Pacific Northwest
NATIONAL LABORATORY

PNNL-33353

Data-driven Mapping of the Mouse Connectome

The utility of transfer learning to improve the performance of deep learning models performing axon segmentation on light-sheet microscopy images

September 2022

Marjolein Oostrom, Rogene Eichler West, Moses Obriri, Michael Muniak, Paritosh Pande, Sarah Akers, Tianyi Mao, Bobbie-Jo Webb-Robertson

U.S. DEPARTMENT OF ENERGY

Prepared for the U.S. Department of Energy under Contract DE-AC05-76RL01830

DISCLAIMER

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <https://www.ntis.gov/about>
Online ordering: http://www.ntis.gov

PNNL-33353

Data-driven Mapping of the Mouse Connectome

September 2022

Marjolein Oostrom, Rogene Eichler West, Moses Obriri, Michael Muniak, Paritosh Pande, Sarah Akers, Tianyi Mao, Bobbie-Jo Webb-Robertson

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Abstract

Light sheet microscopy has made possible the high temporal and spatial 3D imaging of both fixed and live biological tissue, with samples as large as the entire mouse brain. However, segmentation and quantification of that data remains a time-consuming manual process. Machine learning methods promise the possibility of automating this process. This study seeks to advance the performance of prior models through the application of refinements such as transfer learning.

Summary

Light-sheet fluorescence microscopy is an important tool in understanding how the functional brain states are achieved and how they are modulated by experience, genetics, and pharmacological agents require a reference map of the structural connections between neurons. Technical challenges remain to make this high resolution data meaningful for connectome research, specifically: 1) to stitch together aligned serial sections to construct a 3D image volume; 2) to efficiently annotate voxels of interest (for example, labeling the content as belonging to the axon of a particular neuron); and, 3) registering landmarks in an image volume with a reference atlas so that data from multiple studies might be compared in the same coordinate space (Tyson and Margrie, 2021).

TrailMap (Tissue Registration and Automated Identification of Light-sheet Microscope Acquired Projectomes), a 3D U-net model developed for axon segmentation, provides an opportunity to exploit the benefits of transfer learning specifically from other images in this domain (Friedmann et al., 2020). This study endeavored to combine the model optimization benefits of the nnU-Net framework with the domain-specific transfer learning capabilities of TrailMap by comparing the performance in identifying axons in mouse brain slices using the standard settings from the TrailMap package (Friedmann et al., 2020) against the pre-processing and network adjustment suggested by nnU-Net (Isensee et al., 2019). The model will be fine-tuned with our own labeled images, once they are entirely available.

Fine-tuning did not improve the model's accuracy. However, we only have one cube with 160 slices, which includes 5 labeled slices. When it's divided into training and testing data, each dataset only has 2-3 labeled slices. Hopefully when more labeled data become available, the models will improve with fully fine-tuning on the entire dataset. Technologies such as TrailMap contribute to the growing effort to develop templates to integrate data across multiple non-disruptive 3-dimensional image modalities, combining insights from light-sheet microscopy and MRI into a common coordinate framework (Perens et al., 2022). Such co-registration is essential for studying disease progression, developing effective therapies, translating findings into clinical procedures, for example, delivering brain site-specific drugs and implanting electrodes precisely for recording or stimulation.

While the work in this study demonstrates improvements in the ability to automate axon segmentation, the problem of complete mapping of neuronal structures into a co-registered system will likely require many structure-specific solutions in the short-term. The segmentation of tightly packed neurites will not benefit from an adaptive enhancement of the signal-to-noise ratios and may require other kinds of algorithmic slights of hand to feign higher resolution (Zhou et al., 2022).

Acknowledgments

This work is supported by the PNNL Laboratory Directed Research and Development (LDRD) program. It was performed at the Pacific Northwest National Laboratory (PNNL), a facility operated by Battelle for the U.S. Department of Energy. Battelle operates PNNL for the U.S. Department of Energy under contract DE-AC05-76RL01830.

Acronyms and Abbreviations

3D: 3-Dimensional

CNN: Convolutional Neural Network

GPU: Graphical Processing Unit

LSFM: Light Sheet Fluorescence Microscopy

nnU-Net: no-new-U-Net

OHSU: Oregon Health & Science University

PNNL: Pacific Northwest National Laboratory

TrailMap: Tissue Registration and Automated Identification of Light-sheet Microscope Acquired Projectomes

Figures

Tables

1.0 Introduction

Understanding how functional brain states are achieved and how they are modulated by experience, genetics, and pharmacological agents require a reference map of the structural connections between neurons. Such a map will facilitate research into brain disorders but will also generate insights into the information processing mechanisms of healthy brains. Such is the goal of the Mouse and Human Connectome Projects (Bargmann and Marder, 2013; Sporns, 2011; Zingg et al., 2014).

Light-sheet fluorescence microscopy (LSFM) is an important tool in this discovery process. Based on wide-field fluorescence technologies, LSFM enables direct visualization of cellular structure on a scale of a few micrometers. It introduced the use of intrinsic optical sectioning capabilities by automating the movement of a sample through a single plane sheet of light. The resulting fluorescence is detected by an orthogonal-placed camera. By illuminating only a thin section of tissue with each pass, photobleaching of the tissue is minimized, yielding a more intense and complete fluorescence signal. Another important advancement associated with LSFM is optical tissue clearing. Through the use of organic solvents, lipid removal, or immersion in refractive index matching solutions, brain tissue is rendered effectively transparent, allowing greater contrast of the labeled fluorescent structures, including in a living brain (Girkin and Carvalho, 2018; Olarte et al., 2018; Santi, 2011)

Technical challenges remain to make this high resolution data meaningful for connectome research, specifically: 1) to stitch together aligned serial sections to construct a 3D image volume; 2) to efficiently annotate voxels of interest (for example, labeling the content as belonging to the axon of a particular neuron); and, 3) registering landmarks in an image volume with a reference atlas so that data from multiple studies might be compared in the same coordinate space (Tyson and Margrie, 2021). Annotating a volume of images is a labor and time-intensive process that can be subjective due to fatigue-related errors and inhomogeneities in the intensity of stained structures resulting in significant variability between experimenters. The number of images slices that yield a complete volume range from several dozen to well over ten thousand, depending on the microscopy system, species, and orientation of the sectioning method (Agarwal et al., 2018; Kasthuri et al., 2015). A complete dataset of high-resolution data for a whole brain can be on the order of a terabyte, limiting analysis algorithms to those that do not require the full volume of data to be loaded into runtime memory (Seiriki et al., 2017).

Computational methods to automatically stitch, segment, annotate, and register meaningful information in images have made enormous strides in the past decade. Early methods tended to rely on thresholding, clustering, edge detection, region growing, or curve propagation (Kaur and Kaur, 2014; Zhang, 2006). A significant break-through came in 2015 with the development of a class of convolutional networks (CNNs) that use symmetric contractive and expansive passes over augmented data samples. Such an approach allowed for a balance between localization and context, as well as a robustness to deformations and variance within the examples from which the network learned. The first such class of these networks, U-Net, won the ISBI (IEEE International Symposium on Biomedical Imaging) cell tracking challenge that year by a wide margin Ronneberger et al. (2015). This success yielded a plethora of new U-Net architectures and approaches intended to refine and build on this advancement, but also to address some of the limitations. In particular, CNNs require that their hyperparameters be tuned to obtain optimal performance, which can be considerably time consuming. And secondly, the problem of a sufficient number of hand-annotated datasets with which to train the CNN remains.

Network models require optimization because experimental data significantly varies with respect to class ratios, image size, image generation, image pre- and post-processing methods, network architecture, data augmentation, learning rate, and loss functions. nnU-Net (no-new-U-Net) sought to address this issue by establishing a framework with a baseline against which optimizations for model performance might be compared and a methodology for the systematic tuning of parameters against that baseline (Isensee et al., 2019). First fixed parameters values, which will stay constant regardless of the dataset, are set by evaluating model performance while systemically modifying the parameter values for 10 medical segmentation datasets. These parameters include the learning rate and loss function. Secondly, some parameters are discovered experimentally with every run, and these are set by evaluation of cross-validation performance. These experimentally found parameters include the post-processing configuration. Finally, parameters whose optimal value will depend on the dataset, are found with heuristic rules depending on the “fingerprint” of the dataset, which includes details such as image size, voxel spacing, and class balance.

Transfer learning can be a useful method in which the features learned from another domain with a richly annotated dataset are used to segment a new smaller dataset outside of the original domain (Caruana, 1994). Perhaps the best known pre-trained object detection model, AlexNet, was initially trained on over 1 million training images from an extensively studied dataset, ImageNet. Transfer learning leverages all of the features learned from those millions of images as a starting place for reasoning in a new setting with fewer labeled images (Krizhevsky et al., 2012).

TrailMap (Tissue Registration and Automated Identification of Light-sheet Microscope Acquired Projectomes), a 3D U-net model developed for axon segmentation, provides an opportunity to exploit the benefits of transfer learning specifically from other images in this domain (Friedmann et al., 2020). The Trailmap model was trained with 36 patches of a labeled brain volume containing axons of 100 - 300 voxels per dimensions, and 40 patches containing imaging artifacts and backgound, all with 3 - 10 labeled xy planes per patch. Model evaluation accounts for sparsely labeled images (such that only every 80 to 180 μm slice are being labeled) by masking unlabeled slices. Edges of axons are generated within the code; the misclassification of edges is devalued in calculating loss compared to the misclassification of axons, ensuring that a false positive within one pixel of the axon is not strongly penalized, thereby minimizing the effect of pixel-size differences in expert labeling.

This study endeavored to combine the model optimization benefits of the nnU-Net framework with the domain-specific transfer learning capabilities of TrailMap by comparing the performance in identifying axons in mouse brain slices using the standard settings from the TrailMap package (Friedmann et al., 2020) against the pre-processing and network adjustment suggested by nnU-Net (Isensee et al., 2019). As we wish to keep the innovative features of Trailmap (e.g., sparse labeling, edges generation, network weights from pretraining), we will adjust only the parameters that can be changed without changing the network architecture (ex. the number of down-sampling operations). Our modified version of Trailmap will adopt nnU-Net’s approach to data augmentation, patch size selection, data sampling, learning rate, and inference method. The model will be fine-tuned with our own labeled images.

2.0 Materials and Methods

Labeled Image Generation

Labeled images for fine tuning the model were obtained according to the following procedure.

[* The following is template text; PNNL is waiting for a response from OHSU collaborators to correct the steps with their specific procedures.]

"*All animal procedures followed animal care guidelines approved by the Oregon Health Science University's Animal Research Committee. Mice were transferred to a stereotaxic apparatus where they were maintained under 2% isoflurane in oxygen. Following a small skin incision, a dental drill was used to drill through the skulls above the injection targets. For collateralization mapping, 300uL of Addgene was injected unilaterally into either left NAc (AP: 1.3, ML: -1.0, DV: 4.7), left VTA (AP: -3.3, ML: 0.4, DV: -445 4.5), or right PL (AP: 1.8, ML: -0.5, DV: -2.3). For topographical mapping of the origin of each projection type, mice were injected with 300uL of cholera toxin subunit B of Fluorogold at the same coordinates.*

Mice were transcardially perfused with phosphate-buffered saline (PBS) followed by 4% paraformaldehyde (PFA) in PBS. Brains were dissected, post-fixed in 4% PFA for 12–24h and placed in 30% sucrose for 24–48 hours. Sections were then stained with chicken anti-GFP in 0.3% PBST containing 3% donkey serum overnight at 4°C. The following day, sections were washed 3x5min in PBS and then stained with secondary antibody in 0.3% PBST containing 5% donkey serum for 2 hours at room temperature.

Mouse brains were collected and processed based on the published Adipo-Clear protocol with slight modifications. Brains were hemisected approximately 1mm past midline and postfixed overnight in 4% PFA at 4°C. To dehydrate samples, a gradient of washes in MeOH:H2O (20%, 40%, 60% and 80%) were conducted for 30min each, followed by 3x100% MeOH for 30min, 1h, then 1.5h. The next day, samples were washed in 100% DCM 2x1h each. Samples were then cleared in 100% DBE (dibenzyl ether). DBE was changed after 4h. Samples were stored in DBE in a dark place at room temperature. Imaging took place at least 24h after clearing.

Brain samples were imaged on a light-sheet microscope (Ultramicroscope II, LaVision Biotec) equipped with a sCMOS camera (Andor Neo) and a 2x/0.5 NA objective lens (MVPLAPO 2x) equipped with a 6 mm working distance dipping cap. Image stacks were acquired at 0.8x optical zoom using Inspector Microscope v285 controller software. We imaged using 488-nm (laser power 20%) and 640-nm (laser power 50%) lasers. The samples were scanned with a step-size of 3 µm using the continuous light-sheet scanning method with the included contrast adaptive algorithm for the 640-nm channel (20 acquisitions per plane), and without horizontal scanning for the 488-nm channel. The images were then annotated by hand. "

The images used in this study are 160 pixel in the x, y, and z direction. Every 20th section is annotated manually, starting at the 15th slice. After the expert annotation, the images were processed with Trailmap's axon edge annotator scirp to provide the edge labels. These images were subsequently used for training and testing of the modified TrailMap/nnU-Net model (Figure 1).

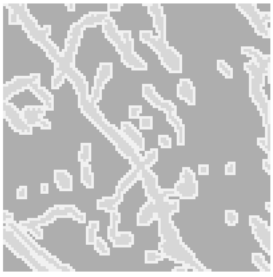

Figure 1. Example of a labeled slice of data

2.1 Hybrid Model Development

2.1.1 Data Augmentation

Data augmentation is the process of increasing the amount of training data by generating new instances from existing data, thereby reducing the tendency of a model to overfit. Common techniques include a combination of affine transformations such as scaling, rotation, reflection, and shearing.

TrailMap contains functions for generating augmented data using scaling and rotation. nnU-Net uses a Python package, Batchgenerators, for augmentation that additional included elastic and spatial transformations (Isensee et al., 2020). With a few modifications, these additional augmentation methods were included in our updated version of into the TrailMap code.

The script for augmentation allows the user to decide what percentage of images will have each type of transformation applied. In addition, the user can adjust the degree of transformation by setting the range the variables for each of the rotations, which are sampled randomly from a uniform distribution for each image. The images are generated by a Tensorflow generator during the training, which creates new images for each “step” (batch size) of the epoch. 720 training patches and 180 test patches were created.

An illustration of the spatial transformations on a input patch and labels can be seen in Figure 2. It includes an artificial line to illustrate that the patch input and label remain correctly mapped to each other after the transformation.

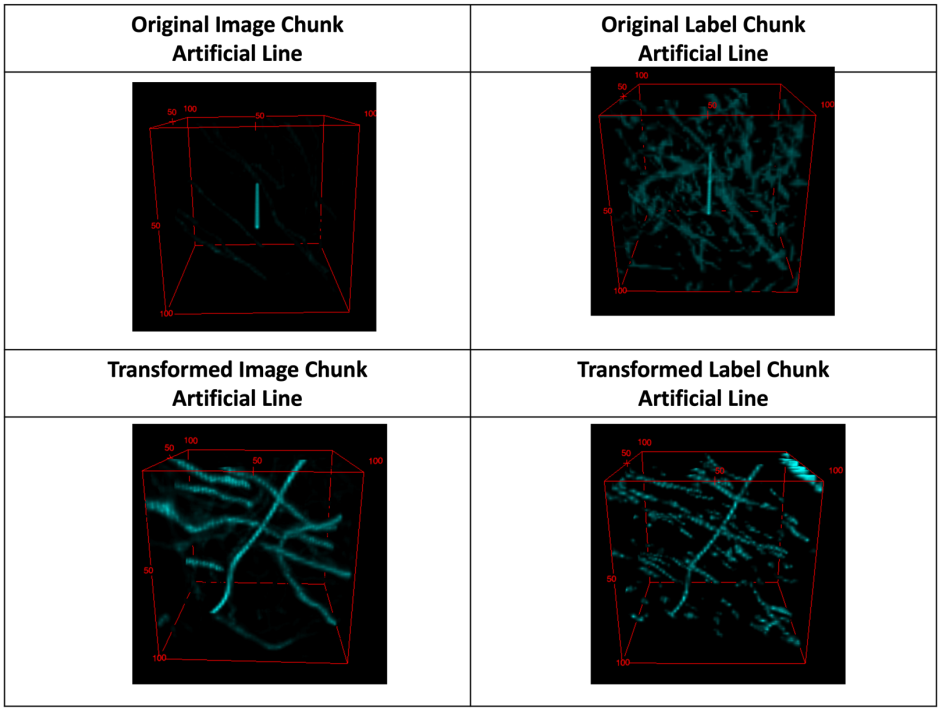

Figure 2. Data Augmentation: The results of spatial transformations (bottom rows) on the image input and output with a line added to illustrate the transformations, created with the TrailMap starting data

2.1.2 Architecture

The architecture of a model describes the number, size, and kinds of layers, as well as their connectivity patterns. Because we are fine-tuning a pre-trained TrailMap model, we are constrained to the architecture of that model.

2.1.3 Patch Size

Patch size in a convolutional neural network refers to a patch of an image retaining local context in an operation. This differs from a kernel or filter, which is used in a convolutional operation to obtain a weighted average of pixel values.

nnU-Net seeks to use the largest patch that will run within the computational environment. nnU-Net determines patch size by the maximum size that can be processed within a GPU constraint. The increasing patch size is accompanied with changes in the network architecture (Isensee et al., 2019). TrailMap uses a standard input of 64 pixels for the images and 36 pixels for the labels.

Trailmap is based on U-net, which allows variable input sizes. Therefore, without changing the model architecture, we can change the input size. In our incorporation of the nnU-net code into the Trailmap, we will use the largest patch can also be manually determined by incrementally increasing the patch size until a memory error is thrown. Our current sample of data is too small to dramatically increase the input size, but we plan to increase the input size when we receive more data.

2.1.4 Data Sampling

TrailMap uses 3D patches from the training data to create the training dataset and 3D patches from the validation data to create the validation dataset. These patches are randomly obtained from the larger 3D training tif file. In our case, the training tif file is 160 pixel for each axis, while the 3D patches used in the training dataset are each 64 pixels for each axis. nnU-Net oversamples the 3D patches from areas in the full training image with foreground features (Isensee et al., 2019). The foreground features for our study are the axon, which are indicated in the labeled data. The nnU-Net code for oversampling simply selects from a subset of data with the foreground image instead of all the data and is straightforward to include in the TrailMap code.

2.1.5 Overlapping Windows

Both nnU-net and Trailmap use sliding windows for inference. The difference between the two models lies in whether their sliding windows overlap: TrailMap uses non-overlapping sliding windows for inference, with the exception of the edge cases where overlapping tiles ensure the entire image is used even when the image dimensions are not evenly divisible by the patch dimensions. nnU-net uses overlapping sliding windows. The sliding windows the same size in both cases. We might except an increased performance with overlapping windows as the no voxel will be predicted using only a window where it is on the edge of the window.

Our modified TrailMap code uses the same logic and the same Gaussian matrix function as nnU-Net. Each overlapping patch is multiplied by a Gaussian matrix before being added to the predicted image. Each resulting predicted values is then divided by the sum of the Gaussian multipliers used at that location in the image. The value at each voxel at the end is a weighted sum, with the weights representing the centrality of the voxel of each patch. The logic for this method is that predictions at the edges of patches will be less accurate due to edges having less surrounding information and thus should be weighted less.

It's important to note that implementing the Gaussian function even without an overlap changes predictions on the edge cases. To compensate for this, a weighted sum is applied when multiple predictions occur for one voxel, regardless of the predicted value of the center tiles.

2.1.6 Learning rate

The learning rate in a neural network controls how quickly the model changes to incorporate the new patterns being presented. A rate that is too slow will progress very slowly, however; a rate that is too high can cause unstable behavior.

TrailMap uses a fixed learning rate, while nnU-Net uses an adaptive rate that decreases when there is a prolonged period of no improvement to the validation loss. It's straight-forward to reduce the learning rate after a metric has stopped improving with a built-in Pytorch learning-rate scheduler.

The TrailMap had a learning rate of .001. For nnU-net we used the ReduceLROnPlateau callback to reduce the learning rate once the model stops improving.

2.1.7 Fine-tuning the modified model

We used 100 epochs at batch-sizes of 6. At each epoch, 720 training patches and 180 test patches were run. The training and test data was created by splitting the labeled cube in half on the z-axis. We used Adam as the optimizer and used cross-entropy as the loss function.

3.0 Results

We predicted the neural segmentation for the test dataset using the original model, the fine-tuned model using the TrailMap code, and the fine-tuned model using the TrailMap code with the nnU-Net augmentation. We used adjusted accuracy, axon precision and axon recall as performance metrics. The recall was the higher than the precision for all groups, with the largest difference between the fine-tuned model with nnU-Net adjustments (99.1% recall and 50% precision. The original model resulted in the highest precision and accuracy while the fine-tuned model with nnU-net adjustments has the highest recall (Table 1).

Table 1. Accuracy and Precision

Metrics	*Original*	Fine-Tuned	nnU-net Fine-Tuned	Overlap - Fine-Tuned
Adjusted Accuracy	91.8%	91.5%	87.2%	88.7%
Axon Precision	64.1%	61.4%	50.0%	61.5%
Axon Recall	81.4%	89.5%	99.1%	30.2%

The visualizations of the predictions of the models using the original model, the fine-tuned model using the TrailMap code, and the fine-tuned model using the TrailMap code with the nnU-Net augmentation shows a dense predicted axon label (Figure 3).

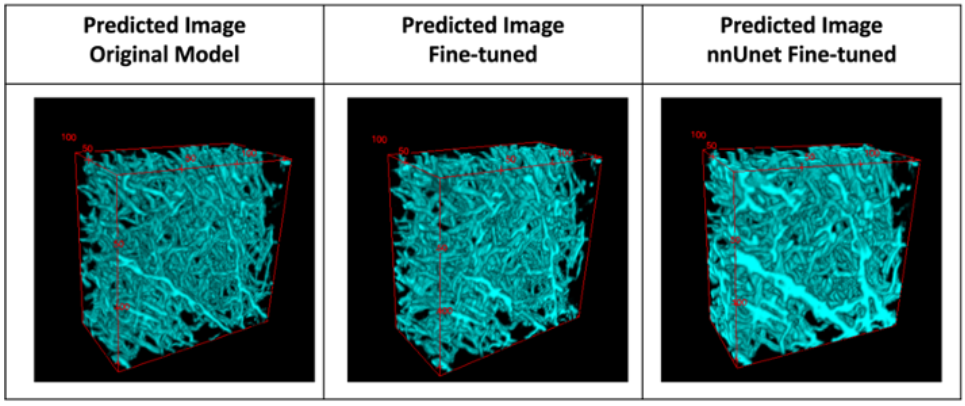

Figure 3. Data Inference: The results of inference using the original model, the fine-tuned model and the fine-tuned model with nnU-Net augmentation

The Gaussian overlapping did not improve results. Figure 4 shows inference with an overlap at 0% and 66% of the patch width. With a 66% width overlap for inference, the results appear patchy.

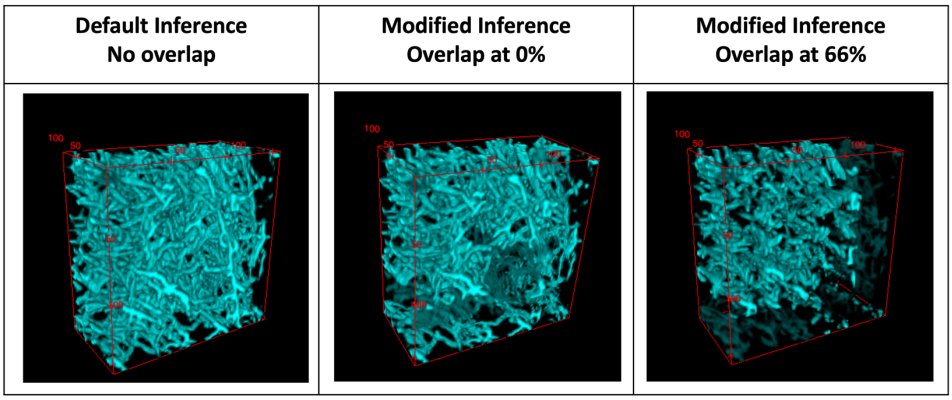

Figure 4. Data Inference

4.0 Discussion

Fine-tuning did not improve the model's accuracy. However, we only have one cube with 160 slices, which includes 6 labeled slices. When it's divided into training and testing data, each dataset only has 3 labeled slices. Hopefully when we have more data, the models will improve with fine-tuning. The images of the predicted axons also look too dense with axons.

We would also hope with more data that the nnU-Net augmentation will improve the accuracy of the model. The increase in performance with data augmentation is sometimes explained by the transformations mimicking the possible real-life transformations. For example, a study using mammogram images justified elastic deformation by the natural elastic deformation of a mammogram imaging process (Castro et al., 2018). It's possible that the neural imaging process also adds deformations, and that the most useful elastic deformation would be one mimics the real-life elastic transformation. Calculating or estimating this would be beyond the scope of this paper, but we will experiment with a range of standard deviations for the Gaussian kernel of the transformation.

When there is an overlap, it seems the Gaussian makes the predictions worse. A possible explanation for this is that TrailMap already has a unique solution to the problem of predicting edge cases. TrailMap models are trained with an image and label pairing where the label has its edges cropped off and the inference is likewise already predicted in patches where the image patch is bigger than the label patch. This means that the input images overlap but the label patches do not. This might be why the improvements were Gaussian were not needed and, in fact, decreased performance. From Figure 4, an overlap at 66% of the patch lead to fragmented neurons.

Technologies such as TrailMap contribute to the growing effort to develop templates to integrate data across multiple non-disruptive 3-dimensional image modalities, combining insights from light-sheet microscopy and MRI into a common coordinate framework (Perens et al., 2022). Such co-registration is essential for studying disease progression, developing effective therapies, translating findings into clinical procedures, for example, delivering brain site-specific drugs and implanting electrodes precisely for recording or stimulation.
While the work in this study demonstrates improvements in the ability to automate axon segmentation, the problem of complete mapping of neuronal structures into a co-registered system will likely require many structure-specific solutions in the short-term. The segmentation of tightly packed neurites will not benefit from an adaptive enhancement of the signal-to-noise ratios and may require other kinds of algorithmic slights of hand to feign higher resolution (Zhou et al., 2022).

5.0 References

Agarwal, Nitin, Xiangmin Xu, and Meenakshisundaram Gopi. “Geometry Processing of Conventionally Produced Mouse Brain Slice Images.” *Journal of Neuroscience Methods* 306 (2018): 45–56. https://doi.org/10.1016/j. jneumeth.2018.04.008.

Anantharaman, Rajaram, Matthew Velazquez, and Yugyung Lee. “Utilizing Mask R-CNN for Detection and Segmentation of Oral Diseases.” In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2197–2204. IEEE, 2018. https://doi.org/10.1109/BIBM.2018.8621112.

Bargmann, Cornelia I, and Eve Marder. “From the Connectome to Brain Function.” *Nature Methods* 10, no. 6 (2013): 483–490. https://doi.org/10.1038/nmeth.2451.

Caruana, Rich. “Learning Many Related Tasks at the Same Time with Backpropagation.” *Advances in Neural Information Processing Systems* 7 (1994): 657–664.

Castro, Eduardo, Jaime S Cardoso, and Jose Costa Pereira. “Elastic Deformations for Data Augmentation in Breast Cancer Mass Detection.” In *2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, 230–234. IEEE, 2018. https://doi.org/10.1109/BHI.2018.8333411.

Friedmann, Drew, Albert Pun, Eliza L Adams, Jan H Lui, Justus M Kebschull, Sophie M Grutzner, Caitlin Castagnola, Marc Tessier-Lavigne, and Liqun Luo. “Mapping Mesoscale Axonal Projections in the Mouse Brain Using a 3D Convolutional Network.” *Proceedings of the National Academy of Sciences* 117, no. 20 (2020): 11068–11075. https://doi.org/10.1073/pnas.1918465117.

Girkin, John M, and Mariana Torres Carvalho. “The Light-Sheet Microscopy Revolution.” *Journal of Optics* 20, no. 5 (2018): 1–20. https://doi.org/10.1088/2040-8986/aab58a.

He, Kaiming, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. “Mask R-CNN.” In *Proceedings of the IEEE International Conference on Computer Vision*, 2961–2969. 2017.

Isensee, Fabian, Paul Jäger, Jakob Wasserthal, David Zimmerer, Jens Petersen, Simon Kohl, Justus Schock, et al. *Batchgenerators - A Python Framework for Data Augmentation*. Software, 2020. https://doi.org/10.5281/zenodo.3632567.

Isensee, Fabian, Jens Petersen, Simon AA Kohl, Paul F Jäger, and Klaus H Maier-Hein. *nnU-Net: Breaking the Spell on Successful Medical Image Segmentation*. ArXiv:1904.08128 preprint published 17 April 2019.

Johnson, Jeremiah W. *Adapting Mask-RCNN for Automatic Nucleus Segmentation*. ArXiv:1805.00500 preprint published 1 May 2018.

Kasthuri, Narayanan, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, Richard Lee Schalek, José Angel Conchello, Seymour Knowles-Barley, Dongil Lee, Amelio Vázquez-Reina, Verena Kaynig, Thouis Raymond Jones, et al. “Saturated Reconstruction of a Volume of Neocortex.” *Cell* 162, no. 3 (2015): 648–661. https://doi.org/10.1016/j.cell.2015.06.054.

Kaur, Dilpreet, and Yadwinder Kaur. “Various Image Segmentation Techniques: A Review.” *International Journal of Computer Science and Mobile Computing* 3, no. 5 (2014): 809–814.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton. “ImageNet Classification with Deep Convolutional Neural Networks.” *Advances in Neural Information Processing Systems* 25 (2012).

Olarte, Omar E, Jordi Andilla, Emilio J Gualda, and Pablo Loza-Alvarez. “Light-Sheet Microscopy: A Tutorial.” *Advances in Optics and Photonics* 10, no. 1 (2018): 111–179. https://doi.org/10.1364/AOP.10.000111.

Perens, Johanna, Casper Gravesen Salinas, Urmas Roostalu, Jacob Lercke Skytte, Carsten Gundlach, Jacob Hecksher-Sørensen, Anders Bjorholm Dahl, and Tim B Dyrby. *Multimodal 3D Mouse Brain Atlas Framework with Skull-Derived Coordinate System*. Research Square preprint published 21 July 2022. https://doi.org/10.21203/rs.3.rs-1832101/v1.

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. “U-net: Convolutional Networks for Biomedical Image Segmentation.” In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer, 2015.

Santi, Peter A. “Light Sheet Fluorescence Microscopy: A Review.” *Journal of Histochemistry & Cytochemistry* 59, no. 2 (2011): 129–138. https: //doi. org/10.1369/0022155410394857.

Seiriki, Kaoru, Atsushi Kasai, Takeshi Hashimoto, Wiebke Schulze, Misaki Niu, Shun Yamaguchi, Takanobu Nakazawa, Ken-ichi Inoue, Shiori Uezono, Masahiko Takada, et al. “High-Speed and Scalable Whole-Brain Imaging in Rodents and Primates.” *Neuron* 94, no. 6 (2017): 1085–1100. https://doi.org/10.1016/j.neuron.2017.05.017.

Sporns, Olaf. “The Human Connectome: A Complex Network.” *Annals of the New York Academy of Sciences* 1224, no. 1 (2011): 109–125. https://doi.org/10.1111/j.1749-6632.2010.05888.x.

Tyson, Adam L, and Troy W Margrie. “Mesoscale Microscopy and Image Analysis Tools for Understanding the Brain.” *Progress in Biophysics and Molecular Biology* 168 (2022): 81–93. https://doi.org/0.1016/j.pbiomolbio.2021.06.013.

Zhang, Yu-Jin. “An Overview of Image and Video Segmentation in the Last 40 Years.” *Advances in Image and Video Segmentation*, 2006, 1–16. https:/doi.org/10.4018/978-1-59140-753-9.ch001.

Zhou, Hang, Tingting Cao, Tian Liu, Shijie Liu, Lu Chen, Yijun Chen, Qing Huang, Wei Ye, Shaoqun Zeng, and Tingwei Quan. “Super-resolution Segmentation Network for Reconstruction of Packed Neurites.” *Neuroinformatics*, 2022, 1–13. https://doi.org/10.1007/s12021-022-09594-3.

Zingg, Brian, Houri Hintiryan, Lin Gou, Monica Y Song, Maxwell Bay, Michael S Bienkowski, Nicholas N Foster, Seita Yamashita, Ian Bowman, Arthur W Toga, et al. “Neural Networks of the Mouse Neocortex.” *Cell* 156, no. 5 (2014): 1096–1111. https://doi.org/10.1016/j.cell.2014.02.023.

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354

1-888-375-PNNL (7665)

www.pnnl.gov