. Speed of training, ease of use, and accuracy improve over existing tools without a need for specialized computing hardware. Given the recent emphasis on genetically and functionally defining cell types in neural circuit analysis, TrailMap will facilitate automated extraction and quantification of axons from these specific cell types at the scale of the entire mouse brain, an essential component of deciphering their connectivity.

neural networks | axons | whole-brain | light-sheet microscopy | tissue clearing

Volumetric imaging to visualize neurons in intact mouse brain tissue has become a widespread technique. Light-sheet microscopy has improved both the spatial and temporal resolution for live samples (1, 2), while advances in tissue clearing have lowered the barrier to imaging intact organs and entire organisms at cellular resolution (3, 4). Correspondingly, there is a growing need for computational tools to analyze the resultant large datasets in three dimensions. Tissue clearing methods such as CLARITY and iDISCO have been successfully applied to the study of neuronal populations in the mouse brain, and automated image analysis techniques have been developed for these volumetric datasets to localize and count simple objects, such as cell bodies of a given cell type or nuclei of recently active neurons (5–8). However, there has been less progress in software designed to segment and quantify axonal projections at the scale of the whole brains.

As single-cell sequencing techniques continue to dissect heterogeneity in neuronal populations (e.g., refs. 9–11) and as more genetic tools are generated to access these molecularly or functionally defined subpopulations, anatomical and circuit connectivity characterization is crucial to inform functional experiments (12, 13). Traditionally, and with some exceptions (e.g., refs. 14, 15), “projectome” analysis entails qualitatively scoring the density of axonal fibers in manually defined subregions selected from representative tissue sections and imaged by confocal microscopy. This introduces biases from the experimenter, including which thin tissue section best represents a large and complicated three-dimensional (3D) brain region; how to bin axonal densities

into high, medium, and low groups; whether to consider thin and thick axons equally; and whether to average or ignore density variation within a target region. Additionally, it can be difficult to precisely align these images to a reference atlas (16). Volumetric imaging of stiff cleared samples and 3D registration to the Allen Brain Institute’s Common Coordinate Framework (CCF) has eliminated the need to select individual tissue sections (17). However, without a computational method for quantifying axon content, researchers must still select and score representative two-dimensional (2D) optical sections (18).

The automated identification and segmentation of axons from 3D images should circumvent these limitations. Recent application of deep convolutional neural networks (DCNNs) and Markov random fields to biomedical imaging have made excellent progress at segmenting grayscale computed tomography and MRI volumes for medical applications (19–23). Other fluorescent imaging strategies including light-sheet, fMOST, and serial two-photon tomography have been combined with software like TeraVR, Vaa3D, Ilastik, and NeuroGPS-Tree to trace and reconstruct individual neurons (24–28). However, accurate reconstruction often requires a technical strategy for sparse labeling. While reconstructions of sparsely labeled axons are informative for measuring projection variability among single cells, they are usually labor-intensive and underpowered when interrogating population-level statistics and rare collateralization patterns.

Significance

Simple, high-resolution methods for visualizing complex neural circuitry in 3D in the intact mammalian brain are revolutionizing the way researchers study brain connectivity and function. However, concomitant development of robust, open-source computational tools for the automated quantification and analysis of these volumetric data has not kept pace. We have developed a method to perform automated identifications of axonal projections in whole mouse brains. Our method takes advantage of recent advances in machine learning and outperforms existing methods in ease of use, speed, accuracy, and generalizability for axons from different types of neurons.

Author contributions: D.F., A.P., and L.L. designed research; D.F., A.P., E.L.A., J.H.L., J.M.K., and C.C. performed research; D.F., A.P., E.L.A., and M.T.-L. contributed new reagents/analytic tools; D.F., A.P., and S.M.G. analyzed data; and D.F. and L.L. wrote the paper.

Reviewers: G.S.X.E.J., Medical Research Council Laboratory of Molecular Biology; and H.P., Institute for Brain and Intelligence, and Southeast University.

The authors declare no competing interest.

Published under the PNAS license.

Data deposition: Example image files, tutorials, and associated code for data handling, model training, and inference, as well as weights from the authors’ best model are available at <https://github.com/AlbertPun/TRAILMAP>.

¹D.F. and A.P. contributed equally to this work.

²To whom correspondence may be addressed. Email: lluo@stanford.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1918465117/-DCSupplemental>.

First published May 1, 2020.

As one of the most successful current DCNNs, the U-Net architecture has been used for local tracing of neurites in small volumes and for whole-brain reconstruction of brightly labeled vasculature (29–33). To identify axons, a similarly regular structural element, we posited that a 3D U-Net would be well suited for the challenges of a much lower signal-to-noise ratio, dramatic class imbalance, uneven background in different brain regions, and difficult annotation strategy. In addition, whole-brain imaging of axons necessitates the inclusion of any artifacts that contaminate the sample. The paired clearing and analysis pipeline we present here mitigates the impact of autofluorescent myelin and nonspecific antibody labeling that interfere with the detection of thin axons. Our network, TrailMap (Tissue Registration and Automated Identification of Light-sheet Microscope Acquired Projectomes), provides a solution to all of these challenges. We demonstrate its generalization to multiple labeling strategies, cell types, and target brain regions. Alignment to the Allen Institute reference atlas allows for visualization and quantification of individual brain regions. We have also made available our best trained model (<https://github.com/AlbertPun/TRAILMAP>) such that any researcher with mesoscale projections in cleared brains can use TrailMap to process their image volumes or, with some additional training and transfer learning, adapt the model to their data.

Results

TrailMap Procedures. To generate training data, we imaged 18 separate intact brains containing fluorescently labeled serotonergic axons (*SI Appendix, Fig. S1A*). From these brains, we cropped 36 substacks with linear dimensions of 100 to 300 voxels (Fig. 1A). Substacks were manually selected to represent the diversity of possible brain regions, background levels, and axon morphology. As with all neural networks, quality training data are crucial. Manual annotation of complete axonal structures in three dimensions is difficult, imprecise, and laborious for each volume to be annotated. One benefit of the 3D U-Net is the ability to input sparsely labeled 3D training data. We annotated 3 to 10 individual XY planes within each substack, at a spacing of 80 to 180 μm between labeled slices (Fig. 1A). Drawings of axons even included single labeled voxels representing the cross-section of a thin axon passing through the labeled slice. In the same XY slice, a second label surrounding the axon annotation (“edges”) was automatically generated, and the remaining unlabeled voxels in the slice were given a label for “background.” Unannotated slices remained without a label as previously indicated.

Training data were derived from serotonergic axons labeled by multiple strategies to better generalize the network (*SI Appendix, Fig. S1A*). As serotonin neurons innervate nearly all forebrain structures, they provide excellent coverage for examples of axons in regions with variable cytoarchitecture and, therefore, variable background texture and brightness. Our focus on imaging intact

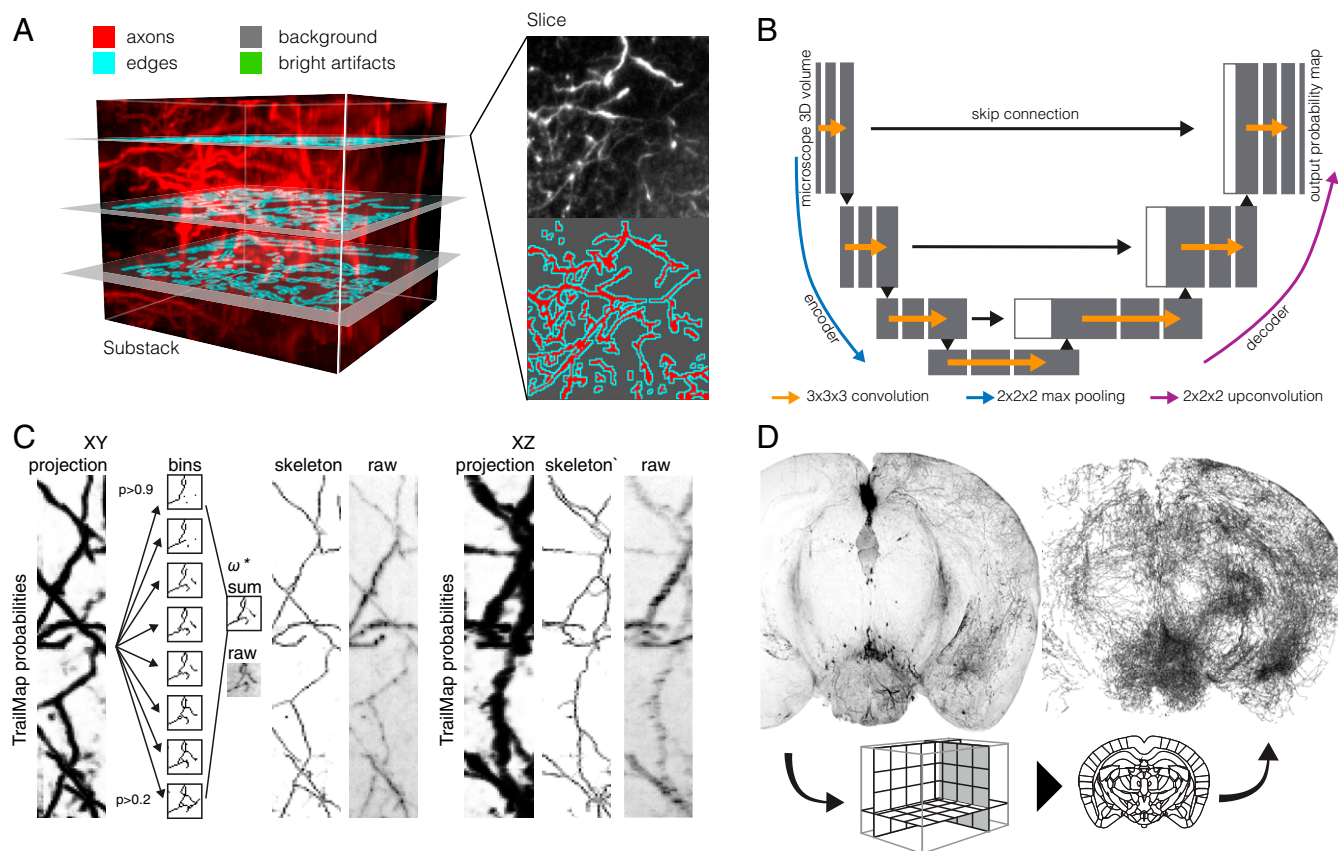


Fig. 1. Overview of the TrailMap workflow to extract axonal projections from volumetric data. (A) Annotation strategy for a single subvolume ($120 \times 120 \times 101$ voxels). Three planes are labeled with separate hand-drawn annotations for background, artifacts, and axons. The one-pixel-width “edges” label is automatically generated. (B) Basic architecture of the encoding and synthesis pathways of the 3D convolutional U-Net. Microscope volumetric data enters the network on the left and undergoes a series of convolutions and scaling steps, generating the feature maps shown in gray. Skip connections provide higher-frequency information in the synthesis path, with concatenated features in white. More details are provided in ref. 30. (C) A network output thinning strategy produces skeletons faithful to the raw data but with grayscale intensity reflecting probability rather than signal intensity or axon thickness. XY and XZ projections of one subvolume are shown ($122 \times 609 \times 122 \mu\text{m}$). (D) A 2-mm-thick volumetric coronal slab, before and after the TrailMap procedure, which includes axon extraction, skeletonization, and alignment to the Allen Brain Atlas Common Coordinate Framework.

brains required the inclusion of contaminating artifacts from the staining and imaging protocol since these nonspecific and bright artifacts are common in cleared brains and interfere with methods for axon identification. We addressed this in two ways. First, we implemented a modified version of the AdipoClear tissue clearing protocol (34, 35) that reduces the autofluorescence of myelin. As fiber tracts composed of unlabeled axons share characteristics with the labeled axons we aim to extract, this reduces the rate of false positives in structures such as the striatum (*SI Appendix, Fig. S1B*). Second, we included 40 substacks containing representative examples of imaging artifacts and nonspecific background and generated a fourth annotation label, “artifacts,” for these structures.

From this set of 76 substacks, we cropped and augmented 10,000 separate cubes of $64 \times 64 \times 64$ voxels to use as the training set. Our validation set comprised 1,700 cubes extracted from 17 separate substacks, each cropped from one of nine brains not used to generate training data. The U-Net architecture included two 3D convolutions with batch normalization at each layer, $2 \times 2 \times 2$ max pooling between layers on the encoder path, and $2 \times 2 \times 2$ upconvolution between layers on the decoder path (Fig. 1B). Skip connections provide information necessary for recovering a high-resolution segmentation from the final $1 \times 1 \times 1$ convolution. The final network was trained for 188 epochs over 20 h, but typically reached a minimum in the validation loss approximately a third of the way into training. Subsequent divergence in the training and validation F1 scores (*Materials and Methods*) indicated overfitting, and, as such, the final model weights were taken from the validation loss minimum (*SI Appendix, Fig. S1 C–E*).

For a given input cube, the network outputted a $36 \times 36 \times 36$ volume containing voxel-wise axon predictions ($0 < P < 1$). Large volumes, including intact brains, were processed with a sliding window strategy. From this output, a thinning strategy was implemented to generate a skeletonized armature of the extracted axons (Fig. 1C). Grayscale values of the armature were the weighted sum of 3D skeletons generated from binarization of network outputs. For visualizations, this strategy maintained axon continuity across low-probability stretches that would otherwise have been broken by a thresholding segmentation strategy. A separate benefit of this skeletonization strategy is that it treats all axons, thin and thick or dim and bright, equally for both visualization and quantification. A second imaging channel was used to collect autofluorescence, which in turn was aligned to the Allen Brain Atlas Common Coordinate Framework (CCF) via sequential linear and nonlinear transformation (<http://elastix.isi.uu.nl/>). These transformation vectors were then used to warp the axon armature into a standard reference space (Fig. 1D).

Comparisons with Random Forest Classifier. One of the most widely used tools for pixelwise classification and image segmentation is the random-forest-based software Ilastik (<https://www.ilastik.org/>). We compared axon identification by TrailMap with multiple Ilastik classifiers and found TrailMap to be superior (best TrailMap model, recall, 0.752; precision, 0.377; one-voxel exclusion zone precision, 0.790; best Ilastik classifier, recall, 0.208; precision, 0.661; one-voxel exclusion zone precision, 0.867), but most notably in the clarity of the XZ projection (Fig. 2A–C). The increase in TrailMap’s precision when excluding “edge” voxels suggested that many false positives are within one voxel of annotated axons. Ilastik examples used for comparison include classifiers trained in 3D with sparse annotations that include or exclude edge information as described for TrailMap, an example trained with images generated by iDISCO+, and a classifier trained in 2D using the exact set of labeled slices used to train TrailMap (*SI Appendix, Fig. S2A*). An alternative deep network architecture trained on similar 3D data did not properly segment axons in our dataset (*SI Appendix, Fig. S2 B, Left*). This is not

surprising given the different imaging paradigm used to generate training data (fMOST). However, TrailMap performed well on example data from fMOST imaging (*SI Appendix, Fig. S2 B, Right*), suggesting that it can generalize to other 3D data.

To better understand which aspects of the TrailMap network contribute to its function, we trained multiple models without key components for comparison. As with all neural networks, we identified a number of hyperparameters that affected the quality of the network output. The most important was the inclusion of a weighted loss function to allow for imperfect annotations of thin structures. The cross-entropy calculation for voxels defined as “edges” were valued $30\times$ less in calculating loss than the voxels annotated as axons. Manually labeling structures with a width of a single pixel introduced a large amount of human variability (*SI Appendix, Fig. S2C*), and deemphasizing the boundaries of these annotations allowed the network’s axon prediction to “jitter” without producing unnaturally thick predictions (*SI Appendix, Fig. S2D*). Relatedly, the total loss equation devalued background (7.5 \times) and artifact examples (1.875 \times) with respect to axon annotations—a set of weights that balances false positives and negatives and compensates for class imbalances in the training data (Fig. 2D and *SI Appendix, Fig. S2D*).

Details of Axonal Projections. To determine TrailMap’s ability to identify axons from intact cleared samples, we tested a whole brain containing axon collaterals of serotonin neurons projecting to the bed nucleus of stria terminalis. With the extracted axonal projectome transformed into the Allen Institute reference space, the axon armature could be overlaid on a template to better highlight their presence, absence, and structure in local subregions (Fig. 3A and *SI Appendix, Fig. S3*). However, it was difficult to resolve steep changes in density or local hotspots of innervation without selectively viewing very thin sections. By averaging axons with a rolling sphere filter, a small hyperdense zone is revealed in amygdala that would have been missed in region-based quantifications (Fig. 3B, arrow). Total axon content in functionally or anatomically defined brain regions could be quantified (*SI Appendix, Fig. S4*) (36) or otherwise projected and visualized to retain local density information in three dimensions (Fig. 3C and D and *Movie 1*).

An added benefit of the transformation into a reference coordinate system is that brain regions as defined by the Allen Institute could be used as masks for highlighting axons in individual regions of interest. As an example, the aforementioned density observed in amygdala is revealed to be contributed by two nearby dense zones in the basomedial and basolateral amygdala, respectively (Fig. 3E and *Movie 2*). These analyses of forebrain innervation by distantly located, genetically identified serotonergic neurons ensure certainty of axonal identity. However, in anatomical zones surrounding the labeled cell bodies in dorsal raphe, TrailMap does not distinguish among local axons, axons in passage, or intermingled dendritic arbors (*SI Appendix, Fig. S5A*). Without anatomical predictions, our imaging cannot morphologically distinguish these fiber types.

Generalization to Other Types of Neurons and Brain Regions. Given that TrailMap was trained exclusively on serotonergic neurons, it may not generalize to other cell types if their axons are of different sizes, tortuosity, or bouton density. The network may also fail if cells are labeled with a different genetic strategy or imaged at different magnification. However, our data augmentation provided enough variation in training data to produce excellent results in extracting axons from multiple additional cell types, fluorophores, viral transduction methods, and imaging setups. TrailMap successfully extracted brain-wide axons from pons-projecting cerebellar nuclei (CN) neurons retrogradely infected with *AAV_{retro}-Cre* and locally injected in each CN with *AAV-DIO-tdTomato* (Fig. 4A). Visualizing the whole brain (Fig. 4A)

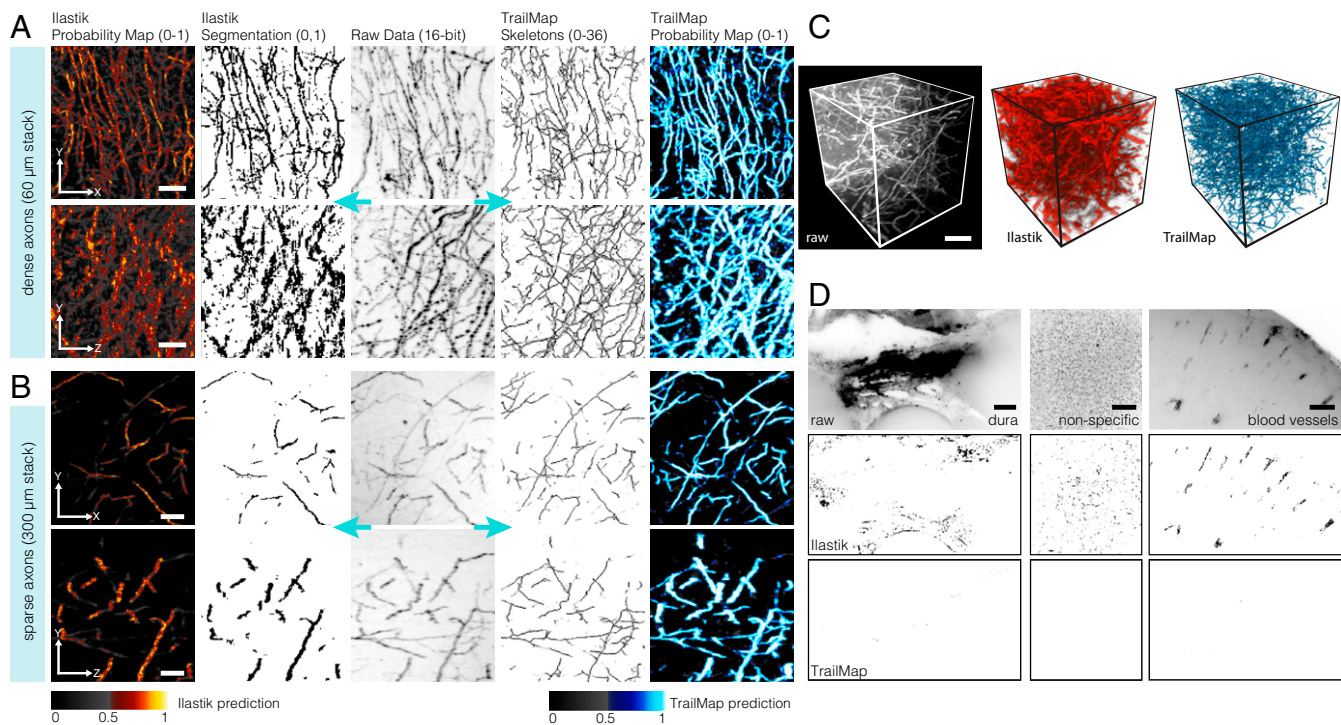


Fig. 2. Comparison of TrailMap to a random forest classifier. (A) From left to right: 60- μm Z-projection of the probability map output of an Ilastik classifier trained on the same 2D slices as the TrailMap network; typical segmentation strategy, at a $P > 0.5$ cutoff; raw data for comparison; skeletonized axonal armature extracted by TrailMap; and probability map output of the TrailMap network. To better indicate where $P > 0.5$, color maps for images are grayscale below this threshold and colored above it. (Scale bar, 100 μm .) Second row shows the same region as above, rotated to a X-projection. (B) Sparse axon identification by TrailMap and Ilastik; images show 300 μm of Z- or X-projection. (Scale bar, 100 μm .) (C) A 3D maximum intensity projection of a raw volume of axons, the resultant Ilastik probabilities, and TrailMap skeletonization. (Scale bar, 40 μm .) (D) Network output from examples of contaminating artifacts from nonspecific antibody labeling, similar to those included in the training set. Raw data, Ilastik segmentation, and TrailMap skeletonization are shown. (Scale bar, 200 μm .)

or the thalamus (Fig. 4B and Movie 3) in 3D reveals the contours of the contralateral projections of these neurons without obscuring information at the injection site or ipsilateral targets.

We also tested TrailMap on functionally identified thalamocortical axons in barrel cortex labeled by conditional *AAV-DIO-CHR2-mCherry* and TRAP2 activity-dependent expression of Cre recombinase (37). The higher density of local innervation of these axons necessitates imaging at higher magnification, resulting in a change in axon scale and appearance in the imaging volume. Buoyed by the spatial scaling included in the data augmentation, TrailMap reliably revealed the dense thalamic innervation of somatosensory layer IV and weaker innervation of layer VI (Fig. 4C). In contrast, though the Ilastik classifier performs moderately well on serotonergic axons, it fails to generalize to thalamocortical axons, perhaps owing to the change in scale of the imaging strategy (SI Appendix, Fig. S5B).

Finally, TrailMap also extracted cortico-cortical projection axons from prefrontal cortex (PFC) labeled by retrograde *CAV-Cre* and *AAV-DIO-mGFP-2A-synaptophysin-mRuby*. Axons from PFC neurons imaged in posterior visual and entorhinal cortices were identified with the exception of the most superficial axons in layer I (Fig. 4D). The failure to identify layer-I axons could be because the serotonergic training set did not include examples of superficial axons; as a result, the trained network used the presence of low-intensity grayscale values outside the brain to influence the prediction for each test cube containing the edge of the sample. Using 17 new training substacks from brains with annotated superficial axons from PFC cortical projections and 5 new validation volumes, we performed transfer learning using our best model as the initial weights. After just five epochs, the model successfully identified these layer-I axons (Fig. 4D).

While TrailMap's training and test data all originated from cleared mouse brains imaged by light-sheet microscopy, the network can also extract filamentous structures from 3D volumes imaged in other species and by other imaging methods. For example, variably bright neuronal processes in a *Drosophila* ventral nerve cord (38) imaged by confocal microscopy were all extracted by our model and equalized in intensity by our thinning method (SI Appendix, Fig. S5C). This normalization reveals the full coverage of these fibers, independent of visual distractions from bright cell bodies or thick branches. This same thinning strategy, applied to a single *Drosophila* antennal lobe interneuron's dendritic arborization (SI Appendix, Fig. S5D), simplified the cell's structure to a wireframe which could aid in manual single-cell reconstructions. Lastly, dye-labeled vasculature (39) in cleared mouse brains also yielded high-quality segmentations of even the smallest and dimmest capillaries (SI Appendix, Fig. S5E), suggesting TrailMap's utility extends beyond neuronal processes. Notably, all of these results were generated without any retraining of the network. With additional training data, we expect that TrailMap may apply to a wider range of cell types, species, and imaging modalities.

Discussion

Here we present an adaptation of a 3D U-Net tuned for identifying axonal structures within noisy whole-brain volumetric data. Our trained network, TrailMap, is specifically designed to extract mesoscale projectomes rather than reconstructions of axons from sparsely labeled individual neurons. For intact brains with hundreds of labeled neurons or zones of high-density axon terminals, we are not aware of a computational alternative that can reliably identify these axons.

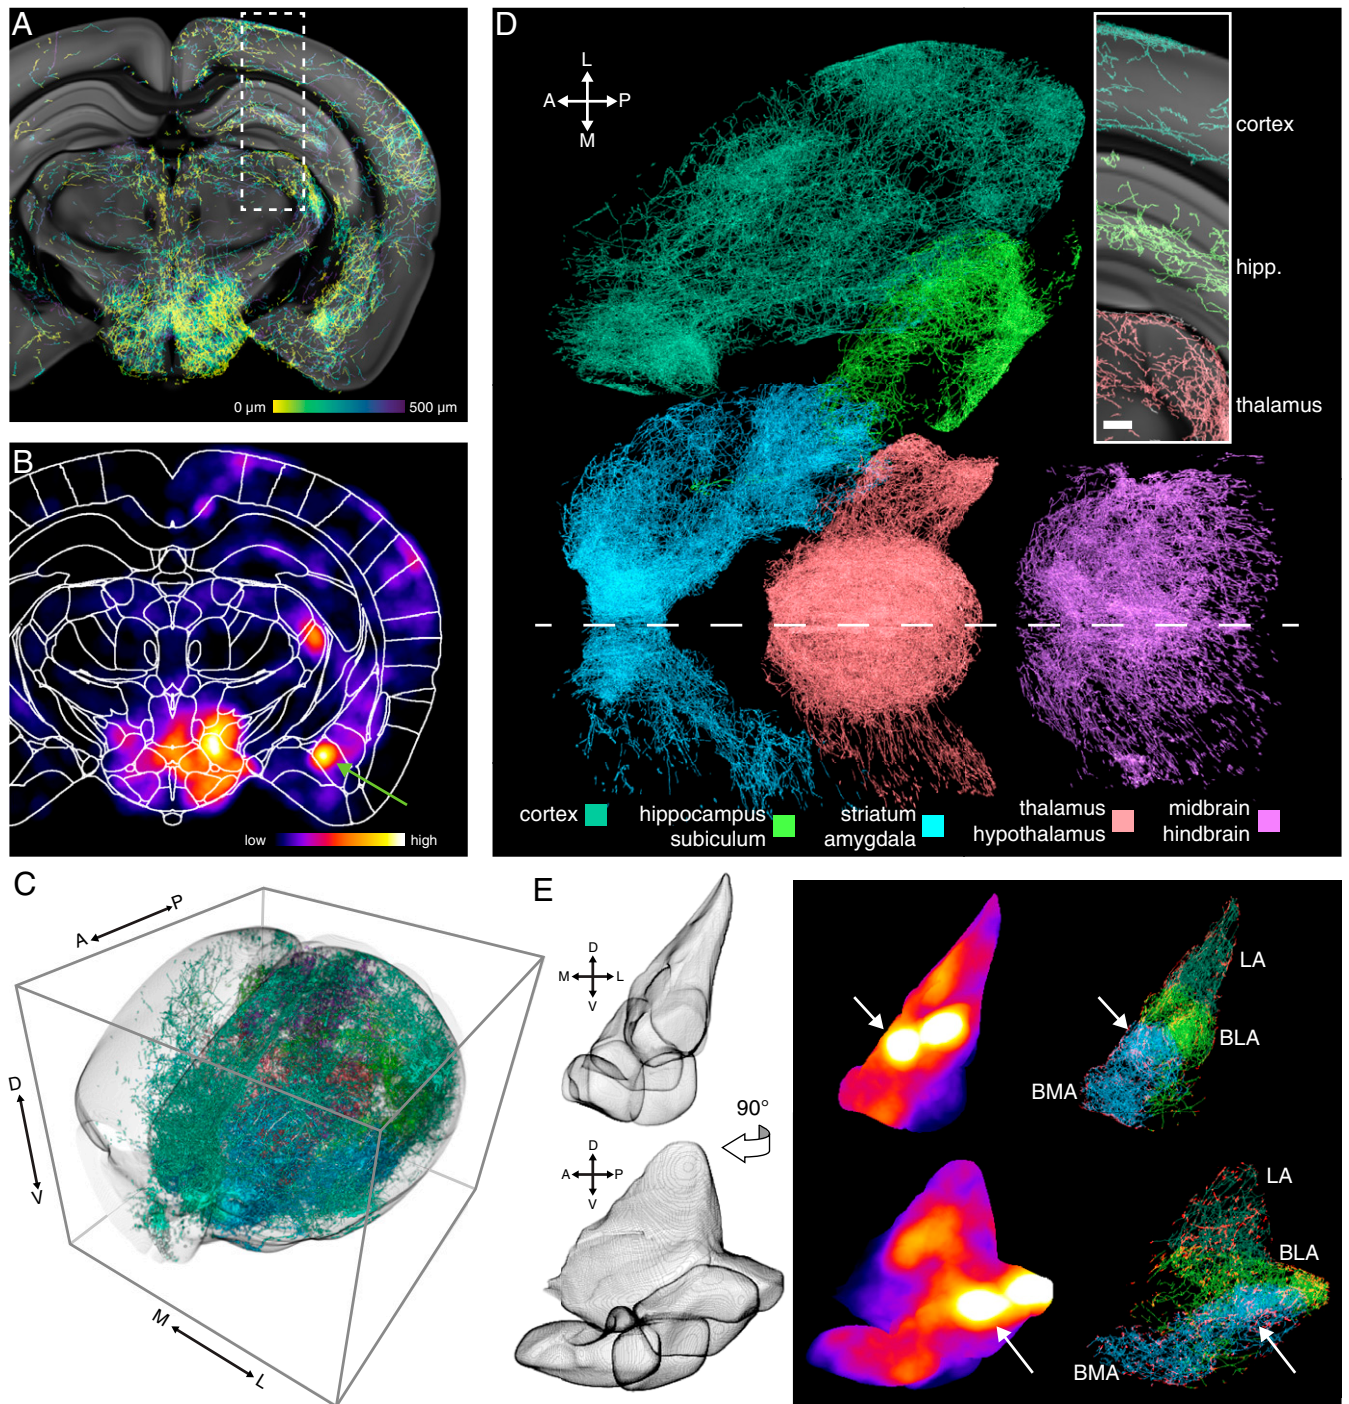


Fig. 3. Volumetric visualizations highlight patterns of axonal innervation. (A) Coronal Z-projection of extracted serotonergic axons, color-coded by depth (0–500 μm) and overlaid on the CCF-aligned serial two-photon reference atlas. *SI Appendix, Fig. S1* details viral-transgenic labeling strategy. (B) The same volumetric slab as in A presented as a density heatmap calculated by averaging a rolling sphere (radius = 225 μm). Green arrow highlights a density hotspot in the amygdala. (C) TrailMap-extracted serotonergic axons innervating forebrain are subdivided and color-coded based on their presence in ABA-defined target regions. (D) Same brain as in C as seen from a dorsal viewpoint, with major subdivisions spatially separated. Midline is represented by a dashed white line. *Inset* highlights the region indicated by the dashed box in A; Z-projection, 500 μm . (Scale bar, 200 μm .) (E, *Left*) Mesh shell armature of the combined structures of the lateral (LA), basolateral (BLA), and basomedial (BMA) amygdala in coronal and sagittal views. (*Right*) Density heatmap and extracted axons of the amygdala for the same views. LA, BLA, and BMA are color-coded in shades of green/blue by structure, and axonal entry/exit points to the amygdala are colored in red. White arrows highlight the same density hotspot indicated in B.

Our clearing and image processing pipeline address a number of challenges that have prevented these mesoscale analyses until now. First, all clearing techniques have at least some issues with background or nonspecific labeling that can interfere with automated

image analysis. Removing myelinated fiber tracts with a modified AdipoClear protocol greatly improved TrailMap's precision in structures such as the striatum. Second, our weighted loss function considers manually annotated, nonspecific, bright signals separately

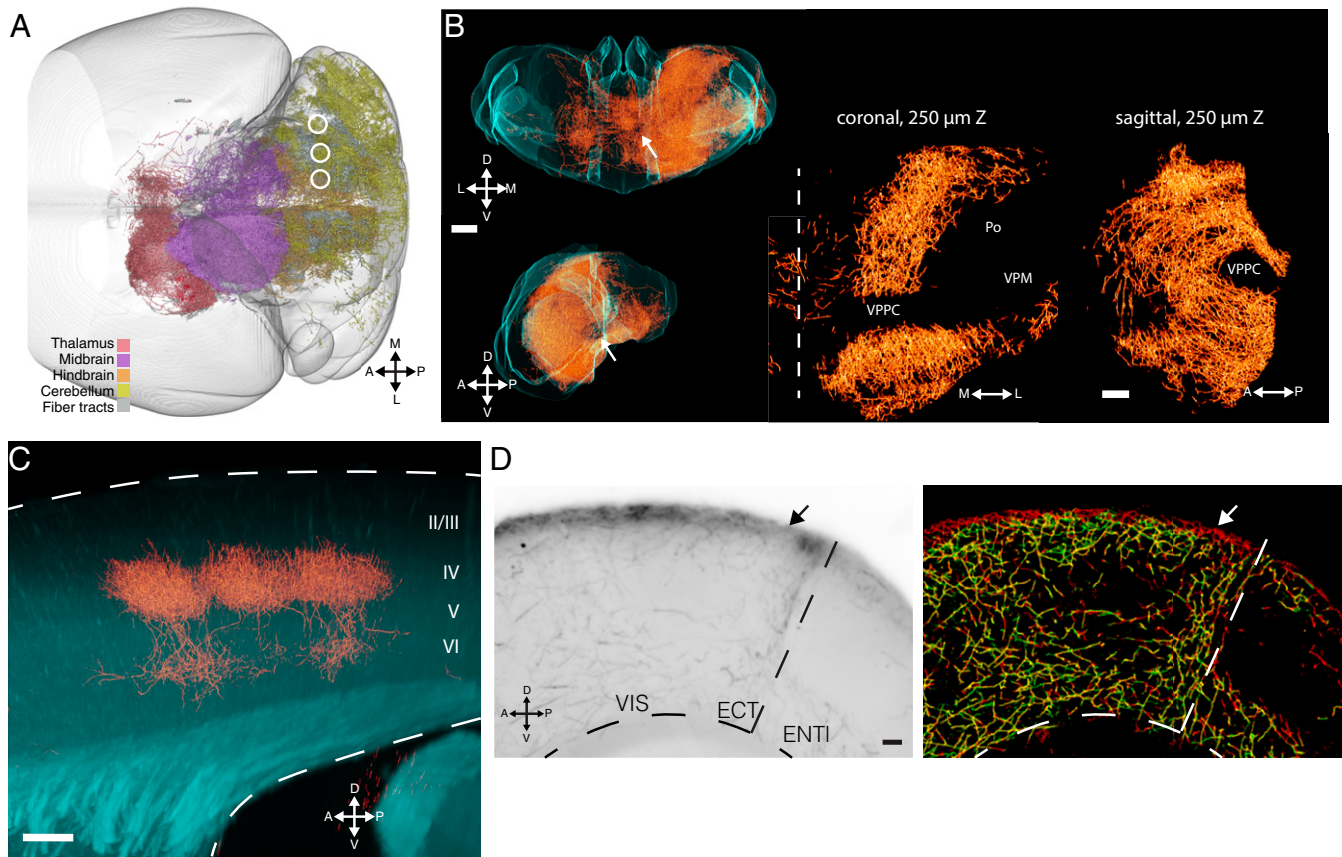


Fig. 4. Generalization to other cell types with and without transfer learning. (A) Dorsal view of posterior brain with extracted axon collaterals from pons-projecting cerebellar nuclei neurons color-coded by their presence in major subregions. Injection sites in the right hemisphere lateral, interposed, and medial CN from the same brain are indicated. (B) Coronal (Top) and sagittal (Bottom) views of the extracted axons (orange) within the structure of thalamus (cyan) from the brain in A. (Scale bar, 500 μm .) Zoomed images (Right) of midline thalamus show exclusion of axons from specific subregions. VPPC, ventral posterior nucleus, parvocellular part; VPM, ventral posteromedial nucleus; Po, posterior thalamic nuclear group. Midline shown as dashed vertical line. (Scale bar, 200 μm .) (C) Extracted thalamocortical axons in barrel cortex. XZ-projection of a 3D volume extracted from a flat-mount imaged cortex. Dashed lines indicate upper and lower bounds of cortex. Z-projection, 60 μm . (Scale bar, 100 μm .) (D, Left) Raw image of axons of prefrontal cortex neurons in posterior cortex. VIS, visual cortex; ECT, ectorrhinal; ENTI, lateral entorhinal. (Right) Axons extracted by TrailMap model before (green) and after (red) transfer learning. Z-projection, 80 μm . (Scale bar, 100 μm .) Labeling procedure, virus, and transgenic mouse for all panels are described in *SI Appendix, Fig. S1 and Methods*.

from other background areas—an essential step in reducing false positives. Relatedly, by devaluing the loss calculated for voxels adjacent to axon annotations, we reduced the rate of false negatives by allowing the network to err by deviations of a single voxel. Third, we present a strategy for thinning TrailMap’s output to construct an armature of predicted axons. One benefit of this thinning technique is a reduction in false breaks along dim axonal segments. Importantly, this benefit relies on the gradients in probability predicted by the cross-entropy-based loss calculation. The resultant armature improves visualizations and reduces biases in analysis and quantification by giving each axon equal thickness independent of staining intensity or imaging parameters.

Aligning armatures and density maps to the Allen Institute’s reference brain highlights which brain regions are preferentially innervated or avoided by a specific projection (36). Given that some brain regions are defined with less certainty, it will be possible to use the axon targeting of specific cell types to refine regional boundaries. TrailMap succeeds at separating thalamocortical projections to individual whisker barrels (Fig. 4C), a brain area with well-defined structure. Thus, it will be interesting to locate other areas with sharp axon density gradients that demarcate substructures within larger brain regions. These collateralization maps will also assist neuroanatomists investigating the efferent projection patterns of defined cell populations.

TrailMap has the added benefit of 3D, intact structures as the basis for quantification, but also the ability to process samples and images in parallel, reducing the active labor required to generate a complete dataset that spans the entire brain. TrailMap code is publicly available, along with the weights for our best model and example data. A consumer-grade GPU is sufficient for both processing samples and for performing transfer learning, while training a new model from scratch benefits from the speed and memory availability from cloud computing services. While 2D drawing tools are sufficient for generating quality training sets and reduce manual annotation time, we suspect that virtual reality-based visualization products would assist labeling strategies in three dimensions, potentially bolstering accuracy even further (24). We hope that TrailMap’s ease of use will lead to its implementation by users as they test brain clearing as a tool to visualize their neurons of interest. We expect that, as neuronal cell types are becoming increasingly defined by molecular markers, TrailMap can be used to map and quantify whole-brain projections of these neuronal types using viral-genetic and intersectional genetic strategies (e.g., ref. 36).

Materials and Methods

Animals. All animal procedures followed animal care guidelines approved by Stanford University’s Administrative Panel on Laboratory Animal Care (APLAC). Individual genetic lines include wildtype mice of the C57BL/6J and

CD1 strains, TRAP2 (*Fos-iCreER^{T2}*; Jackson, stock no. 030323), Ai65 (Jackson, stock no. 021875), and *Sert-Cre* (MMRRC, stock no. 017260-UCD). Mice were group-housed in plastic cages with disposable bedding on a 12-h light/dark cycle with food and water available ad libitum.

Viruses. Combinations of transgenic animals and viral constructs are outlined in *SI Appendix, Fig. S1A*. Viruses used to label serotonergic axons include *AAV-DJ-hSyn-DIO-HM3D(Gq)-mCherry* (Stanford vector core), *AAV8-ef1 α -DIO^{FR1}-loxP-STOP-loxP-mGFP* (Stanford vector core; ref. 40), *AAV-retro-CAG-DIO-Flp* (Salk Institute GT3 core; ref. 40), *AAV8-CAG-DIO-tdTomato* (UNC vector core, Boyden group), *AAV_{retro}-ef1 α -Cre* (Salk Institute GT3 core), *AAV8-ef1 α -DIO-CHR2-mCherry* (Stanford vector core), *AAV8-hSyn-DIO-mGFP-2A-synaptophysin-mRuby* (Stanford vector core, Addgene no. 71760), and *Cav-Cre* (Eric Kremer, Institut de Génétique Moléculaire, Montpellier, France; ref. 41).

AdipoClear Labeling and Clearing Pipeline. Mice were transcardially perfused with 20 mL 1 \times PBS containing 10 μ g/ μ L heparin followed by 20 mL ice-cold 4% PFA and postfixed overnight at 4 °C. All steps in the labeling and clearing protocol are on a rocker at room temperature for a 1-h duration unless otherwise noted. Brains are washed 3 \times in 1 \times PBS and once in B1n before dehydrating stepwise into 100% methanol (20, 40, 60, 80% steps). Two additional washes in 100% methanol remove all of the water before an overnight incubation in 2:1 dichloromethane (DCM):methanol. The following day, two washes in 100% DCM and three washes in methanol precede 4 h in a 5:1 methanol:30% hydrogen peroxide mixture. Stepwise brains are rehydrated into B1n (60, 40, 20% methanol), washed once in B1n, and then permeabilized 2 \times in PTxwH containing 0.3 M glycine and 5% DMSO. Samples are washed 3 \times in PTxwH before adding primary antibody (chicken anti-GFP, 1:2,000; Aves Labs; rabbit anti-RFP, 1:1,000; Rockland). Incubation is for 7 to 11 d, rocking at 37 °C, with subsequent washes also at this temperature. Brains are washed 5 \times in PTxwH over 12 h and then 1 \times each day for 2 additional days. Secondary antibody (donkey anti-rabbit, 1:1,000; Thermo; donkey anti-chicken, 1:2,000; Jackson) is incubated rocking at 37 °C for 5 to 9 d. Washes were performed 5 \times in PTxwH over 12 h and then 1 \times each day for 2 additional days. Samples are dehydrated stepwise into methanol, as before, but with water as the counterpart, then washed 3 \times in 100% methanol, overnight in 2:1 DCM:methanol, and in 2 \times 100% DCM the next morning. The second wash in DCM is extended until the brain sinks, before transfer to dibenzyl ether (DBE) in a fresh tube and incubation with rocking for 4 h before storing in another fresh tube of DBE at room temperature. Solutions were as follows: B1n, 1:1,000 Triton X-100, 2% wt/vol glycine, 1:10,000 NaOH 10N, 0.02% sodium azide; and PTxwH in 1 \times PBS, 1:1,000 Triton X-100, 1:2,000 Tween-20, 2 μ g/ μ L heparin, 0.02% sodium azide. This protocol diverges from the AdipoClear method subtly, by adding and extending the duration of exposure to DCM and by adjusting the duration of some steps. Most importantly, the overnight methanol/DCM steps and the hydrogen peroxide timing are crucial to image quality.

Light-Sheet Imaging. Image stacks were acquired with the LaVision Ultra-microscope II light-sheet microscope using the 2 \times objective at 0.8 \times optical zoom (4.0625 μ m per voxel, XY dimension). Thalamocortical axons were imaged at 1.299 μ m per voxel, XY dimension. Maximum sheet objective NA combined with 20 steps of horizontal translation of the light sheet improves axial resolution. Z-step size was 3 μ m. Axon images were acquired with the 640-nm laser, and a partner volume of autofluorescence of equal dimensions was acquired with the 488-nm laser. No preprocessing steps were taken before entering the TrailMap pipeline.

Datasets. All mouse brain datasets were generated in the Luo Lab as described earlier. *Drosophila* datasets were downloaded from public FlyCircuit (104198-Gal4) and FlyLight (R11F09-Gal4) repositories (38, 42). Mouse vasculature and associated ground truth annotation were downloaded from <http://discotechnologies.org/VesSAP> (39).

TrailMap Annotation Strategy. To create the training set for axons, we used the Fiji Segmentation Editor plugin to create sparse annotations in volumes of ~100 to 300 voxels per side. These were cropped from 18 samples across experimental batches from each of the three serotonergic neuron labeling strategies outlined in *SI Appendix, Fig. S1*. Two experts traced axons in these 36 different substacks by labeling single XY-planes from the stack every ~20 to 30 slices. Additionally, 40 examples of bright artifacts were found from these volumes and labeled as artifacts through an intensity thresholding method. Some examples contained both manually annotated axons and thresholded labels for artifacts. From these labeled volumes, 10,000 training examples were generated by cropping cubes (linear dimensions, 64 \times 64 \times

64) from random locations within each labeled substack. We also introduced an additional independent label, referenced as “edges.” This label was programmatically added to surround the manually drawn “axon” label in the single-voxel thick XY plane. This label was generated and subsequently given less weight in the loss function specifically to help the network converge by reducing the penalty for making off-by-one errors in voxels next to axons. Both edge and artifact labels are only used for weighting in the loss function, but are still counted as background. We created a validation set using the same methodology as the training set; however, to test for resilience against the potential impacts of staining, imaging, and other technical variation, the substacks used for the validation set were from different experimental batches than the training set.

Data Augmentation and Network Structure. Due to the simple cylindrical shape of an average axon segment, Z-score normalization was avoided as it removed the raw intensity information from the original image volume. Without this information, the network could not differentiate natural variability in the background from true axons. To provide robustness to signal intensity variation during training, chunks were augmented in real time through random scaling and summing random constants. We used a 3D-U-Net architecture (30) with input size 64³ and output of a 36³ segmentation. The output dimensions are smaller, accounting for the lack of adequate information at the perimeter of the input cube to make an accurate prediction. To segment large volumes of intact brains, the network was applied in a sliding window fashion.

We used a binary cross-entropy pixel-wise loss function, where axons, background, edges, and artifacts were given static weights to compensate for class imbalances and the structural nature of the axon. The binary loss function is calculated as $(-[\log(?) + (1-?)\log(1-?)]) * w$ for every voxel, with the predicted value (p), the true label (y), and a weight (w) determined from the true label's associated weight. The network was trained on Amazon Web Services using p3.xlarge instances for ~16 h. Training is ~2 \times slower, using 6 volumes per batch (as compared to 8 volumes per batch on AWS) with an NVIDIA GeForce 11GB 1080 Ti GPU. We trained many models with varying weights for the loss function and data scaling factor and picked the model that resulted in the lowest validation loss.

Model Evaluation. To evaluate the network, we compared the validation set output by the model to ground truth human annotations. Each pixel was classified to be TP (true positive), TN (true negative), FP (false positive), or FN (false negative). It is important to note that we did not include pixels which were labeled as edges because these were predetermined to be ambiguous cases and would not be an accurate representation of the network's performance. Using these four classes, we used the following formulas for metrics: precision = $TP / (TP + FP)$; recall = $TP / (TP + FN)$; F1 value = $2 * TP / (2 * TP + FP + FN)$; and Jaccard index = $TP / (FP + TP + FN)$. Human-human comparison was done by having 2 separate experts annotate the same 41 slices from each of 8 separate substacks before proceeding with the same set of evaluations.

Skeletonization and 3D Alignment to CCF. Probabilistic volumes from the output of TrailMap were binarized at eight separate thresholds, from 0.2 to 0.9. Each of these eight volumes was skeletonized in three dimensions (skimage-morphology-skeletonize_3d). The logical binary skeleton outputs were each weighted by the initial probability threshold used to generate them and subsequently were summed together. The resulting armature thus retains information about TrailMap's prediction confidence without breaking connected structures by threshold segmentation. Small, truncated, and disconnected objects were removed as previously described (40). We downsampled this armature from 4.0625 μ m per pixel (XY) and 3 μ m per pixel (Z) into a 5 \times 5 \times 5- μ m space and also downsampled the autofluorescence image into a 25 \times 25 \times 25- μ m space. The autofluorescence channel was aligned to the Allen Institute's 25- μ m reference brain acquired by serial two-photon tomography. These affine and bspline transformation parameters (modified from ref. 8) were used to warp the axons into the CCF at 5- μ m scaling. Once in the CCF, ABA region masks can be implemented for pseudocoloring, cropping, and quantifying axon content on a region-by-region basis. Rolling sphere voxelization with a radius of 45 voxels (225 μ m) operates by summing total axon content after binarization.

Data Availability. Example image files, tutorials, and associated code for data handling, model training, and inference, as well as the weights from our best model are all available at <https://github.com/AlbertPun/TRAILMAP>.

ACKNOWLEDGMENTS. We thank Ethan Richman for helpful discussions and comments on the manuscript. D.F. was supported by National Institute of Neurological Disorders and Stroke T32 NS007280, J.H.L. by National Institute of Mental Health K01 MH114022, J.M.K. by a Jane

Coffin Childs postdoctoral fellowship, and E.L.A. by the Stanford Bio-X Bowes PhD Fellowship. This work was supported by NIH Grants R01 NS104698 and U24 NS109113. L.L. is a Howard Hughes Medical Institute investigator.

1. M. B. Bouchard *et al.*, Swept confocally-aligned planar excitation (SCAPE) microscopy for high speed volumetric imaging of behaving organisms. *Nat. Photonics* **9**, 113–119 (2017).
2. T.-L. Liu *et al.*, Observing the cell in its native state: Imaging subcellular dynamics in multicellular organisms. *Science* **360**, eaaq1392 (2018).
3. P. Ariel, A beginner's guide to tissue clearing. *Int. J. Biochem. Cell Biol.* **84**, 35–39 (2017).
4. C. Pan *et al.*, Shrinkage-mediated imaging of entire organs and organisms using uDISCO. *Nat. Methods* **13**, 859–867 (2016).
5. D. S. Richardson, J. W. Lichtman, Clarifying tissue clearing. *Cell* **162**, 246–257 (2015).
6. K. Chung *et al.*, Structural and molecular interrogation of intact biological systems. *Nature* **497**, 332–337 (2013).
7. W. Menegas *et al.*, Dopamine neurons projecting to the posterior striatum form an anatomically distinct subclass. *eLife* **4**, e10032 (2015).
8. N. Renier *et al.*, Mapping of brain activity by automated volume analysis of immediate early genes. *Cell* **165**, 1789–1802 (2016).
9. B. Tasic *et al.*, Shared and distinct transcriptomic cell types across neocortical areas. *Nature* **563**, 72–78 (2018).
10. A. Zeisel *et al.*, Molecular architecture of the mouse nervous system. *Cell* **174**, 999–1014.e22 (2018).
11. A. Saunders *et al.*, Molecular diversity and specializations among the cells of the adult mouse brain. *Cell* **174**, 1015–1030.e16 (2018).
12. L. Luo, E. M. Callaway, K. Svoboda, Genetic dissection of neural circuits: A decade of progress. *Neuron* **98**, 256–281 (2018).
13. Q. Sun *et al.*, A whole-brain map of long-range inputs to GABAergic interneurons in the mouse medial prefrontal cortex. *Nat. Neurosci.* **22**, 1357–1370 (2019).
14. S. W. Oh *et al.*, A mesoscale connectome of the mouse brain. *Nature* **508**, 207–214 (2014).
15. B. Zingg *et al.*, Neural networks of the mouse neocortex. *Cell* **156**, 1096–1111 (2014).
16. D. FÜRTH *et al.*, An interactive framework for whole-brain maps at cellular resolution. *Nat. Neurosci.* **21**, 139–149 (2018).
17. Allen Institute for Brain Science, *Allen Mouse Common Coordinate Framework and Reference Atlas* (Allen Institute for Brain Science, Washington, DC, 2017).
18. M. Schneeberger *et al.*, Regulation of energy expenditure by brainstem GABA neurons. *Cell* **178**, 672–685.e12 (2019).
19. M. Alegro *et al.*, Automating cell detection and classification in human brain fluorescent microscopy images using dictionary learning and sparse coding. *J. Neurosci. Methods* **282**, 20–33 (2017).
20. M. Dong *et al.*, "3D Cnn-based soma segmentation from brain images at single-neuron resolution" in *2018 25th IEEE International Conference on Image Processing (ICIP) IS.* (IEEE, Piscataway, New Jersey, 2018), pp. 126–130.
21. P. Frasconi *et al.*, Large-scale automated identification of mouse brain cells in confocal light sheet microscopy images. *Bioinformatics* **30**, i587–i593 (2014).
22. B. Mathew *et al.*, Robust and automated three-dimensional segmentation of densely packed cell nuclei in different biological specimens with Lines-of-Sight decomposition. *BMC Bioinformatics* **16**, 187 (2015).
23. K. Thierbach *et al.*, "Combining deep learning and active contours opens the way to robust, automated analysis of brain cytoarchitectonics" in *Machine Learning in Medical Imaging*, Shi Y, Suk H-I, Liu M, Eds. (Springer, Basel, Switzerland, 2018), pp. 179–187.
24. Y. Wang *et al.*, TeraVR empowers precise reconstruction of complete 3-D neuronal morphology in the whole brain. *Nat. Commun.* **10**, 3474 (2019).
25. H. Peng, Z. Ruan, F. Long, J. H. Simpson, E. W. Myers, V3D enables real-time 3D visualization and quantitative analysis of large-scale biological image data sets. *Nat. Biotechnol.* **28**, 348–353 (2010).
26. J. Winnubst *et al.*, Reconstruction of 1,000 projection neurons reveals new cell types and organization of long-range connectivity in the mouse brain. *Cell* **179**, 268–281.e13 (2019).
27. T. Guan *et al.*, NeuroGPS-Tree: Automatic reconstruction of large-scale neuronal populations with dense neurites. *Nat. Methods* **13**, 51–54 (2016).
28. Z. Zhou, H.-C. Kuo, H. Peng, F. Long, DeepNeuron: An open deep learning toolbox for neuron tracing. *Brain Inform.* **5**, 3 (2018).
29. O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation. arXiv:1505.04597 (May 15, 2015).
30. Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, O. Ronneberger, 3D U-Net: Learning dense volumetric segmentation from sparse annotation. arXiv:1606.06650 (June 21, 2016).
31. J. Gornet *et al.*, Reconstructing neuronal anatomy from whole-brain images. arXiv:1903.07027 (March 17, 2019).
32. T. Falk *et al.*, U-Net: Deep learning for cell counting, detection, and morphometry. *Nat. Methods* **16**, 67–70 (2019).
33. A. P. Di Giovanna *et al.*, Whole-Brain vasculature reconstruction at the single capillary level. *Sci. Rep.* **8**, 12573 (2018).
34. J. Chi *et al.*, Three-Dimensional adipose tissue imaging reveals regional variation in beige fat biogenesis and PRDM16-dependent sympathetic neurite density. *Cell Metab.* **27**, 226–236.e3 (2018).
35. A. Branch, D. Tward, J. T. Vogelstein, Z. Wu, M. Gallagher, An optimized protocol for iDISCO+ rat brain clearing, imaging, and analysis. bioRxiv:10.1101/639674 (September 17, 2019).
36. J. Ren *et al.*, Single-cell transcriptomes and whole-brain projections of serotonin neurons in the mouse dorsal and median raphe nuclei. *eLife* **8**, e49424 (2019).
37. L. A. DeNardo *et al.*, Temporal evolution of cortical ensembles promoting remote memory retrieval. *Nat. Neurosci.* **22**, 460–469 (2019).
38. A. Jenett *et al.*, A GAL4-driver line resource for Drosophila neurobiology. *Cell Rep.* **2**, 991–1001 (2012).
39. M. I. Todorov *et al.*, Automated analysis of whole brain vasculature using machine learning. bioRxiv:10.1101/613257 (April 18, 2019).
40. J. Ren *et al.*, Anatomically defined and functionally distinct dorsal raphe serotonin sub-systems. *Cell* **175**, 472–487.e20 (2018).
41. L. A. Schwarz *et al.*, Viral-genetic tracing of the input-output organization of a central noradrenergic circuit. *Nature* **524**, 88–92 (2015).
42. A. S. Chiang *et al.*, Three-dimensional reconstruction of brain-wide wiring networks in Drosophila at single-cell resolution. *Curr. Biol.* **21**, 1–11 (2011).