



Published in final edited form as:

Nat Mach Intell. 2020 October ; 2(10): 585–594. doi:10.1038/s42256-020-0227-9.

Semantic segmentation of microscopic neuroanatomical data by combining topological priors with encoder-decoder deep networks

Samik Banerjee¹, Lucas Magee², Dingkan Wang², Xu Li¹, Bing-Xing Huo¹, Jaikishan Jayakumar³, Katherine Matho¹, Meng-Kuan Lin¹, Keerthi Ram³, Mohanasankar Sivaprakasam³, Josh Huang¹, Yusu Wang², Partha P. Mitra¹

¹Cold Spring Harbor Laboratory, NY, USA 11724.

²Computer Science and Engineering Department, The Ohio State University, Columbus, OH, USA 43210.

³Center for Computational Brain Research, Indian Institute of Technology, Chennai, Tamil Nadu, India 600036.

Abstract

Understanding of neuronal circuitry at cellular resolution within the brain has relied on neuron tracing methods which involve careful observation and interpretation by experienced neuroscientists. With recent developments in imaging and digitization, this approach is no longer feasible with the large scale (terabyte to petabyte range) images. Machine learning based techniques, using deep networks, provide an efficient alternative to the problem. However, these methods rely on very large volumes of annotated images for training and have error rates that are too high for scientific data analysis, and thus requires a significant volume of human-in-the-loop proofreading. Here we introduce a hybrid architecture combining prior structure in the form of topological data analysis methods, based on discrete Morse theory, with the best-in-class deep-net architectures for the neuronal connectivity analysis. We show significant performance gains using our hybrid architecture on detection of topological structure (e.g. connectivity of neuronal processes and local intensity maxima on axons corresponding to synaptic swellings) with precision/recall close to 90% compared with human observers. We have adapted our architecture to a high performance pipeline capable of semantic segmentation of light microscopic whole-brain image data into a hierarchy of neuronal compartments. We expect that the hybrid architecture incorporating discrete Morse techniques into deep nets will generalize to other data domains.

Author Contributions

The idea of using topological priors in the pipeline was conceptualized by Y.W. and P.M. Algorithmic design and development was performed by S. B. and L. M.. Proof reading assistance and neuroanatomical expertise for neuroanatomical ground truth data was provided by J. J. and K. M.. Data preparation, including quality control and acquisition were performed by B. H., J. J. and K. M. under the supervision of J. H. and P. M.. The ALBU baseline was tested by D. W.. Evaluation of the algorithm was conducted by S. B., D. W., L. M. and X. L.. Data preparation, including design of an online proofreading interface and hosting was done by A. L., M. S. and K. R.. S. B., L. M., J. J. and P. M. prepared and edited the paper.

Competing Interests Statement

The authors declare no competing interests.

Code Availability

The Process Detection and Semantic Segmentation code for reproduction of results in the paper and limited data with their links and documentation are available on Github at https://github.com/samik1986/ML_Semantic_Segmentation_NMI.

Introduction

Understanding the morphology and the connectivity of neurons is an important step in determining the neural circuitry of the brain. A primary method used for this purpose in larger vertebrate brains involves sparsely labelling individual neurons or groups of neurons using neuronal tracer injections. These methods have been the gold standard in neuroanatomical circuit tracing studies because they directly visualize neurons with high SNR and no inferential process is needed. These data sets have traditionally been painstakingly examined under a microscope by expert neuroanatomists, or more recently using digital microscopy, but still employing human labor intensive methods requiring close and time-consuming interactions with a human expert. Modern advances in optical imaging have allowed the digitization of whole-brain data sets, either 2D image stacks or 3D image volumes, with subcellular resolution¹⁻³. These imaging techniques have been incorporated into integrated neurohistology pipelines to generate large-scale, high-resolution data sets in a high-throughput fashion⁴⁻⁶. The resulting multi-Terabyte to Petabyte scale data sets are not possible to study using purely manual detection and quantification. Thus there is a pressing need for efficient automated algorithms to achieve this goal, with high precision and recall suitable for scientific data analysis.

Machine Learning techniques have facilitated the processing and analysis of large neuroanatomical data sets. In particular, computer algorithms for detecting neuronal fibers from various images with high accuracy have become increasingly important in automating computational characterization of neuronal morphology and circuitry. Previous works include reconstructions of neuronal morphology⁷, neuronal data analysis with an emphasis on reconstructing neurons from EM data cubes⁸; digital reconstruction of the 3D morphology of neurons from image stacks⁹. Recent advancements include a large-scale automated server-based biomedical-image analysis in FARSIGHT¹⁰, a free and open-source toolkit of image analysis methods for quantitative studies of complex and dynamic tissue micro-environments. The BigNeuron project (<https://alleninstitute.org/bigneuron/>) is a community effort to advance the state of the art of single-neuron reconstruction¹¹. There is also a large literature on the related but broader field of semantic segmentation of biomedical image data, including histopathological data, which we briefly review below.

In this manuscript, we develop a computational framework for a systematic treatment of the semantic segmentation problem for neuroanatomical image data. The basis of our framework is a hierarchical set of semantic categories suitable for neuroanatomical images (Figure 1). A key component of this framework is a method for automated segmentation of neurites (axons and dendrites), combining topological data analysis, based on Discrete Morse (DM) theory (Figure Extended Data 1), together with encoder-decoder Deep Net architectures to detect and quantify neuronal processes and boutons in whole brain histology sections. We call this method DM++ (Figure 2). Usage of the Morse-theoretic topological prior allows us to retain the fitting flexibility of deep-learning frameworks, while at the same time incorporating non-trivial prior structure. The resulting method shows both quantitative and qualitative improvements over the existing state of the art encoder-decoder networks for the same task (Table 1). Although neurite detection is a key step in the

semantic segmentation task, there are also other components to the framework (Figure 3, Figure Supplementary 5) to detect other relevant compartments such as somata, or to further sub-categorize detected neurites (Figure 5). We have deployed our methodology into a semi-automated data analysis pipeline including human proofreaders to obtain high quality semantic segmentation of the neuroanatomical image data (Figure Supplementary 2), suitable for further neuroscientific analysis.

Background

Automated computer vision algorithms have been used extensively for the delineation of anatomical structures and other regions of interest in specific radiological and histopathological images. These image segmentation algorithms, play an important role in numerous biomedical-imaging applications, such as the quantification of tissue volumes¹², diagnosis¹³, localization of pathology¹⁴, study of anatomical structure¹⁵, treatment planning¹⁶, and computer-integrated surgery¹⁷.

Prior to the popularization of Deep CNN methods¹⁸⁻²⁰, several machine learning and image processing techniques were available for medical image segmentation²¹. These include amplitude segmentation based on histogram features²², the region based segmentation²³, and the graph-cut approach²⁴. However, semantic segmentation approaches that utilize DL have become popular in recent years in the field of medical image segmentation, lesion detection, and localization²⁵. DCNN architectures have generally provided state-of-the-art performance for image classification²⁶, segmentation²⁷, detection and tracking²⁸, and captioning²⁹ in standardized datasets. Efficient optimization techniques are available for training DCNN models²⁶. However, in most cases, models are explored and evaluated on classification tasks in large-scale datasets like ImageNet²⁶, where the outputs of the classification tasks are a single label or probability value for a given image. Alternatively, smaller fully-connected convolutional neural networks (FCN)³⁰ and Segnet³¹ have been used successfully for semantic image segmentation tasks. State-of-the-art DCNN architectures for image segmentation tasks like U-Net³² and ALBU³³ (Figure 2A) employ an encoder-decoder architecture. While these networks are effective they employ double the number of model parameters thus raising resource issues.

Deep Learning techniques rely on very large volumes of annotated training data and may in some sense be training-data interpolation techniques³⁴. To address the lack of large annotated corpora for new tasks, data transformation or augmentation techniques^{32, 35, 36} including data whitening, rotation, translation, and scaling have been applied to increase the number of labeled samples available. The class imbalance problem is often solved by using tile based approaches rather than entire image sets³⁷. However, these kinds of methodological fixes do not ultimately address the weakness of DL techniques in lacking domain-specific priors, such as the topological prior that can capture a better connectivity of the neuronal architecture in the brain for our study.

To incorporate topological prior structure suitable to neurites, we employ Discrete Morse Graph Reconstruction (Figure Extended Data 1), a method that has been developed and studied thoroughly over the past several years³⁸⁻⁴¹. This algorithm combines persistent

homology⁴² and Discrete Morse Theory⁴³ to extract underlying graph structures from density fields. The algorithm considers the global structure of the data rather than looking strictly at local information and can handle many negative qualities seen in imperfect data, such as noise and non-uniform sampling. The method has been used in many applications, such as reconstructing road networks from GPS traces⁴¹ and extracting filament structures from simulated dark matter density fields⁴⁴.

Our Approach

In this paper, we propose a comprehensive framework for the semantic segmentation tasks for light-microscopic neuroanatomical data and compare with state of the art approaches. Gigapixel sized images are handled by breaking up into tiles of manageable size. The problem of class imbalance, when the pixel-based approach adapted for the tiles is systematically scaled to the entire image, is reduced by using tiles with signals for training and assigning a label for the image backgrounds and a target class for the fore-ground during semantic segmentation. In addition, two methods including binary cross-entropy loss and dice similarity were used for efficient training of classification and segmentation tasks^{35, 36}.

To detect neuronal processes, we propose a hybrid CNN based architecture DM++ (Figure 2B) which incorporates a topological prior using discrete Morse theory to give additional weight to quasi-one dimensional connected structures such as neurites. DCNNs do not have any built in mechanisms that can exploit the topological connectivity of neurites. For example, variations in the data acquisition process may cause some parts of a neurite to be less intensely labelled than other parts. This "signal gap" may pose challenges for a DCNN which might suppress the label on the weakly labelled part of a neurite (refer to the leftmost tile in the middle row for each modality in Figure 4). Nevertheless, prior knowledge that neurites form connected branches of tree-shaped neurons allow a human observer to easily trace through such regions of low intensity. We employ the Discrete Morse technique to incorporate such global connectivity information, as shown in the rightmost tile in the middle row for each modality in Figure 4.

We tested the proposed DM++ algorithm on three different types of image modalities used to image sections of the brain. Examples of the brain section images are shown in the top panels of columns A-C in figure 1, in each case from brains in which neuronal tracer injections have been used to visualize a subset of neurons and their processes. Figure 1(A) shows colored (fluorescent) Whole Slide Images (WSI). Figure 1(B) shows the Bright Field Images (BFI) of a brain section with an immunohistological stain. Both these images were captured using a digital slide scanner (0.46 μ m/pixel). Figure 1(C) shows monochrome images captured using Serial Two Photon Tomography (STPT). The second row (Column 1 A-C) shows examples of labelled neuronal cell bodies (soma) in these different modalities. The next third row shows neuronal processes labelled by red Fluorescent protein (Column 1 A), immunohistochemistry using the ABC-DAB process (Column 1 B) and high intensity eGFP label from the STPT data set (Column 1 C). The bottom row (Column 1 A-C) show examples of putatively terminating neuronal axons with the synaptic swellings/boutons. The cartoons in the second, third and fourth row (Column 1 A-C) beside each image tiles (the right subcolumn) shows masks corresponding to the object categories of interest in each of

the tiles shown in the left subcolumn. The semantic category hierarchy is detailed in the Column 1 D.

In the following sections, we first elaborate the pipeline for process detection, followed by the the rest of the semantic category tree. Subsequently we present results, followed by discussion and conclusions.

Neuronal Process detection using DM++ Architecture

The DCNN component of the DM++ pipeline is based on the ALBU network³³ which is in turn based on U-Net³². Figure 2(A) illustrates U-net and ALBU. 2(C) shows a diagrammatic representation of our proposed architecture DM++ for process detection. A more extensive description can be found in the Methods Section.

Briefly, the DM++ network accepts an image containing relevant objects of interest (e.g. processes) as input, tiles the image, processes these tiles using three processing paths forming a Y-shaped network, and outputs a binary mask indicating the detected neuronal processes. Figure 2(B) shows the three paths, (i) A "topological path" for detecting connected neural processes in the tile, (ii) a "DCNN path" (ALBU) for detecting the processes (iii) the final "common path" for combining the DCNN and topological paths.

Semantic Segmentation

Brain section images not only contain information about one object of interest (processes as detected by the DM++ pipeline) but also other objects such as dendrites, soma. Axonal processes themselves can show heterogeneity depending on their terminating nature. To account for these detections, we have incorporated the mask-RCNN architecture²⁷. This modification enables us to determine the exact pixels belonging to a sub-process or somata. The detections for the sub-categories of processes use the process detection masks generated by the DM++ algorithm. The architecture of Mask-RCNN, developed on FasterRCNN with an instance segmentation module, takes in the output neuronal mask from the DM++, to detect the axon arbors, passing axons and dendrites (for detailed discussion on the Mask-RCNN architecture, please refer to the Methods section). In contrary, the cell detection is performed by the semantic segmentation pipeline independent of the neuronal processes, though the boutons are detected as a sub-category using the output vertices of the graph from the DM reconstruction algorithm within the DM++ pipeline. Thus, the training set for the soma detection on the Mask-RCNN architecture takes in the raw image as input, while the others, viz. dendrites, passing axons and axons arbors, which are the sub-categories of process, are trained on the neuronal mask obtained from DM++. A brief discussion on the bouton detection is given in the methods section. A pictorial representation of the the overall flow diagram for the semantic segmentation methodology determined for the neuronal segmentation with different class labels and the cell detection has been depicted in figure 3 (B), while figure 3 (A) shows the architecture for the Mask-RCNN used to generate masks for each of the sub-categories of the neuronal processes as well as cell detection.

Qualitative and quantitative evaluations

The outputs from the proposed algorithms of neuronal process detection and semantic segmentation discussed above have been evaluated both qualitatively and quantitatively. This section provides a discussion of these evaluations, when compared with other state-of-the-art techniques. The proposed methodology shows improved performance metrics, and in addition also shows the ability to better detect connected process fragments. Finally, we also discuss the time complexities of the proposed pipeline, in terms of training and testing, followed by a discussion on the improvement in annotation time induced by our auto-detects.

Neuronal Process detection

The results are presented in table 1. The table shows the Precision, Recall and F1-measure for the different modalities. We have compared DM++ with an Unsupervised, UNet and ALBU for neuronal process detection. Each of these techniques are discussed in the Methods section. As evident from the sub-tables, our proposed technique outperforms the state of the art architectures across performance measures.

Figure 4 shows example results for different image modalities. Of particular interest is portions of neuronal processes connected by regions of low signal intensity, where DM++ is able to connect the processes through these low-intensity regions, but ALBU does not do so. In each of the three modalities, the middle row depicts a small patch from the image, where the intensity captured by the digital scanners used for imaging these sections is significantly low, but the connectivity can be detected in the high resolution image by visual inspection (refer to the leftmost patch in the middle row of results for each modality).

In each of these modalities, we also show comparisons with the manually annotated data (demarcated as "Annotation" in the figure) with TP (yellow), FP (magenta) and FN (cyan). The rectangular box indicates the location for each image from where the exemplar patch is extracted and the dotted lines establishes the correspondences of each path to the middle row of results for precise inspection. This kind of preservation of connectivity was evident throughout the sections used for testing. This is consistent with the expectation that the topological prior would improve the detection of connected processes.

To evaluate the extension of our framework to 3D, we extended the 2D DM++ pipeline to 3D in a natural manner, and performed a 3D Process Detection task on volumetric fMOST imaging data collected for single-neuron tracing. We compared the results with 3D UNet as a baseline and found that 3D DM++ outperformed the 3D UNet baseline. Details are presented in a supplementary section.

Semantic Segmentation

The semantic segmentation pipeline illustrated in figure 3 was used for detecting the sub-neuronal classes and somata. Performance metrics are presented in table 1. The numbers in bold indicate the results from our overall architecture, compared with a state-of-the-art encoder-decoder based CNN architecture, SegNet³¹ (details in Methods section).

Figure 5 illustrates each of the classes discussed in figure 1. Results for Cell Detection are shown in figure 5(A), which includes the MBA and the BFI dataset. The different sub-compartments of neuronal processes are displayed for the MBA dataset in the subsequent figures. Figures 5(B-D) illustrate the results of the semantic segmentation pipeline for different sub-categories of neuronal processes. The output from the DM++ algorithm is provided as input to the MaskRCNN network, as shown in the top tile. The output mask from the segmentation pipeline is shown just below, while the bottom row shows the performance of the MaskRCNN when compared to the manually annotated ground truth. Precision-Recall-F1 scores are calculated and reported in table 1.

Since the DM++ method specifically improves the connectivity of the detected neural processes, we developed a performance metric to directly measure connectivity in the detected processes as compared to the ground truth. In brief, the metric gives the fraction of point-pairs that are connected in the ground truth, that are also connected in the detected output (please see the Methods section for details). Notably, this connectivity-based performance metric showed a large $\sim 30\%$ advantage of the DM++ method over its nearest competitors depicted in table 1. We thank an anonymous reviewer for the suggestion that we develop such a metric

As an additional usage of the Discrete Morse approach, we detect putative "boutons" on axons by finding intensity maxima along the detected processes. The resulting "bouton" detections are shown in 5(E). The first image is an input tile, the second image shows what our method outputs as putative boutons, and the third image identifies whether these are true positives or false positives (compared with expert judgment), as well as where false negatives occur. Modified precision, recall, and F1 scores for three tiles that were manually annotated are included in 1. These modified scores were obtained with using a radius of 5, meaning a bouton labeled by our method would be labeled as a true positive if the nearest true bouton was within a neighborhood of 5 pixels (2.5μ).

Computational metrics:

We trained the DM++ network for 500 epochs with the training data. The training time was 200 seconds per epoch. For process detections, the algorithm took 4 minutes per mouse brain section (each section is a $24K \times 24K$ image with $0.46\mu\text{m}/\text{pixel}$ resolution) with processing times of 12% pre-processing, 15% neuronal pathway, 48% Topological pathway, 15% combined pathway and 10% post-processing, on a dual GPU (NVIDIA GTX2080 Ti) machine. The topological pathway was run in parallel on a 36 core CPU with 40 parallel threads. We used the pyTorch library for the python code for the neuronal pathway and Keras with TensorFlow backend for the common pathway. We used Keras with Tensorflow backend for the MaskRCNN network for cell detection and mask-generation. The Cell Detection MaskRCNN network was trained for 200 epochs, with a training time of ~ 1000 seconds per epoch. The detects took an average of 3 minutes per section for the cell (soma) detections on the mouse brain dataset. The architecture used for mask generation follows the one-vs-all strategy and approximately takes 3 minutes on an average for the three classes combined into an unified framework. There is a significant improvement in the proof-reading times for the manual annotations. Based on the report of an experienced

neuroanatomist, the proof-reading time for the process detection is no more than 25% compared to annotating the processes denovo. A similar trend was seen for cell-detection where we can reduced the annotation time by ~70%. In each case, proofreading was done with equally high quality criteria as de-novo manual annotation.

Discussion and Conclusions

We demonstrate significant performance improvements for an image segmentation task over state-of-the-art encoder-decoder networks through the incorporation of prior topological structure using methods from Morse theory (Table 1 and Figures 4 & 5). We show that such priors may be incorporated using an intuitively plausible Y-shaped ("Siamese") network architecture we entitle DM++. We further incorporate this process-detection network as an important component in an integrated framework for semantic segmentation of light microscopic neuroanatomical image data.

While we present significant performance improvements over state of the art, we are not yet at the point where the automation is at a sufficiently high quality level for scientific purposes, and a human-in-the loop component is required (although the trained networks lead to a large reduction of the required human labor).

Our post-hoc analysis of the data shows that abrupt changes in the dynamic intensity range within a data set is challenging for DM++, so that training datasets containing a full spectrum of intensity dynamic range is required. To verify robustness and transitivity of the algorithm, we carried out a set of experiments spanning imaging modalities and also studied the effect of test and training sets being drawn from different brains. The results (details in the supplementary materials) show that the separating test and training sets across brains from different animals causes only a minimal drop (1–4%) in performance when we test our algorithm within the same contrast and imaging modality. We also find that the DM++ algorithm shows good performance across imaging modalities (fluorescent whole-slide imaging, serial two-photon tomography, fluorescent micro-optical sectioning tomography), fluorescent tracer colors (red and green AAV tracers, Fast Blue retrograde tracers) and species (mouse, marmoset). We also studied a 3D version of DM++ and found it to provide performance improvements over 3D UNet.

We also tested our algorithm on image tiles near tracer injection sites, which show abrupt changes in the intensity levels, e.g., high background with auto-fluorescence, noise, etc. There was a larger drop in performance levels of approximately 8 – 10%. In our current implementation we separate these tracer-injection regions for analysis, but this shows the room for further improvements in our algorithmic framework.

We found however that our networks iteratively increase performance based on feedback from manual curation, which can be used to repeatedly fine-tune the network to improve generalization. Despite these residual shortcomings, our algorithm (DM++) incorporating topological priors, using discrete Morse theory in combination with DCNNs, significantly outperforms existing techniques for the detection and segmentation of neuronal objects of interests in light-microscopic neuroanatomical data.

Methods

This section describes methods used in this study including the fundamentals of the algorithm design, the architecture, training, testing and evaluation of our technique when compared to the existing conventional methods. We have compared our pipeline to existing techniques include an unsupervised process detection technique, U-Net and ALBU for process detection and SegNet for semantic segmentation, which are also discussed at the end of the Methods.

Discrete Morse Graph Reconstruction:

DM++ takes masks of the Discrete Morse Graph Reconstruction output as input. To generate this, we first need a density function over the domain to feed into the Discrete Morse framework. Image intensity in the color channel relevant to the tracer serves as the density function after applying a Gaussian filter to smooth out the image. We can interpret our density function as the higher the pixel value, the more likely it is to be part of true neuron signal.

Next we input the density function into the Discrete Morse Graph reconstruction framework. If one were to imagine the density function graphed on top of the domain, the Discrete Morse framework outputs the "mountain ridges" of the function, connecting the maxima through intervening saddle points. These ridges are made up of "flow lines" (also known as formally integral lines are curves whose local tangents follow the direction of steepest descent, *i.e.*, the gradient direction, of the density function) between the maxima and saddle points of the function. The reason for selecting these is that the function values are locally maximal along the ridges. Thus if we were to move slight off of a ridge, the probability of being part of the neuronal process decreases.

Noise Removal: Each ridge is assigned a persistence value⁴² which can be interpreted as an importance score. The noise-removal framework adopts a persistence threshold and filters out ridges below that threshold. We provide a low threshold value in order to remove ridges caused by noise while minimizing the number of ridges removed that actually mark true neuron signal.

Given the Morse output, we then create a mask for the domain. The mask is grayscale - where each path in the Discrete Morse output is assigned a constant value. The output of this stage is a simple grayscale image which can then be used at the next step of the DM++ algorithm. The higher the intensity along the path, the higher the gray scale value. After the persistence refined Discrete Morse graph reconstruction step, this mask, along with the original imaging data, is an input to the process detection pipeline (explained under the subsequent subsection DM++).

The standard output of discrete Morse graph reconstruction is a connected graph. In applications where the final output is built directly off of the Morse graph, false edges, such as those along the boundary, will need to be removed. However, in our DM++, the advantages of both ALBU and discrete Morse are combined during the co-training stage: In

particular, ALBU output helps remove such false positives from the discrete Morse output, while the discrete Morse output helps improve the connectivity of the ALBU output.

DM++:

The combination of the topological component with a DCNN convolutional network is aimed at capturing the connected nature of the neuronal processes. The "Siamese" architecture of the DM++ network takes the outputs from the discrete Morse algorithm and the ALBU architecture and combines them using a "Siamese" stacking into two channels. The common path emerging out of the Siamese stack combines features of the convolutional network along with the topological priors of the Discrete Morse.

Architecture: The architecture of DM++ consists of four components: the topological path, the ALBU path, the Siamese stack and the common path. The DCNN paths in this network consist of two-dimensional convolutional layers with ReLU activations, except the final layers in each which have Sigmoid activations. A dropout layer of 50% was introduced in each DCNN stack for regularization purposes. The loss function used in this network is a combined loss of soft dice loss³⁶ and binary cross entropy optimized using the Adam optimizer.

Inputs: The inputs to the DM++ architecture are the likelihood maps obtained from the discrete Morse algorithm and the ALBU models. The input images are assumed to be single channel images with the resolution of 512×512 pixels with varying pixel resolutions across different modalities of the images. The input image is passed through the discrete Morse algorithm to obtain a likelihood image containing the connected graphs with their edge weights as their likelihood. The ALBU model also produces output for the likelihood maps generated for segmentation of the processes.

Outputs: The output of DM++ produces a likelihood map of the processes in 512×512 pixel tiles. The outputs are in the range $[0, 255]$, with a higher value indicating greater confidence that a pixel belongs to a process present in the data. These likelihood maps were thresholded at a predefined threshold (STP: 100, MBA: 75 and BFI: 120) to generate a mask for the process in the section.

Pipeline for Process Detection:

The overall flow of the brain image data from the input to output is denoted as a pipeline for neuronal process detection. The pipeline for Process Detection is illustrated in figure 2(B), which was designed for two-dimensional process detection for a tracer injection in the brain, run section-wise for each brain. Each of these sections are first broken down into tiles of 512×512 pixels. These tiles are then individually passed to the discrete Morse algorithm and the trained ALBU models. The gray-scale maps generated out of the discrete Morse algorithm are fed into the topological pathway of the DM++ and the segmentation likelihood predicted by the ALBU models is fed to the neuronal segmentation pathway. These priors pass through these pathways and then merge into a dual stack (can be thought of as a two-channel image). This stack is further propagated through the common pathway deeper into the network to capture the high-level features of the neuronal process based on both the

priors. The trained DM++ network outputs a likelihood map for the process which is further thresholded to a binary mask for identifying the processes. This mask is also used to localize the detection of boutons, and the sub-categories of the neuronal processes discussed later in the semantic segmentation pipeline. These output tiles are eventually stitched together to form a cumulative mask for the whole section. The discrete Morse algorithm also outputs the peaks in the images as the vertices in the graph, which are seen in along the processes as the high-intensity peaks, identified as boutons or synaptic swellings on the neurons.

Variations in the pipeline for Process Detection: The pipeline has some variations across the different imaging modalities. The input to the pipeline is the raw image of a section from a brain for each data modality. The input is not the same across the different modalities, as seen in figure 1(A-C). The images from the Serial Two-Photon images are 16-bit gray scale images, while the images from the brains in Mouse Brain Architecture project are three-channel 12-bit images and the images from the brightfield WSI images were three-channel 8-bit images. The tracer injections in the STP and the Rabies data were of a single color but the MBA data had two injection colors.

The input to DM++ are single-color images. So, for the STP data the pipeline works on the raw data, while for the Brightfield images the 8-bit color data were converted into a single grayscale image before passing into the pipeline. The MBA data is treated channel-wise. Since the injections are green and red only, we assume that the green and red channels are separate likelihood images for the two tracers, and pass them individually through the pipeline, finally combining their outputs for the whole section. The BFI dataset is also treated as a likelihood image of intensities for the DM reconstruction but ALBU is performed on the three-channel color images.

Ground-truth Annotation: Manual annotation was used to generate training labels our supervised algorithms. The training set for the STP data consists of a coarse annotation of a whole brain (267 sections) and four section of high-quality precise annotations from another brain. The training data for the MBA images included four sections each from two brains with different labels marked as described in our class-hierarchy. The manually annotated BFI dataset consisted of 72 tiles of 1024×1024 pixels from a single section, which was divided for training and testing. We also aim to make the groundtruth annotations available to the research community upon acceptance, via brainarchitecture.org

Preparation of the Data for training and testing: The manually labelled sections were broken down into tiles of size (512×512) pixels. Since the processes are really sparse in the whole brain, the fully annotated STP was randomly tiled to incorporate tiles with varying densities of annotations ranging down to no annotations. Hence, we had a training set with 5749 tiles from the fully annotated STP brain and two high-quality annotated sections entirely tiled to obtain another 750 tiles. The testing for the STP data was done on another two sections fully tiled of the precise annotation data. The eight sections from the two brains in the MBA data were carefully annotated. We used 7 sections for training, resulting in ~ 2000 tiles and testing was done on a section with ~ 500 tiles. The BFI data had 180 tiles in the training set and 108 tiles in the testing set. The total number of tiles for

training and testing has been tabulated in table Supplementary 1, which gives an overview of the limited data available for training and testing our supervised networks.

Evaluation Metrics: The metrics used throughout the manuscript are Precision (P), Recall (R), and the F1-measure. For the evaluation of these metrics we determine the True Positives (TP), False Positives (FP), True Negatives (TN) and False Negatives (FN) (refer figure 4, bottom row leftmost tile for each modality). Every detected pixel was evaluated with a disk of radius of 5 pixels. if an annotated groundtruth existed within this disk, the pixel was considered as TP, otherwise it is considered as FP. If an annotation did not have detected pixel within a 5-pixel radius, then we considered that as TN. Utilizing all these values, Precision and recall were computed as $P = \frac{TP}{TP + FP}$ and $R = \frac{TP}{TP + FN}$, while the F1-score was calculated as harmonic mean of precision and recall, i.e., $F_1 = \frac{2 \cdot P \cdot R}{P + R}$. Also, we did an iterative refinement of annotations, as shown in the supplementary figure Supplementary 2, based on the FP detected by the algorithms on a first pass to double check for the human errors in the annotation. Each of our automatic detections were proof-read to ensure the authenticity of the automatic detects as shown in supplementary figure Supplementary 5.

Connectivity based Performance Metric: We define a connectivity score as follows: Given the ground-truth and the corresponding hard-thresholded prediction (both are 2D binary images), we first sample N random pairs of signals (pixels), each of which is path-connected in the ground-truth. Now, for each one of those pairs, we check the connectivity of their nearest neighbors in the predicted segmentation. Our *connectivity-score* is defined as the percentage of those sampled pairs whose nearest neighbors in the prediction are still path-connected. Obviously, a low connectivity-score means that the prediction is more fragmented than the ground-truth. In our experiment, we set $N = 1000$, and for each pixel we considered its 8 neighbors when deciding connectivity. The final connectivity-score was averaged over 704 STP tiles (images) each of size 512×512 . The connectivity scores for DM++, ALBU, and UNET are 0.9024, 0.6041, and 0.6169 respectively. Our DM++ framework achieves a ~30% advantage over both UNET and ALBU.

Bouton Detection:

The Discrete Morse Graph Reconstruction algorithm was combined with the finalized process detection outputs to detect putative synaptic boutons. A persistent homology⁴² computation was performed within the Discrete Morse Graph Reconstruction. Boutons appear in the image data as small balls with relatively high intensity along neuronal structures. This means that they will be vertices along the Morse skeleton graph with relatively high persistence. Thus to perform bouton detection, we took the vertices of the Morse skeleton output with persistence greater than a user provided threshold. Then from these points, we only include those that fell within the final process detection mask. Portions of the Morse skeleton graph lie outside of true process signal, so such a filter is needed to avoid labeling boutons outside of true process signals.

Pipeline for Semantic Segmentation:

The semantic segmentation architecture included a MaskRCNN network (see figure 3(A)) which performed two segmentation tasks: (i) cell detection, to identify the somata in the injection region as well as in other parts of the brain; and (ii) classification of the different process classes (axon arbors, fasciculated axons, and passing axons). For the semantic segmentation of the before-mentioned classes, the Process Detection is taken as a pre-processing step to mask out the background. MaskRCNN was employed on this pre-processed image to determine all the classes denoted in figure 1(D).

Architecture of MaskRCNN: MaskRCNN architecture (figure 3(A)) included a FasterRCNN network, which has two outputs for each candidate object, a class label, and a bounding-box offset. To this network, we add a third branch that outputs the object mask — which is a binary mask that indicates the pixels where the object is in the bounding box. The additional mask output is distinct from the class and box outputs, requiring extraction of much finer spatial layout of an object. To do this MaskRCNN uses the Fully Convolution Network (FCN). FasterRCNN is a good algorithm that is used for object detection. FasterRCNN consists of two stages. The first stage called a Region Proposal Network (RPN), which proposes candidate object bounding boxes. The second stage extracts features using RoIPool from each candidate box and performs classification and bounding-box regression. The features used by both stages can be shared for inference. So, in short, we can say that MaskRCNN combines the two networks — FasterRCNN and FCN - into one joint architecture. The loss function for the model is the combined loss in doing classification, generating a bounding box and generating the mask.

Soma Detection: The soma detection strategy is two-fold: (i) soma detection within the injection region, termed as "Injection soma" and (ii) soma detection outside the injection region. The injection region is either determined previously by their higher intensities for the different colors in the tracers or they are manually demarcated by an expert. At the injection region, we perform a contrast enhancement and intensity normalization using histogram equalization. All of the soma detected within this region are demarcated as the injection soma, while the others are referred as projection-region soma. Both of these are detected as masks by the MaskRCNN architecture. These masks are determined as individual connected components in the brain slices and the centroids of these connected components are treated as cell-centers. The soma detection algorithm was also applied to all the sections for the MBA (please refer to table Supplementary 1 for the size of the training and testing set); both injection and non-injection sections. For the BFI dataset, limited data availability (please refer to table Supplementary 1 for the size of the training and testing set) allowed us only to categorize the data as soma (where we assume to be outside any injection-like region). The evaluation of these cell centers are also based on the same evaluation technique described earlier.

Semantic Detection: The semantic detection deals with the second layer of the hierarchy after Process Detection. The processes consist mostly of axons and dendrites, but for the semantic detections we are targeting the different intrinsic properties of these processes - viz. the axon arbors and the passing axons. Similarly we sub-categorize dendritic processes.

On a training set of four manually annotated sections, we trained the MaskRCNN network individually on each of these classes using one-vs-other classification strategy, with data augmentation. The results for the sub-categories from the axons and dendrites are masked using the process detection mask generated earlier. This semantic segmentation process was carried on the MBA data. The overall pipeline of the semantic segmentation is described in figure 3(B) where the process detection acts as one of the important modules within the semantic segmentation pipeline. The evaluations for the semantic detections also follow the same procedure as discussed earlier for the process detection.

First, we detect the dendrites. The training is done on four sections. The rest is classified as axons. The manual annotation had 31 tiles of size 512×512 pixels to contain data. Out of which 8 tiles have been separated out for the testing and rest for training. Secondly, on the detected axons, we train our MaskRCNN on the passing axons. The rest are classified as axon arbors. In total, from the 4 sections we had 28 tiles with axonal data. The model trained on dendrites throws out the dendrites from the process detection masks. Six of these axonal tiles have been used for testing while the rest are used for training the MaskRCNN (please refer to table Supplementary 1 for the size of the training and testing set). The quantitative values are reported in table 1.

Extended Data




Figure Extended Data 1:

The Discrete Morse algorithm is given an input image (A). A Gaussian filter is applied to the image (B) - a density function is defined at the pixels. Then the algorithm extracts the ridges of the function across the domain (C) - these ridges form the 1-stable manifold. Finally, each path in the 1-stable manifold is assigned an grayscale value based on intensity along the path, and a grayscale mask is outputted in (D).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The authors gratefully acknowledge IHC Brightfield Imaged data from Dr Peter Strick at U Pitt, and wish to thank Jaimi Nagashima and Mitsutoshi Hanada also at U Pitt for annotating these images. We acknowledge the effort from annotators at the Center for Computational Brain Research at IIT Madras for the bulk of the data annotation and proofreading for this project. This work was supported by the NIH (EB022899, MH114824, MH114821, NS107466, AT010414), the Crick-Clay Professorship (Cold Spring Harbor Laboratory), the Mathers

Charitable Foundation, and H. N. Mahabala Chair (IIT Madras). Work at Ohio State University was in addition partly supported by the NSF under grants CCF-1740761, RI-1815697, and DMS-1547357.

Data Availability

The STP data was collected as a part of Brain Initiative Cell Census Network and shared online. The raw images of the STP dataset are available from: ftp://download.brainimagelibrary.org:8811/biccn/huang/connectivity/anterograde/180830_JH_WG_Fezf2LSLflp_CFA_female_processed/. The sample MBA data can be viewed <http://brainarchitecture.org/viewer4/mouse/map/3611>. The trained models and the annotated data are available on the <http://brainarchitecture.org/semantic-segmentation-in-brain-images>.

References

1. Bohland JW et al. A proposal for a coordinated effort for the determination of brainwide neuroanatomical connectivity in model organisms at a mesoscopic scale. *PLoS computational biology* 5 (2009).
2. Dodt H-U et al. Ultramicroscopy: three-dimensional visualization of neuronal networks in the whole mouse brain. *Nature methods* 4, 331–336 (2007). [PubMed: 17384643]
3. Ragan Tet et al. Serial two-photon tomography for automated ex vivo mouse brain imaging. *Nature methods* 9, 255 (2012). [PubMed: 22245809]
4. Oh SW et al. A mesoscale connectome of the mouse brain. *Nature* 508, 207–214 (2014). [PubMed: 24695228]
5. Pinskiy Vet et al. High-throughput method of whole-brain sectioning, using the tape-transfer technique. *PloS one* 10 (2015).
6. Lin MK et al. A high-throughput neurohistological pipeline for brain-wide mesoscale connectivity mapping of the common marmoset. *Elife* 8, e40042 (2019). [PubMed: 30720427]
7. Halavi M, Hamilton KA, Parekh R & Ascoli G Digital reconstructions of neuronal morphology: three decades of research trends. *Frontiers in neuroscience* 6, 49 (2012). [PubMed: 22536169]
8. Helmstaedter M & Mitra PP Computational methods and challenges for large-scale circuit mapping. *Current opinion in neurobiology* 22, 162–169 (2012). [PubMed: 22221862]
9. Peng H, Meijering E & Ascoli GA From diadem to big neuron (2015).
10. Rey-Villamizar Net et al. Large-scale automated image analysis for computational profiling of brain tissue surrounding implanted neuroprosthetic devices using python. *Frontiers in neuroinformatics* 8, 39 (2014). [PubMed: 24808857]
11. Peng Het et al. Bigneuron: large-scale 3d neuron reconstruction from optical microscopy images. *Neuron* 87, 252–256 (2015). [PubMed: 26182412]
12. Lawrie SM & Abukmeil SS Brain abnormality in schizophrenia: a systematic and quantitative review of volumetric magnetic resonance imaging studies. *The British Journal of Psychiatry* 172, 110–120 (1998). [PubMed: 9519062]
13. Taylor RH, Lavealle S, Burdea GC & Mosges R Computer-integrated surgery: technology and clinical applications (Mit Press, 1995).
14. Zijdenbos AP & Dawant BM Brain segmentation and white matter lesion detection in mr images. *Critical reviews in biomedical engineering* 22, 401–465 (1994). [PubMed: 8631195]
15. Worth AJ, Makris N, Caviness VS Jr & Kennedy DN Neuroanatomical segmentation in mri: technological objectives. *International Journal of Pattern Recognition and Artificial Intelligence* 11, 1161–1187 (1997).
16. Khoo VSet et al. Magnetic resonance imaging (mri): considerations and applications in radiotherapy treatment planning. *Radiotherapy and Oncology* 42, 1–15 (1997). [PubMed: 9132820]
17. Grimson WEL et al. Utilizing segmented mri data in image-guided surgery. *International Journal of Pattern Recognition and Artificial Intelligence* 11, 1367–1397 (1997).

18. LeCun Y, Yoshua bengio, and geoffrey hinton. Deep learning. *nature* 521, 436–444 (2015).
19. LeCun Y, Bottou L, Bengio Y & Haffner P Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86, 2278–2324 (1998).
20. Fukushima K Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics* 36, 193–202 (1980). [PubMed: 7370364]
21. Das S et al. High precision automated detection of labeled nuclei in gigapixel resolution image data of mouse brain. *BioRxiv* 252247 (2019).
22. Ramesh N, Yoo J-H & Sethi I Thresholding based on histogram approximation. *IEE Proceedings-Vision, Image and Signal Processing* 142, 271–279 (1995).
23. Sharma N & Ray AK Computer aided segmentation of medical images based on hybridized approach of edge and region based techniques. In *Proc. Int. Conf. Math. Biol* (2006).
24. Boykov YY & Jolly M-P Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In *Proceedings eighth IEEE international conference on computer vision. ICCV 2001*, vol. 1, 105–112 (IEEE, 2001).
25. Litjens G et al. A survey on deep learning in medical image analysis. *Medical image analysis* 42, 60–88 (2017). [PubMed: 28778026]
26. Krizhevsky A, Sutskever I & Hinton GE Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105 (2012).
27. He K, Gkioxari G, Dollár P & Girshick R Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969 (2017).
28. Redmon J, Divvala S, Girshick R & Farhadi A You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788 (2016).
29. Vinyals O, Toshev A, Bengio S & Erhan D Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence* 39, 652–663 (2016). [PubMed: 28055847]
30. Sabour S, Frosst N & Hinton GE Dynamic routing between capsules. In *Advances in neural information processing systems*, 3856–3866 (2017).
31. Badrinarayanan V, Kendall A & Cipolla R Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 39, 2481–2495 (2017). [PubMed: 28060704]
32. Ronneberger O, Fischer P & Brox T U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241 (Springer, 2015).
33. Buslaev A, Seferbekov SS, Iglovikov V & Shvets A Fully convolutional network for automatic road extraction from satellite imagery. In *CVPR Workshops*, 207–210 (2018).
34. Belkin M, Hsu DJ & Mitra P Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in neural information processing systems*, 2300–2311 (2018).
35. Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T & Ronneberger O 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, 424–432 (Springer, 2016).
36. Milletari F, Navab N & Ahmadi S-A V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, 565–571 (IEEE, 2016).
37. Johnson JM & Khoshgoftaar TM Survey on deep learning with class imbalance. *Journal of Big Data* 6, 27 (2019).
38. Delgado-Friedrichs O, Robins V & Sheppard A Skeletonization and partitioning of digital images using discrete morse theory. *IEEE transactions on pattern analysis and machine intelligence* 37, 654–666 (2014).
39. Gyulassy A, Bremer P-T, Hamann B & Pascucci V A practical approach to morse-smale complex computation: Scalability and generality. *IEEE Transactions on Visualization and Computer Graphics* 14, 1619–1626 (2008). [PubMed: 18989018]

40. Robins V, Wood PJ & Sheppard AP Theory and algorithms for constructing discrete morse complexes from grayscale digital images. *IEEE Transactions on pattern analysis and machine intelligence* 33, 1646–1658 (2011). [PubMed: 21576736]
41. Dey TK, Wang J & Wang Y Road network reconstruction from satellite images with machine learning supported by topological methods. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 520–523 (2019).
42. Edelsbrunner H & Harer J Persistent homology-a survey. *Contemporary mathematics* 453, 257–282 (2008).
43. Forman RA user's guide to discrete morse theory. *Sém. Lothar. Combin*48, 35pp (2002).
44. Sousbie T, Pichon C, Colombi S, Novikov D & Pogosyan D The 3d skeleton: tracing the filamentary structure of the universe. *Monthly Notices of the Royal Astronomical Society* 383, 1655–1670 (2008).
45. Ding L & Goshtasby A On the canny edge detector. *Pattern Recognition* 34, 721–725 (2001).
46. Menze BH et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging* 34, 1993–2024 (2014). [PubMed: 25494501]
47. Kainmüller D, Lange T & Lamecker H Shape constrained automatic segmentation of the liver based on a heuristic intensity model. In *Proc. MICCAI Workshop 3D Segmentation in the Clinic: A Grand Challenge*, 109–116 (2007).
48. He Jet al. Deep reinforcement learning with a natural language action space. *arXiv preprint arXiv:1511.04636* (2015).
49. Wang Det al. Detection and skeletonization of single neurons and tracer injections using topological methods. *arXiv preprint arXiv:2004.02755* (2020).
50. Özgün Çiçek, Abdulkadir A, Lienkamp SS, Brox T & Ronneberger O 3d u-net: Learning dense volumetric segmentation from sparse annotation (2016). 1606.06650.






Figure 1:

Three prevalent modalities of light-microscopic data acquisition where we have applied our algorithm: **(A)** Colored fluorescent Whole Slide Images (WSI) of brain sections (mouse) acquired using digital slide scanners, **(B)** Bright Field images of immunohistochemically brain sections (marmoset) and **(C)** monochrome Serial Two Photon Tomographic images (STPT) (mouse). The top panels show the whole section images and below are the neuronal soma, neuronal processes and putative terminating axon arbors with putative synaptic boutons. **(D)** shows the class hierarchy that we devised to categorize different image components (objects of interests) that we targeted for semantic segmentation.



A. Architecture for Mask Generation & Cell Detection



B. Workflow for Semantic Segmentation





Figure 3:

(A) The figure shows the architecture of the Mask-RCNN network, which we use for Cell Detection as well as for mask generation for different process sub-classes (*e.g.*, passing vs arborized axons). (B) Semantic segmentation for tracer-injected brains is preceded by an Injection Detection step. This is needed since the injection region has a different signal dynamic range than the projection regions. The injection region is suppressed for the Process Detection task, and is also used to separate “injection region somata” from distal somata. The Process Detection output (manually proof-read) acts as a mask for discarding the other background pixels while generating masks for process sub-classes. We are utilizing a high resolution image viewer available in the web portal (brainarchitecture.org) to visualize the output of the semantic segmentation pipeline.

**Figure 4:**

Illustrations of results from the process detection pipeline are shown here for single-color STP images (**3 top left columns**), 3-color fluorescent WSI images from the Mouse Brain Architecture project data set (MBA) (**3 top right columns**), and the Brightfield WSI images (**3 bottom columns**). The top left row shows the original image, human annotation of the processes, and the discrete Morse output. The bottom left shows evaluation of DM++ output compared with the human annotation, with true positives (TP) in yellow and "false positives" (FP) in red. The FP connect the human-annotated neurite fragments, and examination of the original image data shows low intensity signal connecting the annotated fragments. The final panel in the bottom left row is the ALBU output. In each of the three modalities, the middle row depicts a small patch from the image, where the intensity captured by the digital scanners used for imaging these sections is significantly low, but the connectivity can be observed in the high resolution image by naked human eye, when inspected meticulously (leftmost patch in the middle row). As evident in the detections by ALBU, a state-of-the-art CNN technique, they are generally missed while we generate the mask. These faint connections are well captured by Discrete Morse as evident in the patch from the same location in the output of the Discrete Morse reconstruction algorithm (rightmost patch in the middle row). With the introduction of the topological prior the patch from the same location of our output preserves the connectivity (central patch in the middle row). These patches demonstrate that the ALBU output misses faint connections that are picked up by the DM++ algorithm. The bottom right row shows the evaluation of DM++

output compared to manual annotations with TP (yellow), FP (magenta) and FN (cyan). In the MBA modality, the top right row shows red and green process fragments corresponding to neurites labelled with red and green fluorescent protein corresponding to anterograde AAV tracer injections, human annotations (double labelled process fragments are labelled green), and discrete morse output. Evaluations, DM++ and ALBU outputs for the colored MBA image data are shown next, followed by the similar structure in the BFI dataset. The rectangular box indicates the location for each image from where the exemplar patch is extracted and the dotted lines establishes the correspondences of each path to the middle row of results for better visibility.




Figure 5: Semantic Segmentation results:

(A) the results for cell detection for the two modalities along with the injection soma detection for the MBA dataset. Soma segmentations are shown to the immediate right of the original image shown in the left column. The right column shows the detected cells in a different color, for the corresponding left patch. The top row shows a region where the tracer is injected ("Injection Region"). The detected cells in this region are labelled "Injection Soma". The bottom two rows show detections outside this injection region for the two modalities, Fluorescent WSI images (MBA data) and Brightfield WSI images. **(B)** The results for the semantic category ("dendrites") for MBA data, a sub-category of processes. **(C)** The results for the semantic category ("passing axons") for MBA data, a sub-category of processes. **(D)** The results for the semantic category ("axon arbors") of the MBA data. The top tile in each of these semantic categories shows the detected processes from the Process Detection Pipeline for the green/red tracer. The middle tile illustrates the resulting semantic segmentation mask, while the bottom tile shows the comparison of our result with the Groundtruth (GT) annotation, depicting the TP (yellow), FP (magenta) and FN (cyan) for the input tile. **(E)** The results for bouton detection on the MBA dataset.

Table 1:

shows the Precision, Recall, F1 values for the different tasks. The unsupervised technique for Process Detection involves intensity-based thresholding with hard-coded parameters for threshold and morphological operations. Cell Detection Results are shown for the WSI fluorescent images (MBA) and WSI Bright Field images (BFI) compared to a baseline technique (Seg-NET). The (I) on the side of the MaskRCNN indicates that we omit Injection Region while reporting our metrics. In computing these metrics we utilize a 5-pixel neighborhood of a detected point to judge detects. The values in bold indicate the best performance.

Model	Precision	Recall	F1
<i>Process Detection - STP</i>			
Unsupervised	0.61	0.63	0.62
UNET	0.85	0.88	0.86
ALBU	0.90	0.88	0.89
DM++	0.92	0.93	0.92
<i>Process Detection - MBA</i>			
Unsupervised	0.41	0.51	0.46
UNET	0.67	0.73	0.70
ALBU	0.79	0.82	0.81
DM++	0.83	0.84	0.84
<i>Process Detection - BFI</i>			
Unsupervised	0.59	0.61	0.60
UNET	0.73	0.75	0.74
ALBU	0.79	0.80	0.80
DM++	0.81	0.83	0.82
<i>Cell Detection - BFI</i>			
SegNET	0.79	0.74	0.76
MaskRCNN(I)	0.85	0.89	0.87
<i>Cell Detection - MBA</i>			
SegNET	0.72	0.75	0.73
MaskRCNN	0.81	0.82	0.82
MaskRCNN(I)	0.85	0.90	0.87
<i>Dendrite Detection - MBA</i>			
SegNET	0.67	0.71	0.69
MaskRCNN	0.79	0.86	0.83
<i>Passing Axon Detection - MBA</i>			
SegNET	0.72	0.76	0.74
MaskRCNN	0.89	0.93	0.91
<i>Axon Arbor Detection - MBA</i>			
SegNET	0.72	0.75	0.74
MaskRCNN	0.95	0.97	0.96
<i>Bouton Detection - MBA</i>			

Model	Precision	Recall	F1
Discrete Morse	0.76	0.31	0.44

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table Methods 1:

The number of manually annotated tiles used for training and testing the different semantic categories in the different modalities of the brain. These manually annotated tiles are of size 512×512 pixels. They are used to train the different pipelines in the brain, with data augmentation, to cope with the limited availability of the training data. These tiles account for the tiles where the annotators have at least annotated any of the neuronal structures.

Image modality	Neuronal classes	Tiles for training	Tiles for testing
MBA	Processes	2000	494
STPT	Processes	6449	704
BFI	Processes	180	108
MBA	Soma Detection	26	9
BFI	Soma Detection	48	24
MBA	Dendrites	23	8
MBA	Passing Axons	22	6
MBA	Axon Arbors	22	6