Error Correction for Connectomics

Brian Matejek

Harvard John A. Paulson
School of Engineering and
Applied Sciences

Visual
Computing
Group

Connectomics

Goal: Extract the wiring diagram from a brain

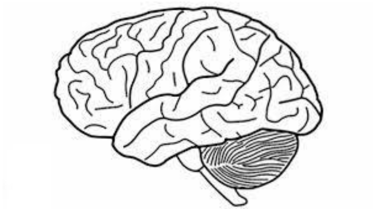

Connectomics

Goal: Extract the wiring diagram from a brain

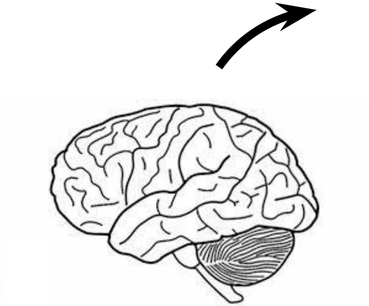

Nano-resolution Imaging

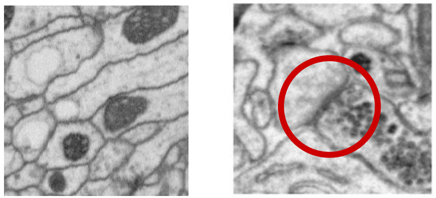

Connectomics

Goal: Extract the wiring diagram from a brain

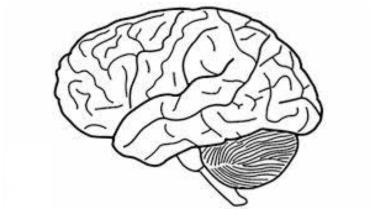

Nano-resolution Imaging

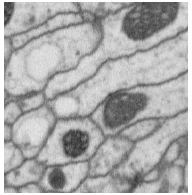

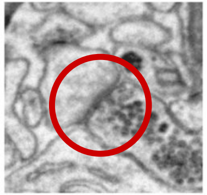

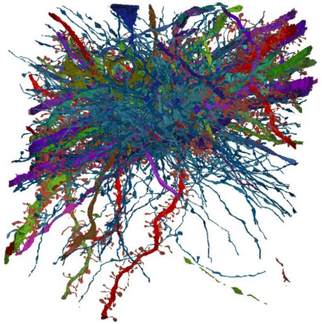

Circuit Reconstruction

Connectomics

Goal: Extract the wiring diagram from a brain

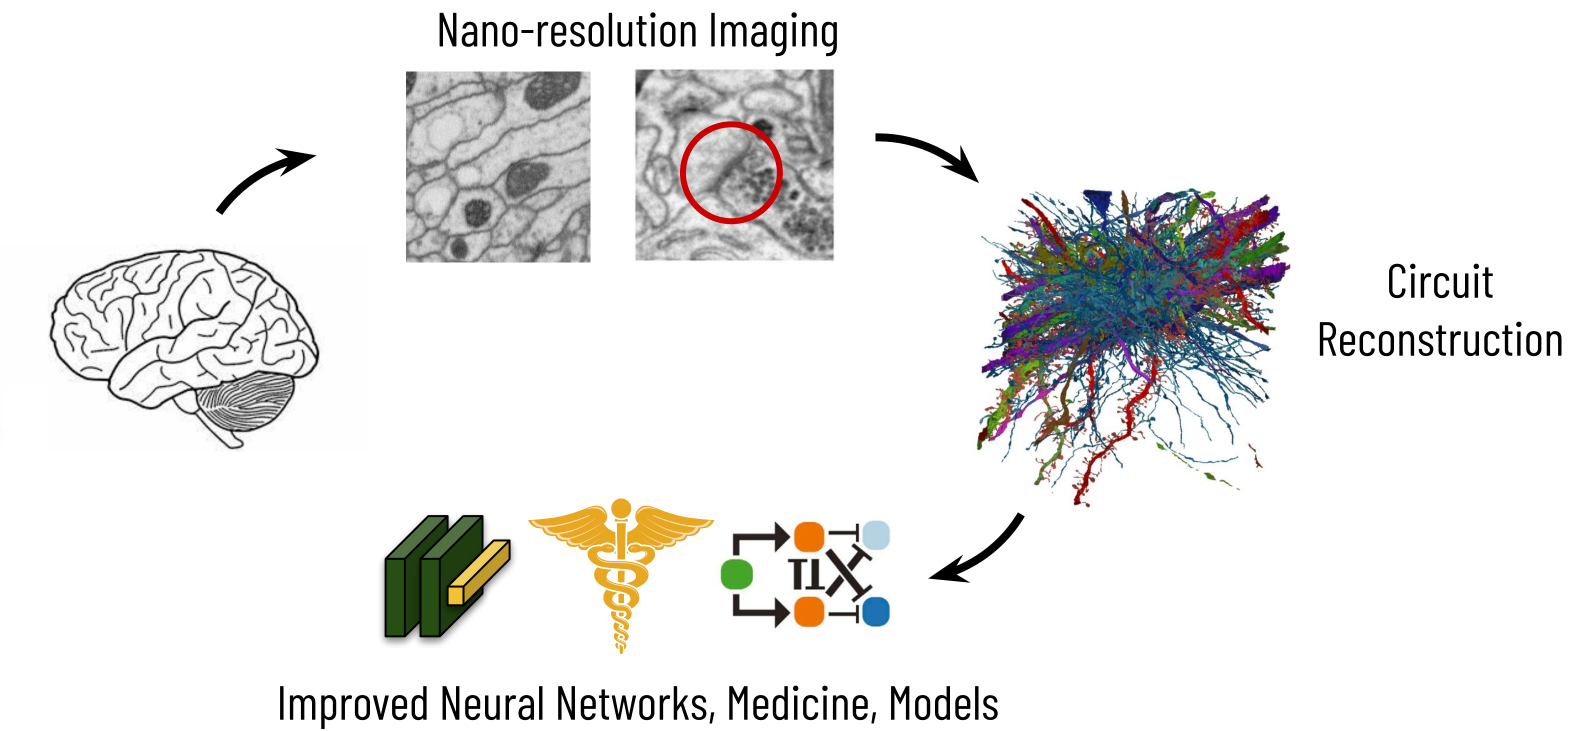

Connectomics

Goal: Extract the wiring diagram from a brain

Nano-resolution Imaging

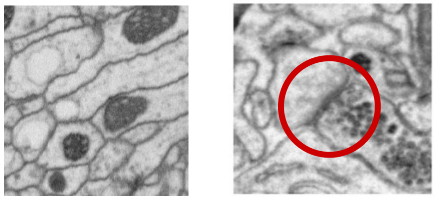

Circuit Reconstruction

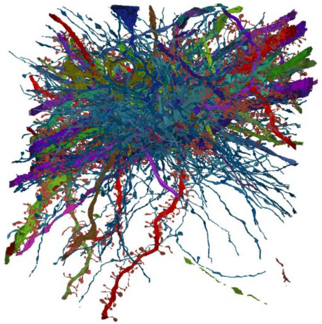

Improved Neural Networks, Medicine, Models

Better Understanding

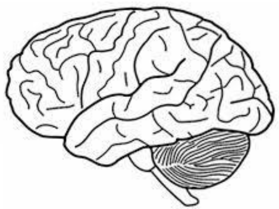

Connectomics

Goal: Extract the wiring diagram from a brain

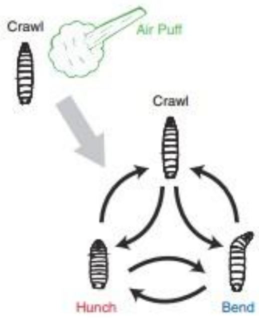

Behavior

Jovanic et al., Competitive Disinhibition Mediates Behavioral Choice and Sequences in Drosophila, Cell 2016

Connectomics

Goal: Extract the wiring diagram from a brain

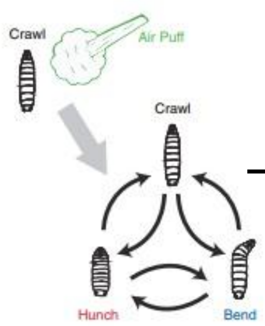

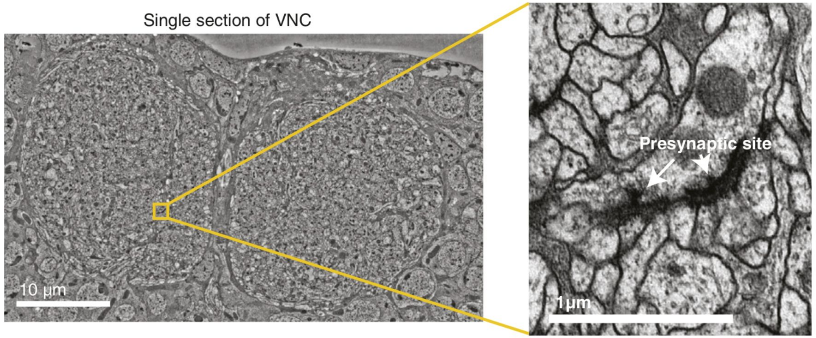

Behavior

Structure

Jovanic et al., Competitive Disinhibition Mediates Behavioral Choice and Sequences in Drosophila, Cell 2016

Connectomics

Goal: Extract the wiring diagram from a brain

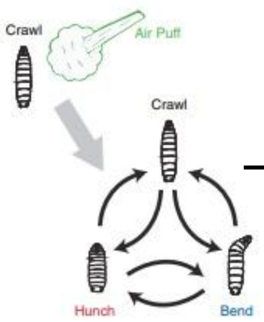

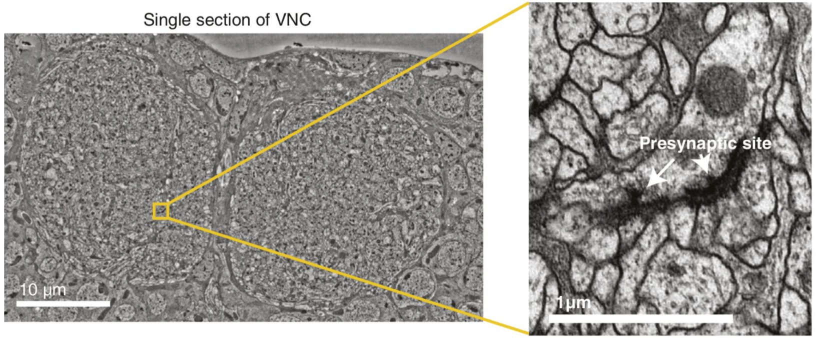

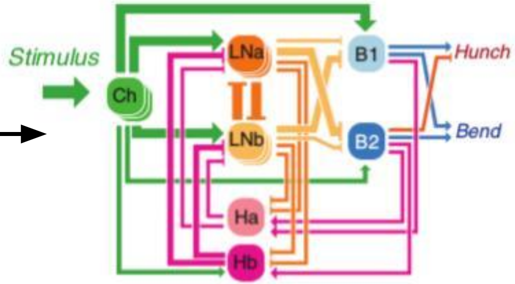

Behavior

Structure

Function

Jovanic et al., Competitive Disinhibition Mediates Behavioral Choice and Sequences in Drosophila, Cell 2016

Connectomics

Goal: Extract the wiring diagram from a brain

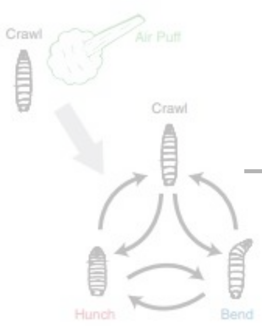

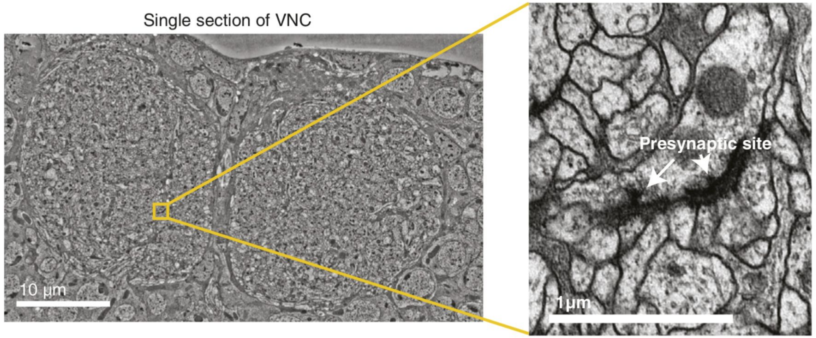

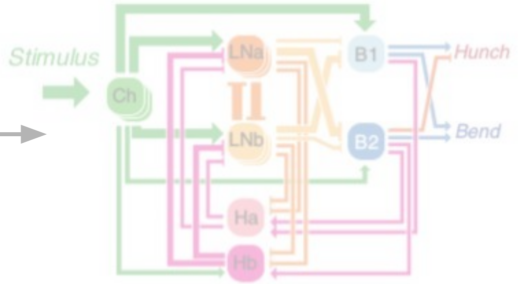

Behavior

Structure

Function

Jovanic et al., Competitive Disinhibition Mediates Behavioral Choice and Sequences in Drosophila, Cell 2016

Connectomics Pipeline

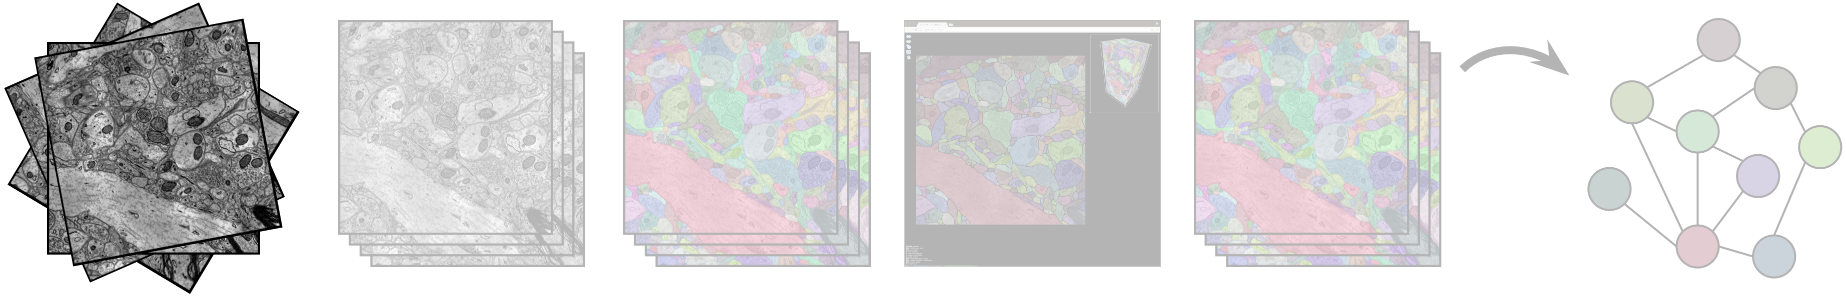

Acquisition Registration Segmentation Proofreading Network Analysis

Suissa-Peleg et al., Automatic Neural Reconstruction from Petavoxel of Electron Microscopy, Microscopy and Microanalysis 2016
Schalek et al., Imaging a 1 mm^3 Volume of Rat Cortex Using a MultiBeam SEM, Microscopy and Microanalysis, 2016

Image Acquisition

Multi-beam electron microscopes collect 1 TB of raw image data every hour

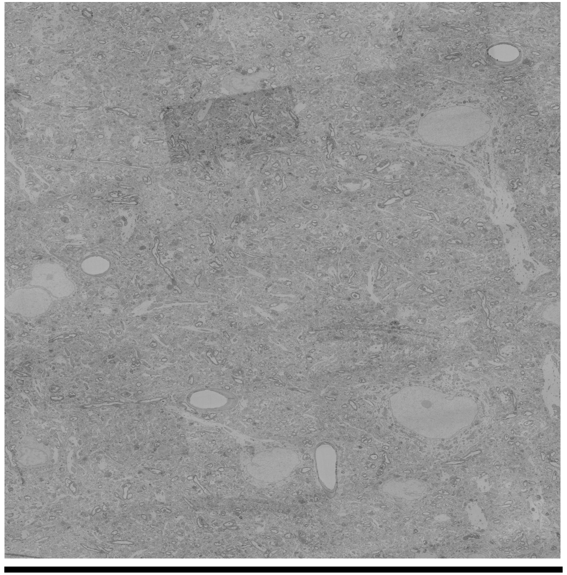

100 μm

Image Acquisition

Multi-beam electron microscopes collect 1 TB of raw image data every hour

Can image 1 mm^3 of image data (2 PB) in 6 months

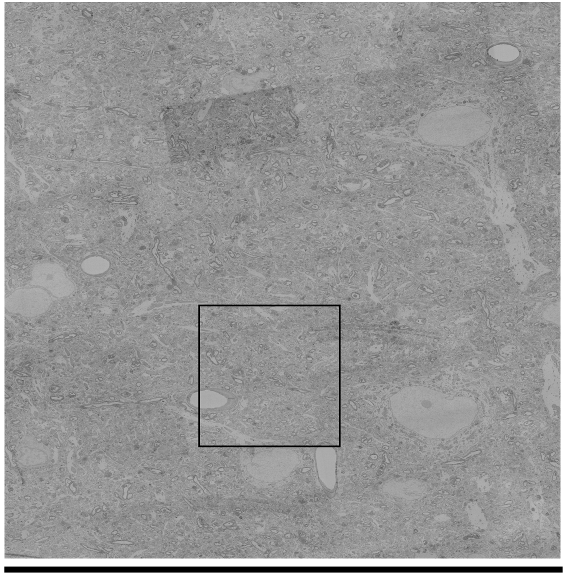

100 μm

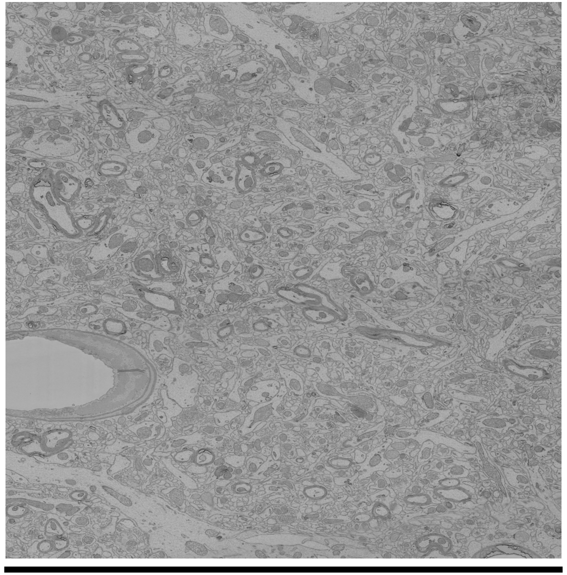

25 μm

Image Acquisition

Multi-beam electron microscopes collect 1 TB of raw image data every hour

Can image 1 mm^3 of image data (2 PB) in 6 months

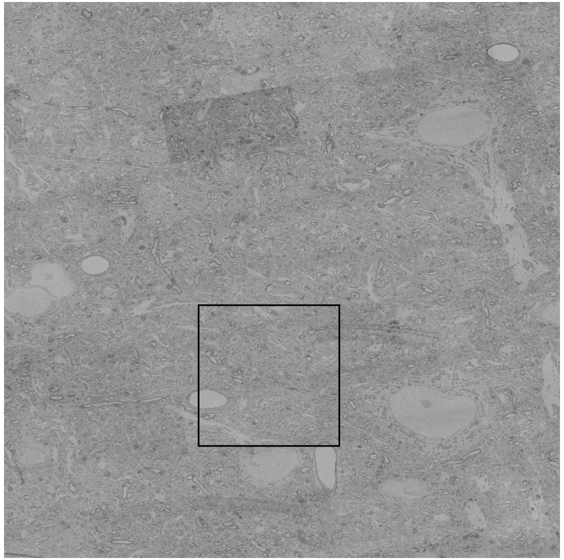

100 μm

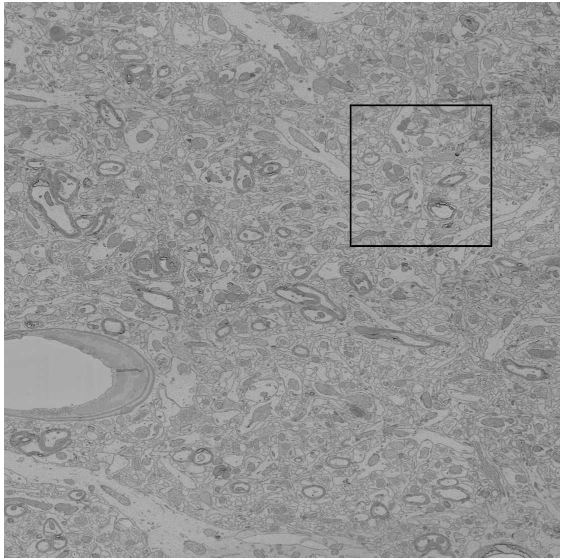

25 μm

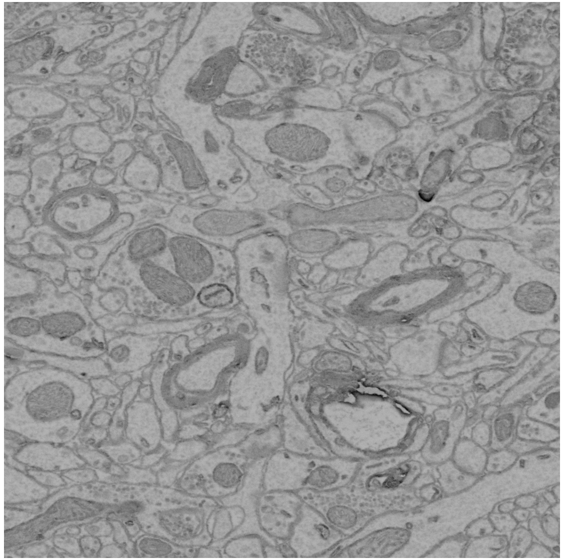

6250 nm

Connectomics Pipeline

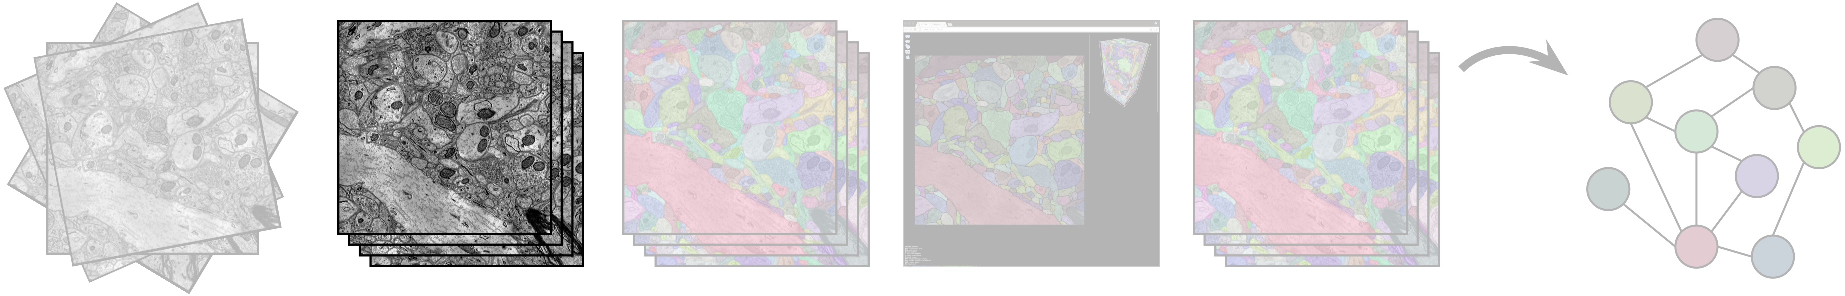

Acquisition Registration Segmentation Proofreading Network Analysis

Saalfield et al., Elastic Volume Reconstruction from Series of Ultra-thin Microscopy Sections, Nature 2012
Khairy et al., Joint Deformable Registration of Large EM Image Volumes: A Matrix Solver Approach, 2018

Registration

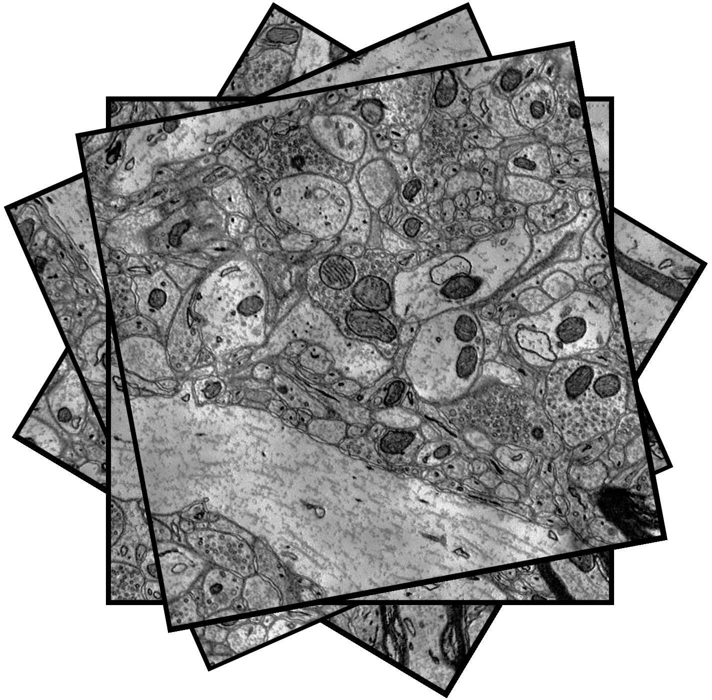

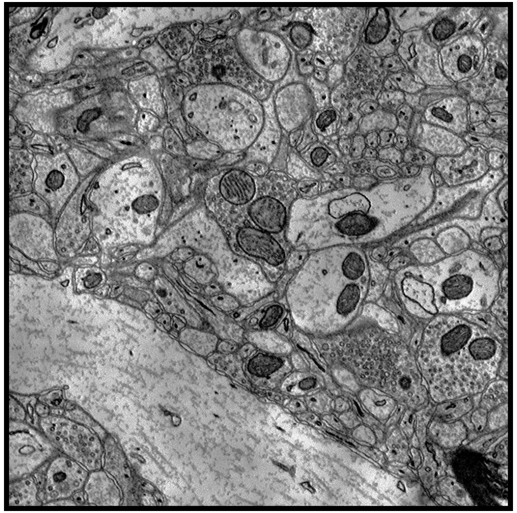

Connectomics Pipeline

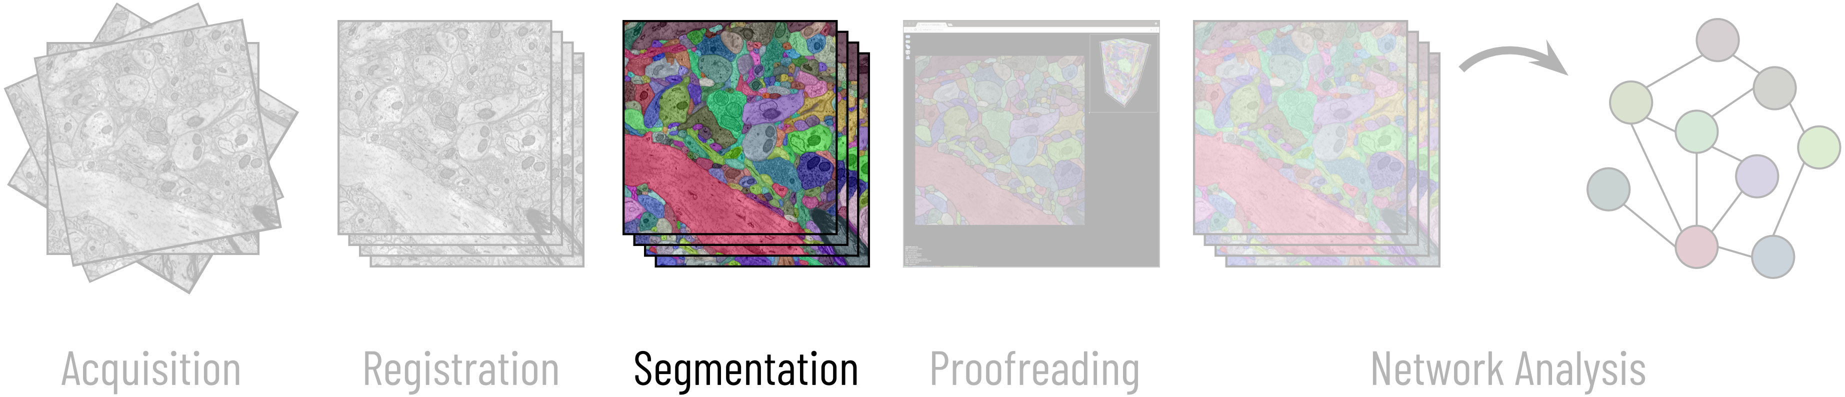

Nunez-Iglesias et al., Machine Learning of Hierarchical Clustering to Segment 2D and 3D Images, PLoS ONE 2014
Cicek et al., 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation, MICCAI 2016
Januszewski et al., Flood-Filling Networks, 2016
Zeng et al., DeepEM3D: Approaching Human-Level Performance on 3D Anisotropic EM Image Segmentation, Bioinformatics 2017
Pape et al., Solving Large Multicut Problems for Connectomics via Domain Decomposition, ICCV 2017
Lee et al., Superhuman Accuracy on the SNEMI3D Connectomics Challenge, 2017

Label Volumes

Two voxels have the same label only if they belong to the same neuron

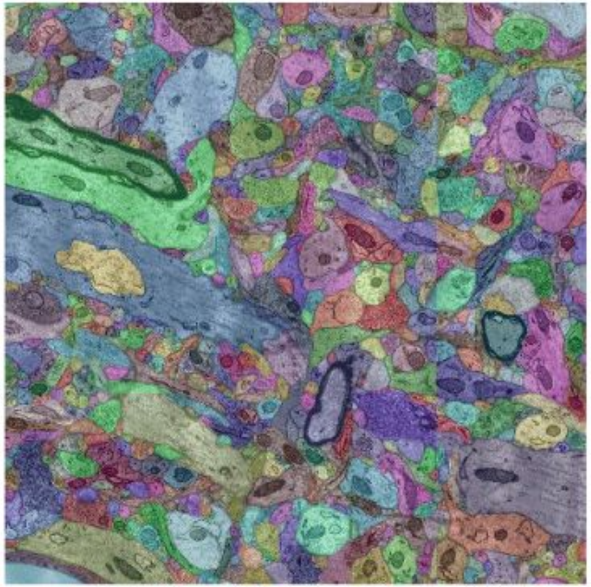

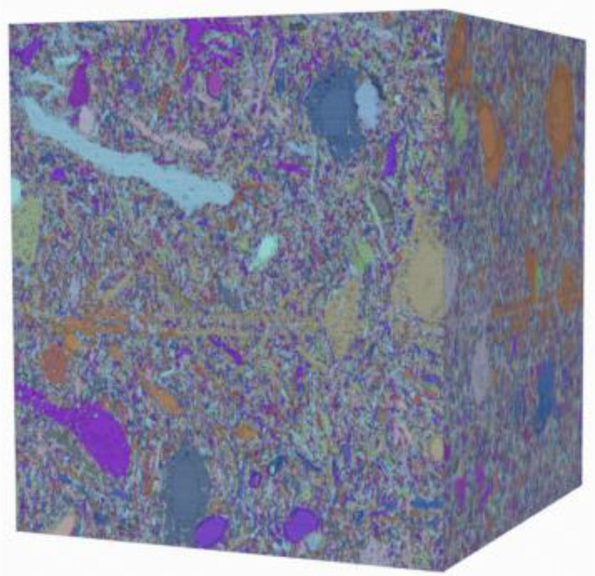

Label Volumes

Two voxels have the same label only if they belong to the same neuron

Typically use 32 or 64 bits per voxel to label each segment uniquely

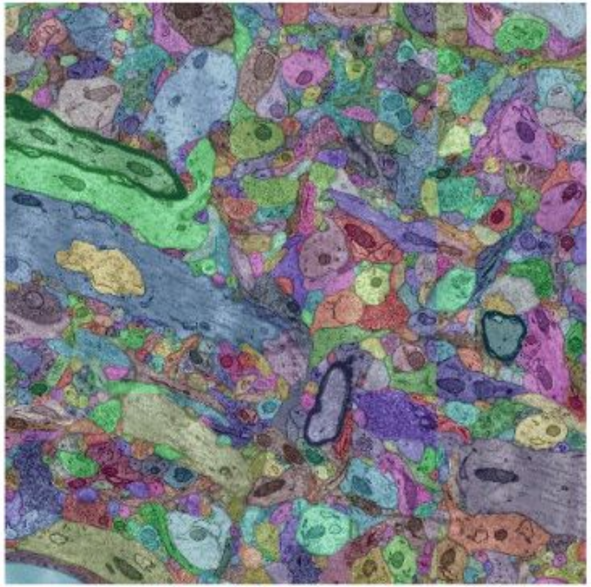

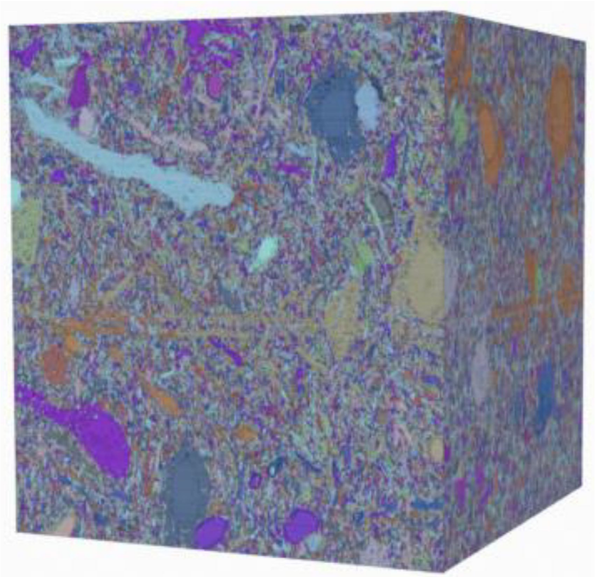

Connectomics Pipeline

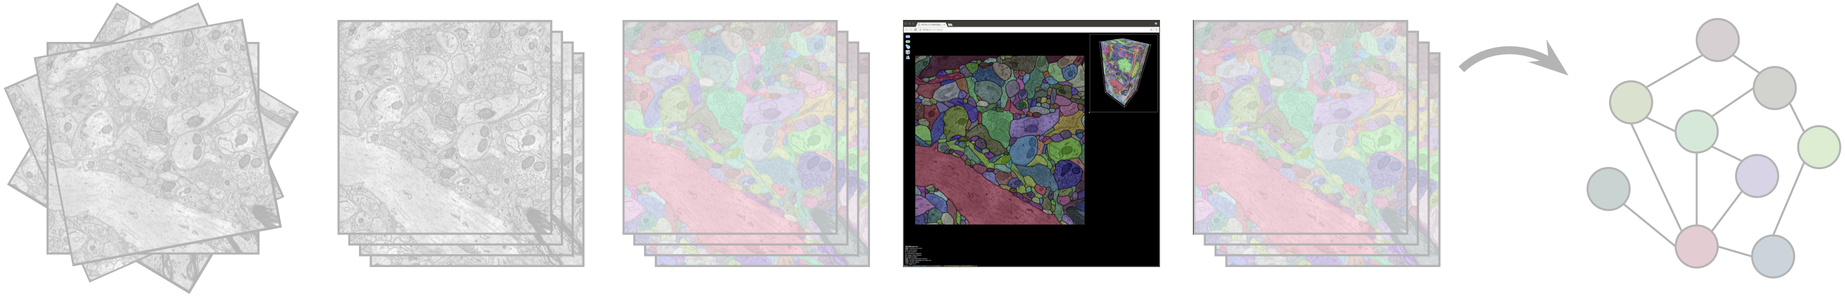

Acquisition Registration Segmentation **Proofreading** Network Analysis

Haehn et al., Design and Evaluation of Interactive Proofreading Tools for Connectomics, IEEE VIS 2014
Zung et al., An Error Detection and Correction Framework for Connectomics, NIPS 2017
Haehn et al., Guided Proofreading of Automatic Segmentations for Connectomics, CVPR 2018

Merge Errors

Automatic Segmentation

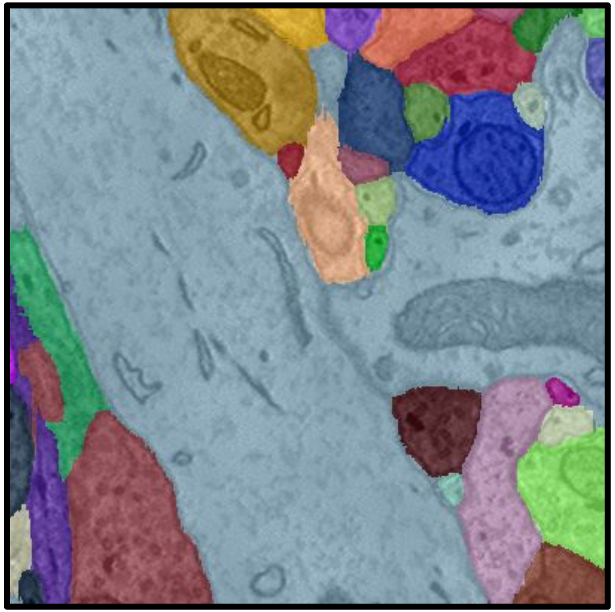

Ground Truth

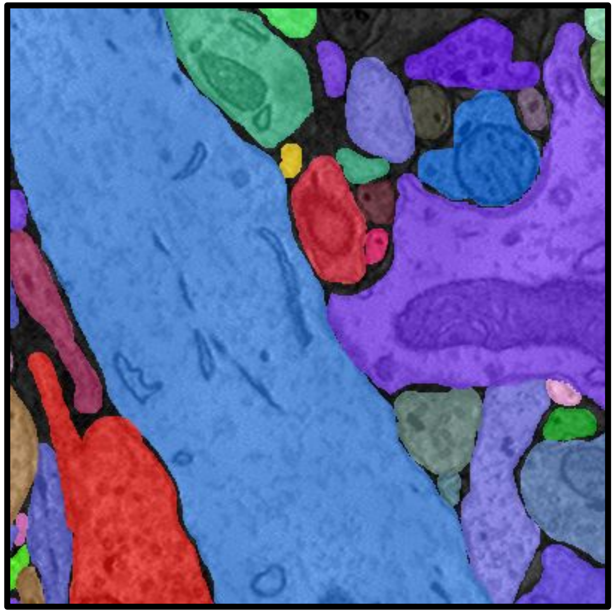

Merge Errors

Automatic Segmentation

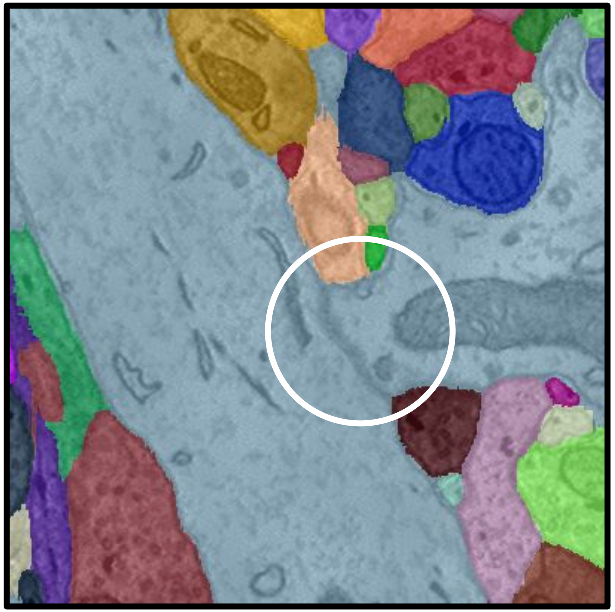

Ground Truth

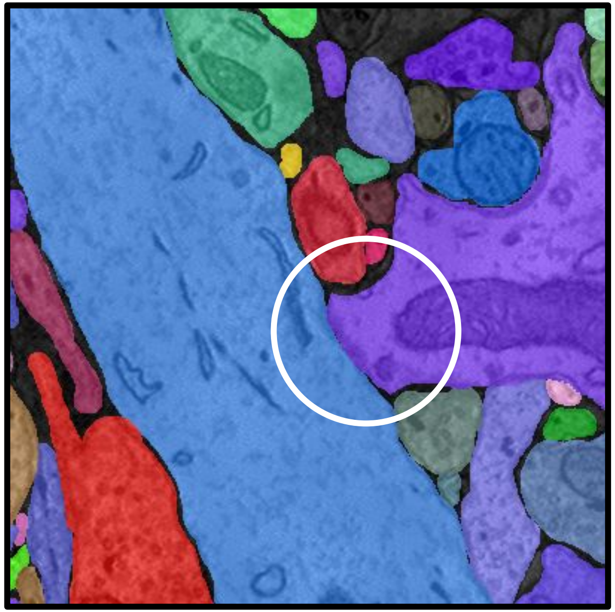

Split Errors

Automatic Segmentation

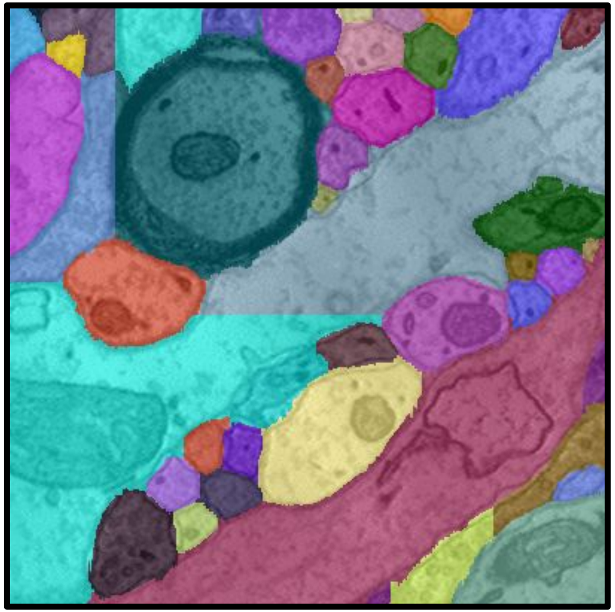

Ground Truth

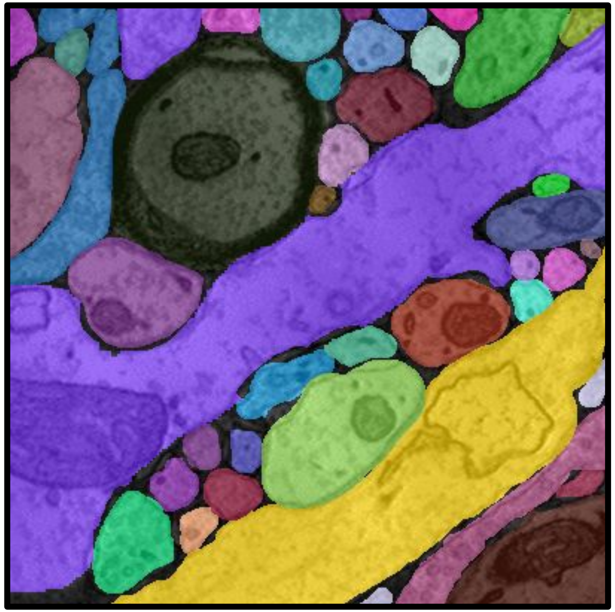

Split Errors

Automatic Segmentation

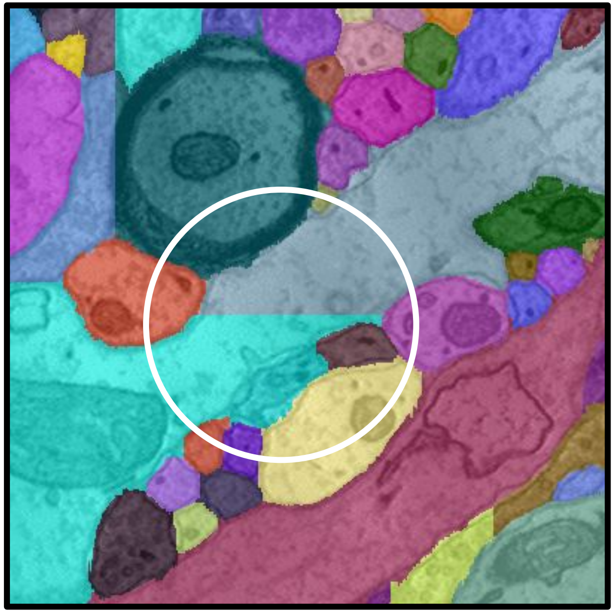

Ground Truth

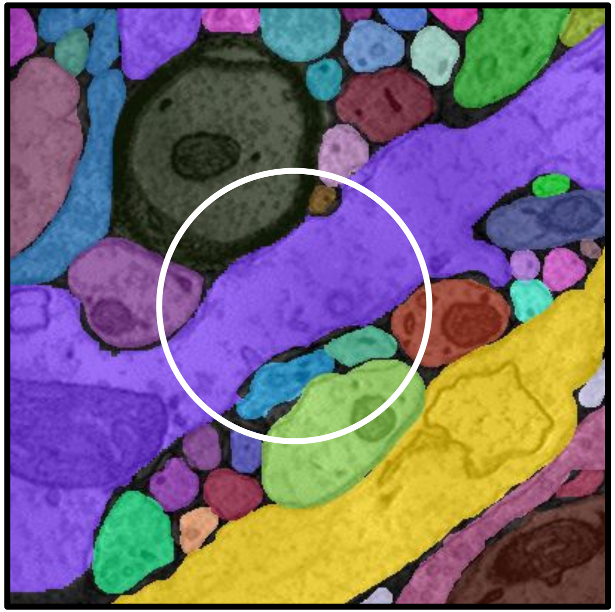

Connectomics Pipeline

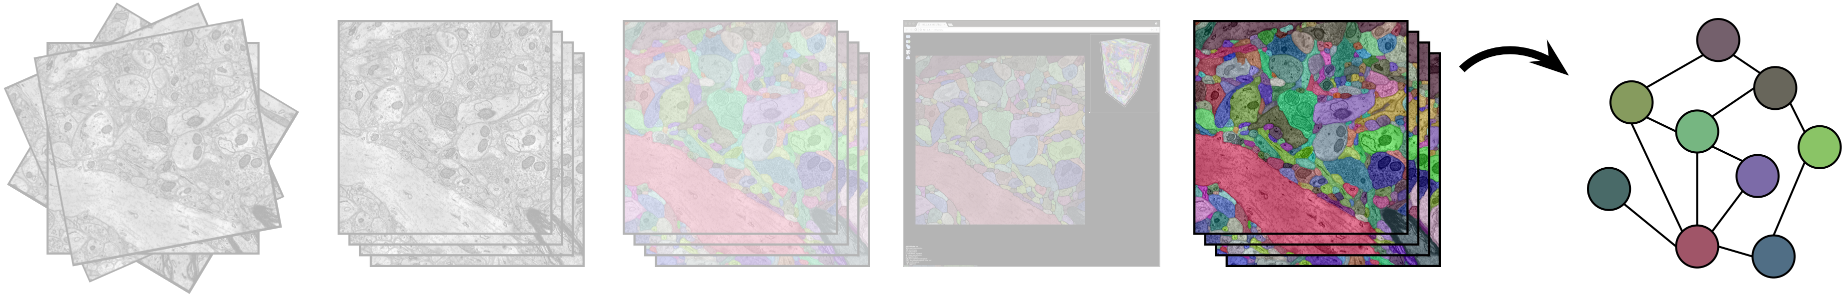

Acquisition Registration Segmentation Proofreading Network Analysis

Sorger et al., neuroMAP - Interactive Graph-Visualization of the Fruit Fly's Neural Circuit, BioVIS 2013
Al-Awami et al., NeuroLines: A Subway Map Metaphor for Visualizing Nanoscale Neuronal Connectivity, IEEE VIS 2014
Haehn et al., Scalable Interactive Visualization for Connectomics, MDPI Informatics 2017

Network Analysis

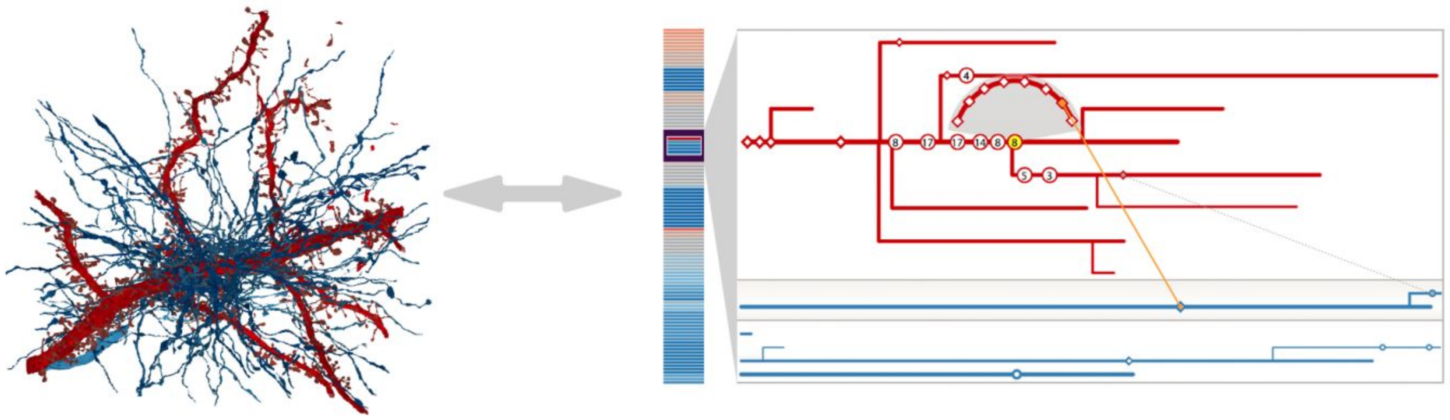

Connectomics Pipeline

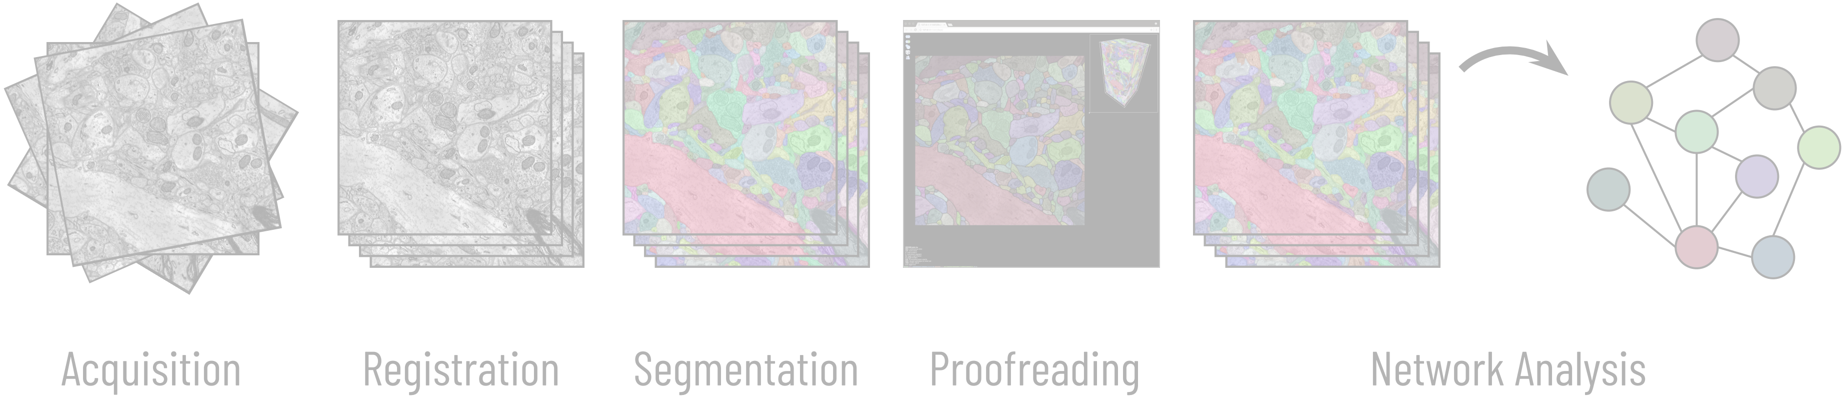

Connectomics Pipeline

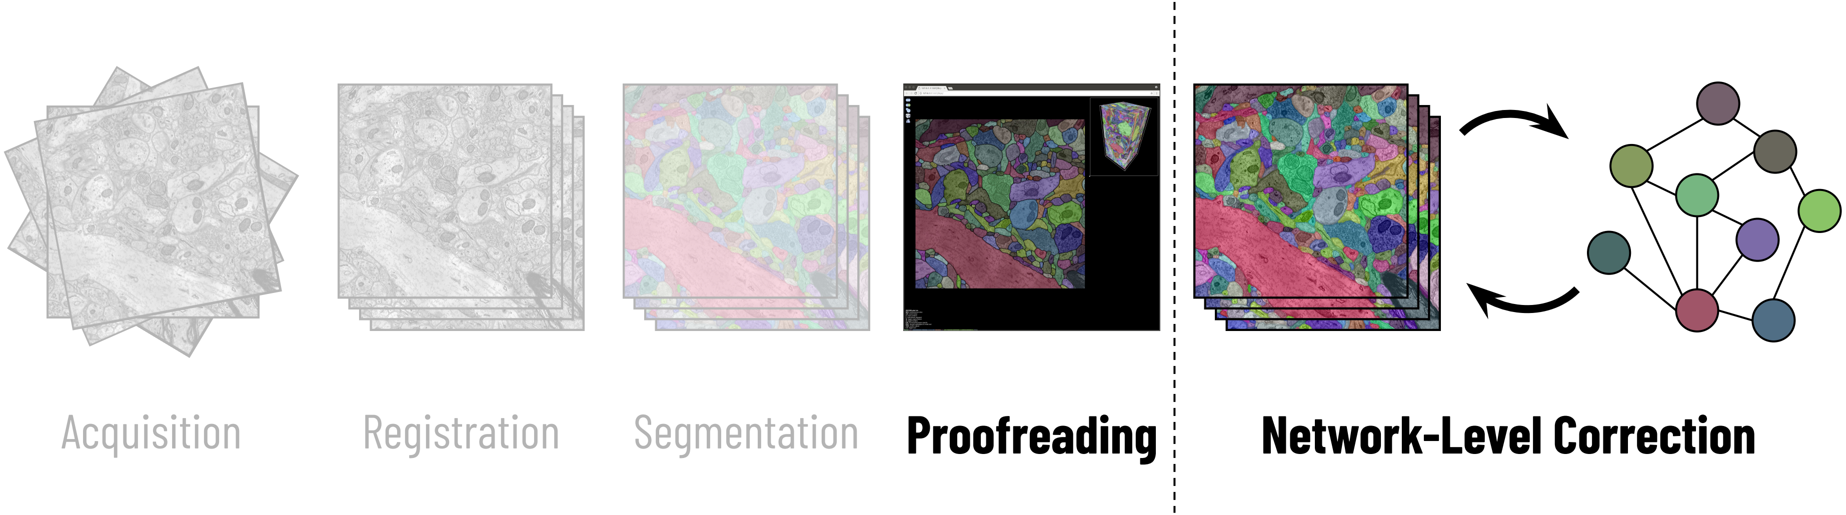

Biologically-Constrained Graphs for Global Connectomics Reconstruction

Brian Matejek[1], Daniel Haehn[1], Haidong Zhu[2], Donglai Wei[1], Toufiq Parag[3], Hanspeter Pfister[1]

[1]Harvard University

[2]Tsinghua University

[3]Comcast Research

Harvard John A. Paulson
School of Engineering and
Applied Sciences

Visual
Computing
Group

Affinity Generation

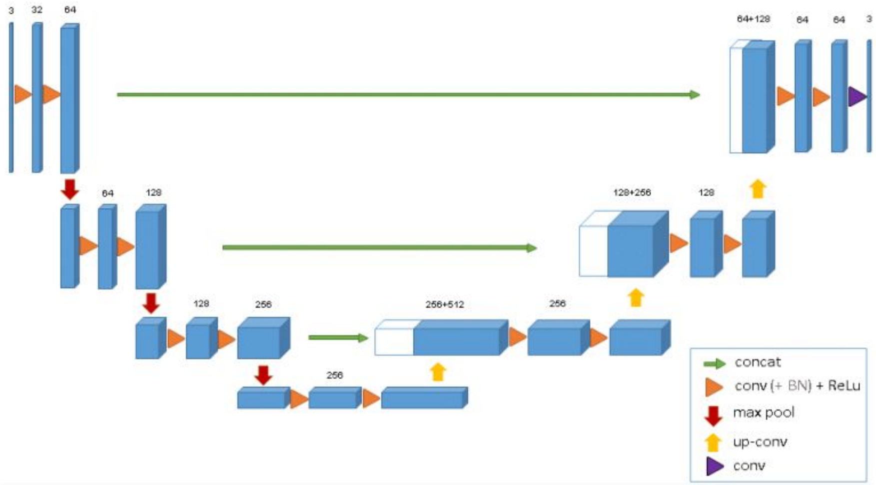

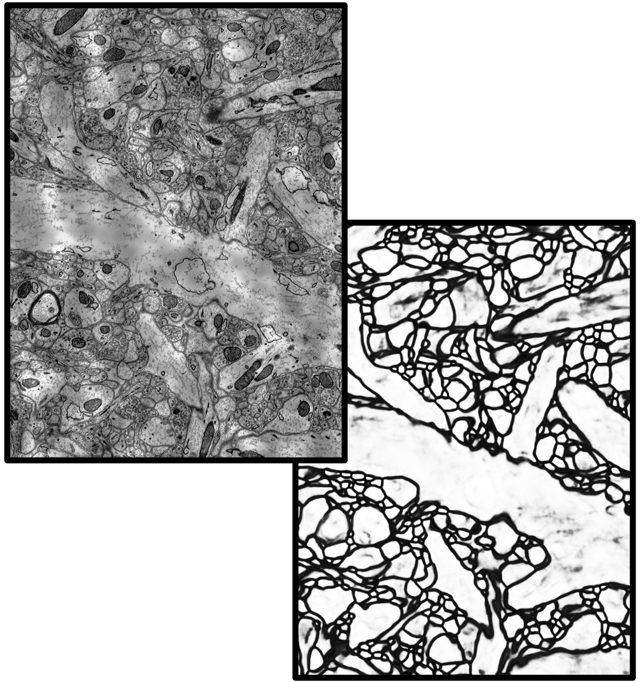

Ronneberger et al., U-Net: Convolutional Networks for Biomedical Image Segmentation, MICCAI 2015
Cicek et al., 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation, MICCAI 2016

3D Watershed on Affinities

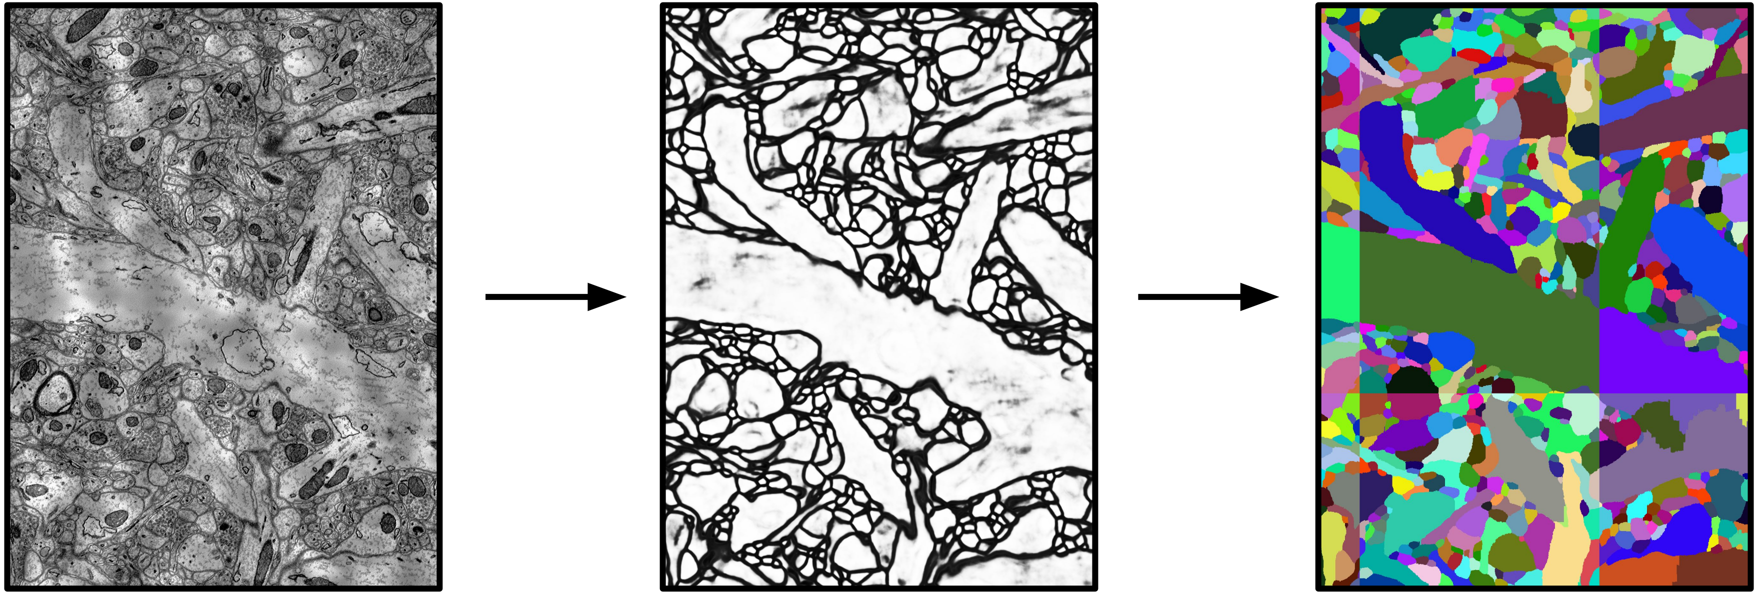

Zlateski et al., Image Segmentation by Size-Dependent Single Linkage Clustering of a Watershed Basin Graph, 2015
Funke et al., A Deep Structured Learning Approach Towards Automating Connectome Reconstruction from 3D Electron Micrographs, 2017
Zeng et al., DeepEM3D: Approaching Human-Level Performance on 3D Anisotropic EM Image Segmentation, Bioinformatics 2017

Agglomeration

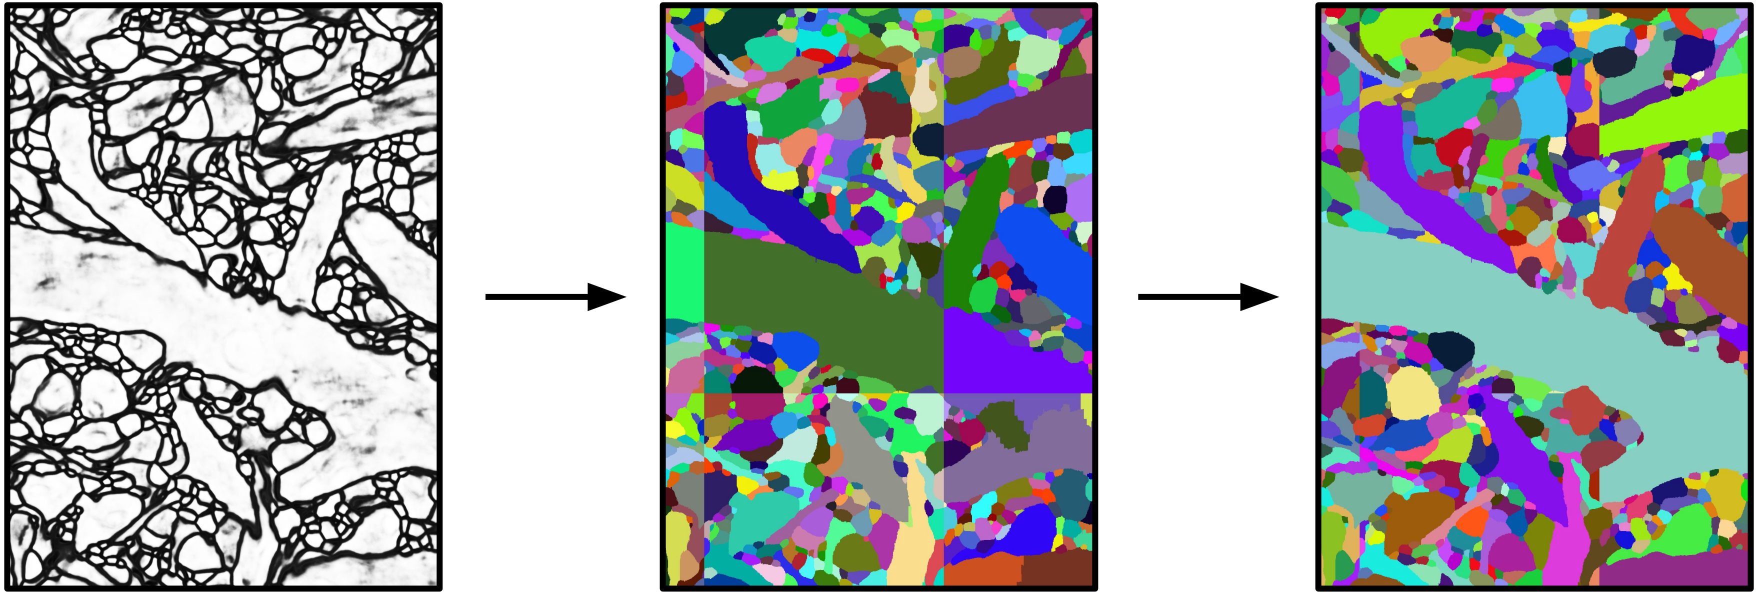

Nunez-Iglesias et al., Machine Learning of Hierarchical Clustering to Segment 2D and 3D Images, PLoS ONE, 2013
Parag et al., A Context-Aware Delayed Agglomeration Framework for Electron Microscopy Segmentation, PLoS ONE 2015
Funke et al., A Deep Structured Learning Approach Towards Automating Connectome Reconstruction from 3D Electron Micrographs, 2017

Lifted Multicuts

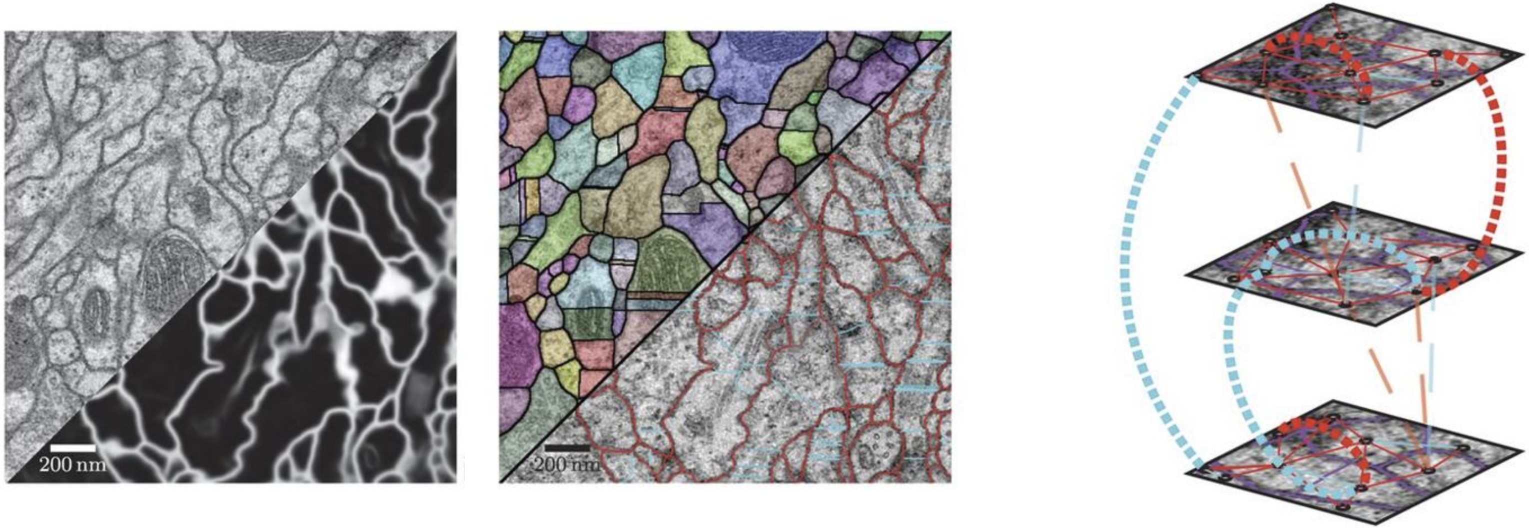

Beier et al., Multicut Brings Automated Neurite Segmentation Closer to Human Performance, Nature 2017

Errors

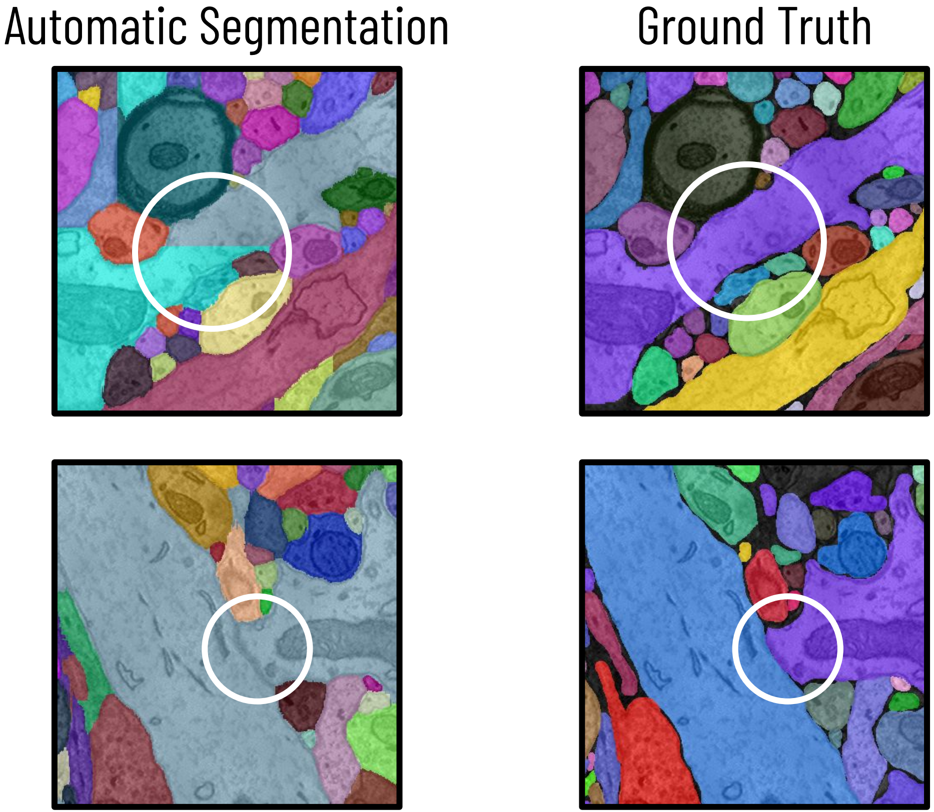

Guided Proofreading

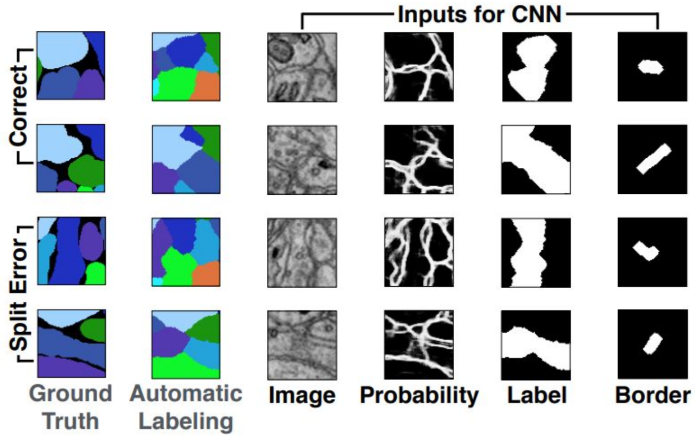

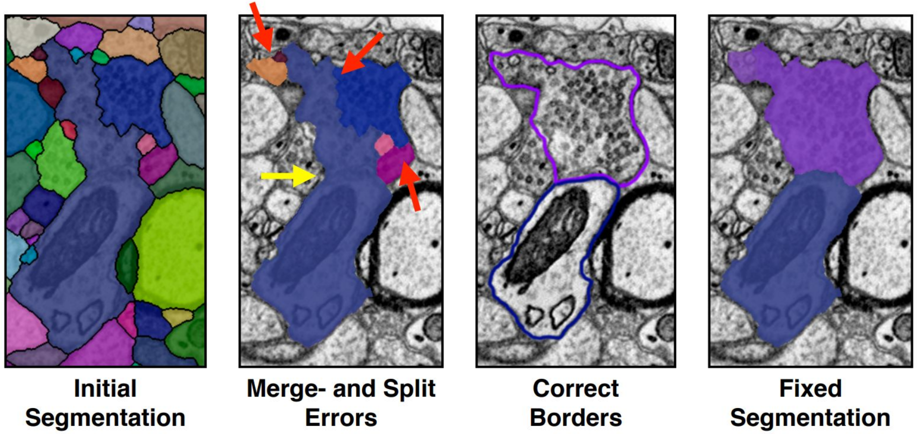

Haehn et al., Guided Proofreading of Automatic Segmentations for Connectomics, CVPR 2018

Automatic Proofreading

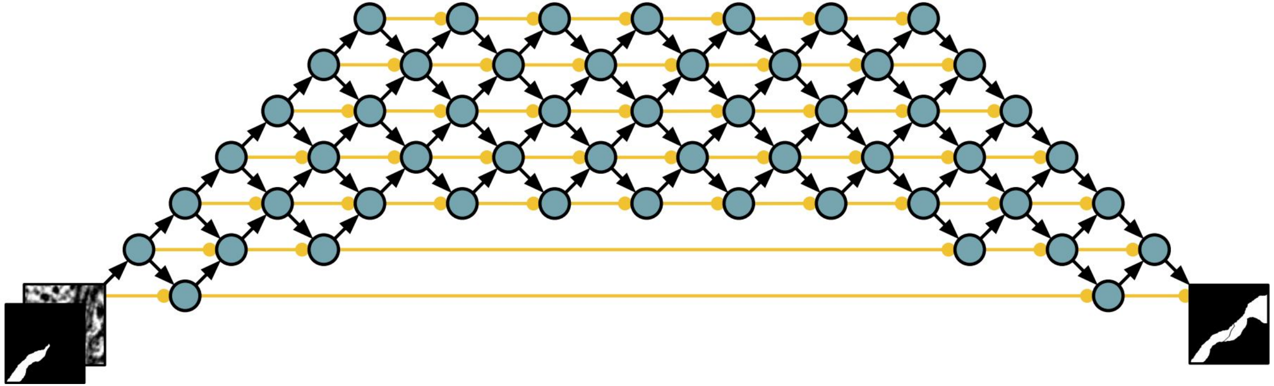

Zung et al., An Error Detection and Correction Framework for Connectomics, NIPS 2017

Proposed Automatic Error Correction

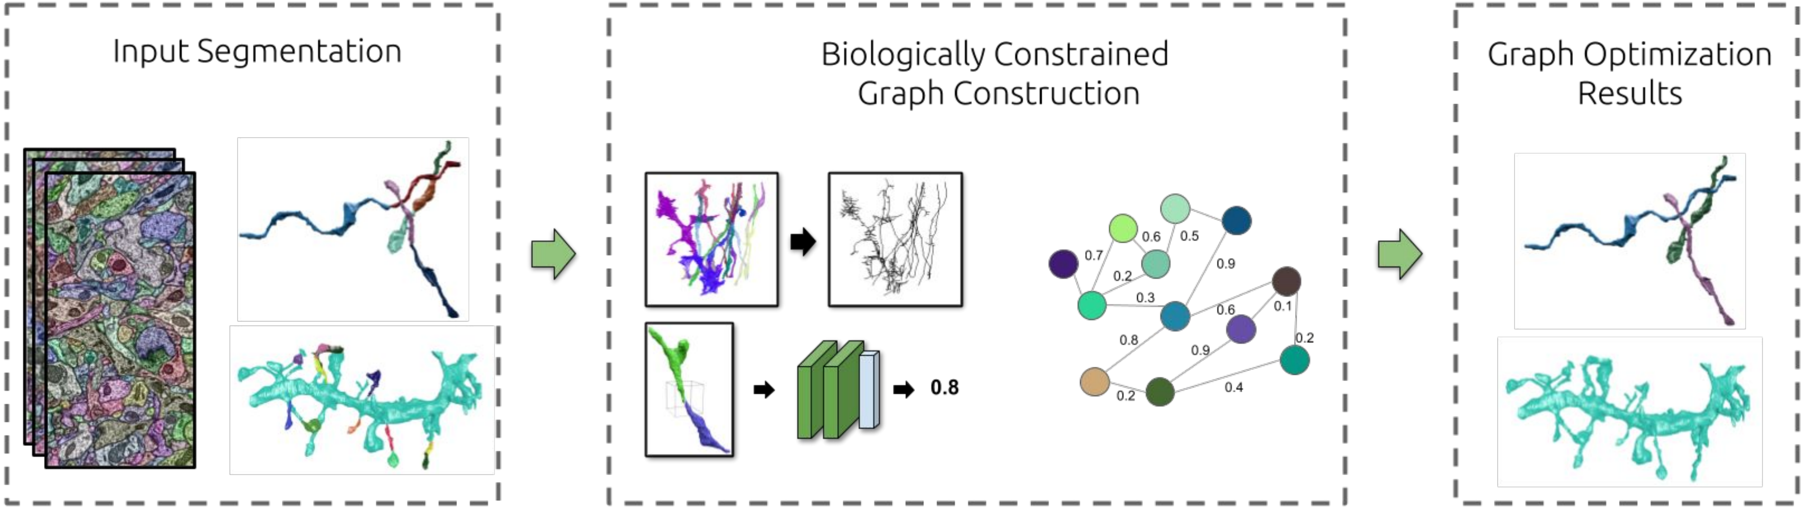

Input

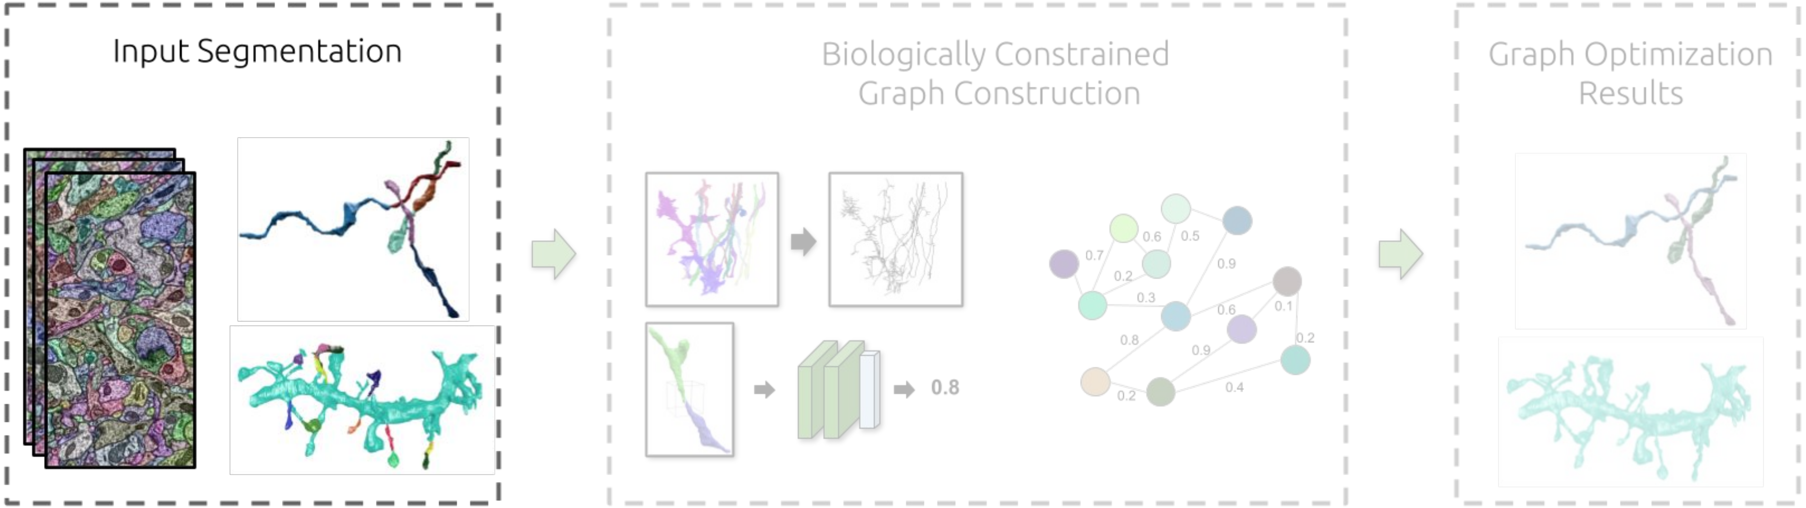
Input Segmentation
Biologically Constrained
Graph Construction
Graph Optimization
Results
0.8

Traditional Two-Stage Frameworks

Existing segmentation strategies typically produce over-segmentations

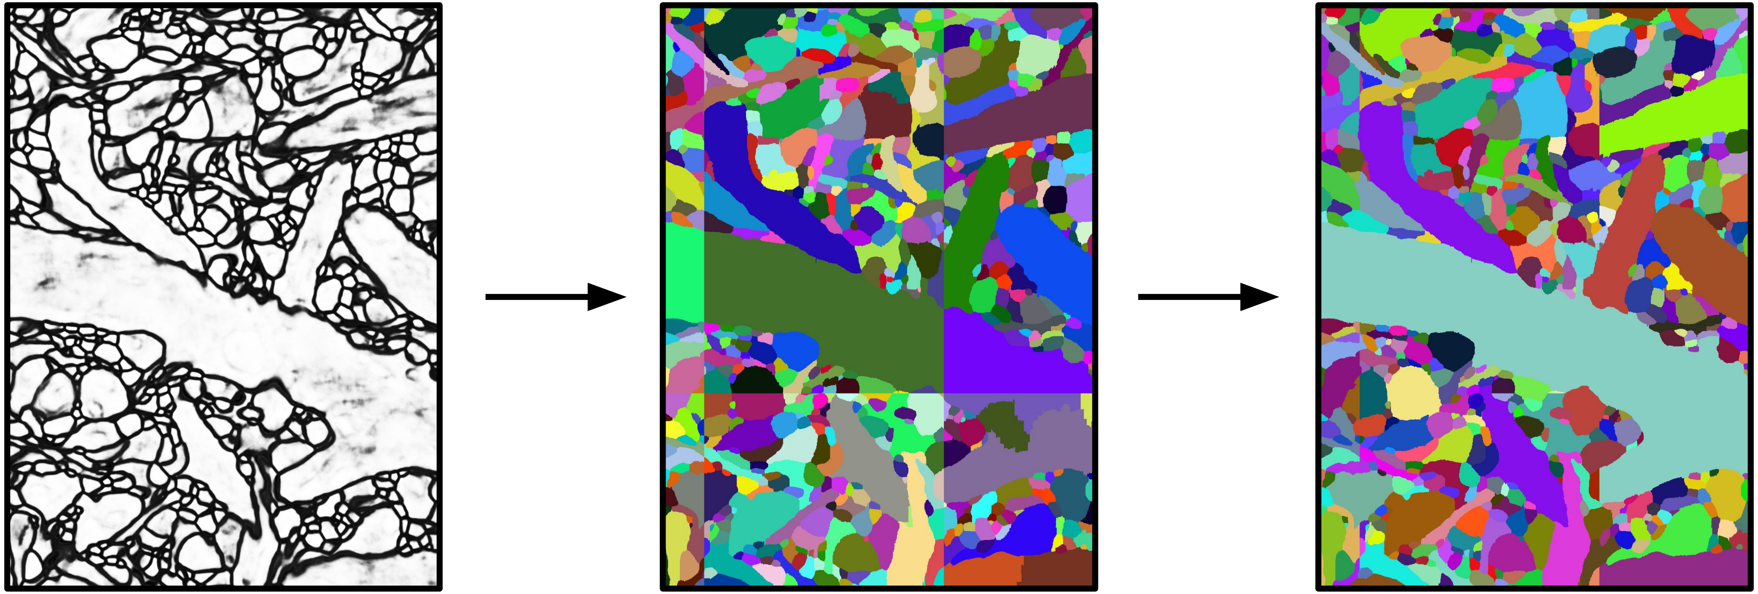

Affinity Generation | Watershed Transform | Agglomeration

Traditional Two-Stage Frameworks

Existing segmentation strategies typically produce over-segmentations

We use the result from an existing strategy as our input

Traditional Two-Stage Frameworks

Existing segmentation strategies typically produce over-segmentations

We use the result from an existing strategy as our input

Allows us to leverage larger local context when forming our graph

Goal: Construct a graph with as few nodes and edges as possible

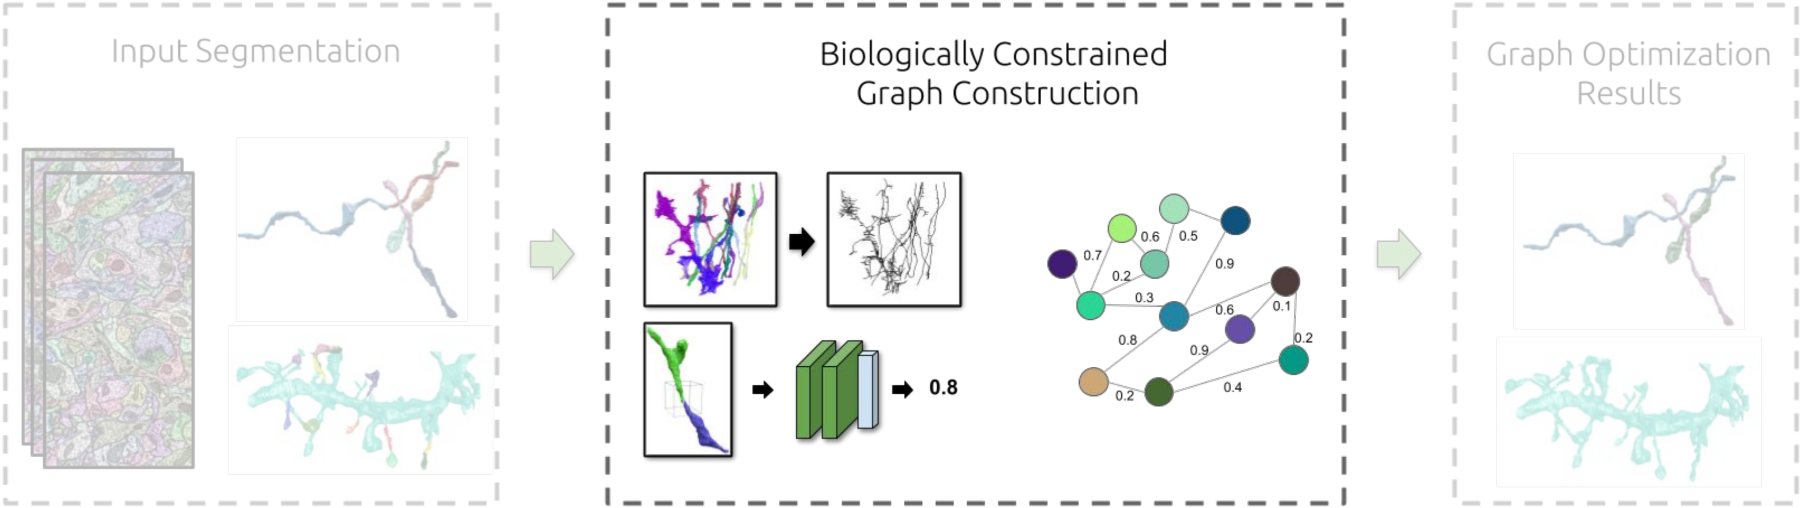

Graph Construction with Biological Constraints

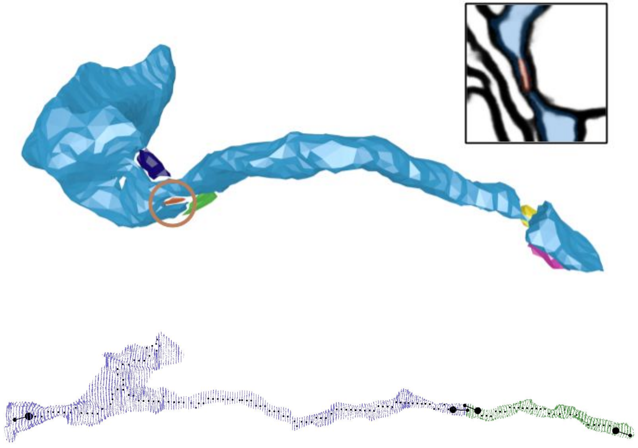

Hand-Designed
Geometric Constraints

Graph Construction with Biological Constraints

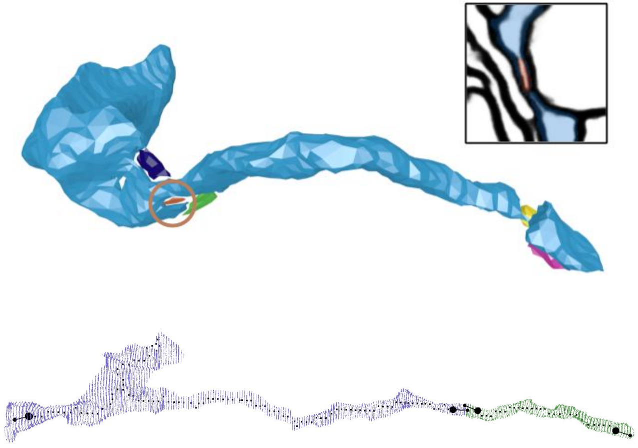

Hand-Designed
Geometric Constraints

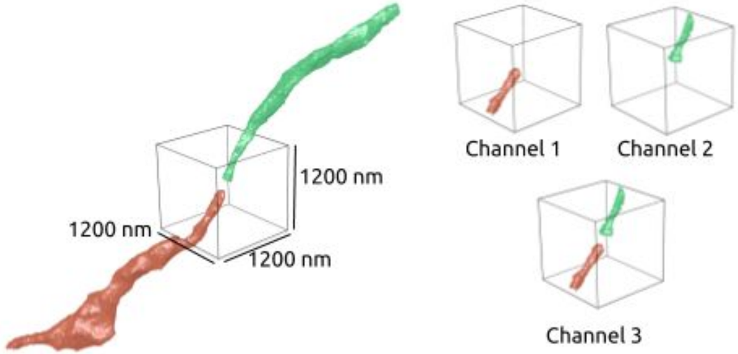

Machine-Learned
Morphologies

Node Generation

Existing segmentation strategies produce a large number of small segments

Node Generation

Existing segmentation strategies produce a large number of small segments

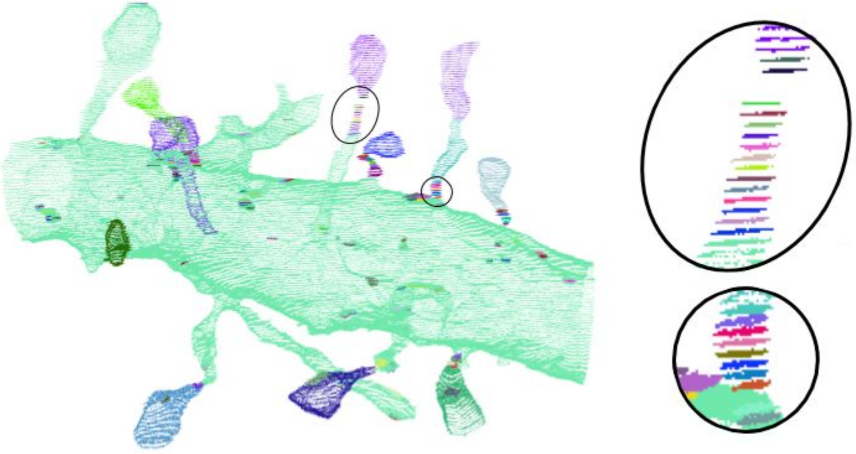

Singleton Slices

Node Generation

Existing segmentation strategies produce a large number of small segments

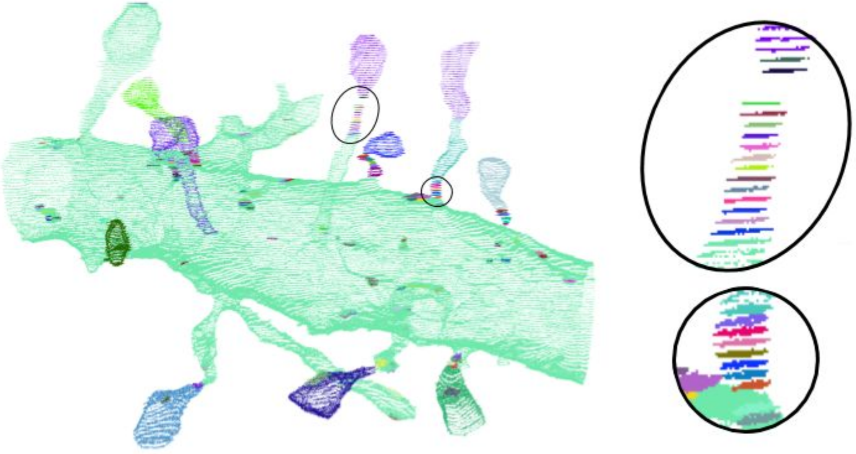

Singleton Slices

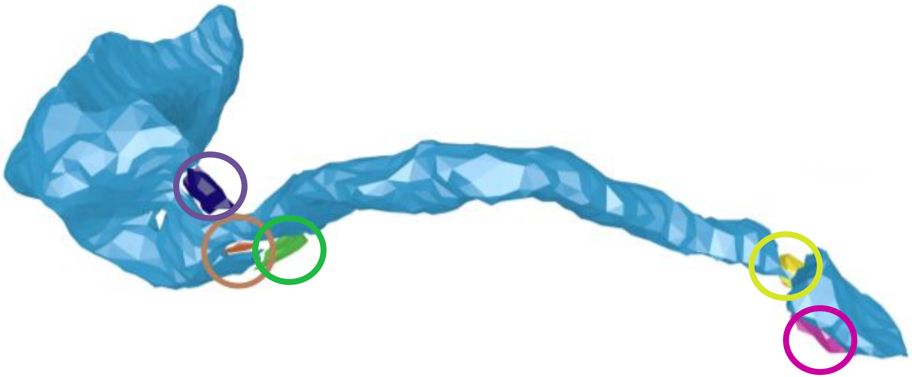

5 Small Segments

Singleton Removal

Merge adjacent singleton slices that have an Intersection-over-Union above 0.30

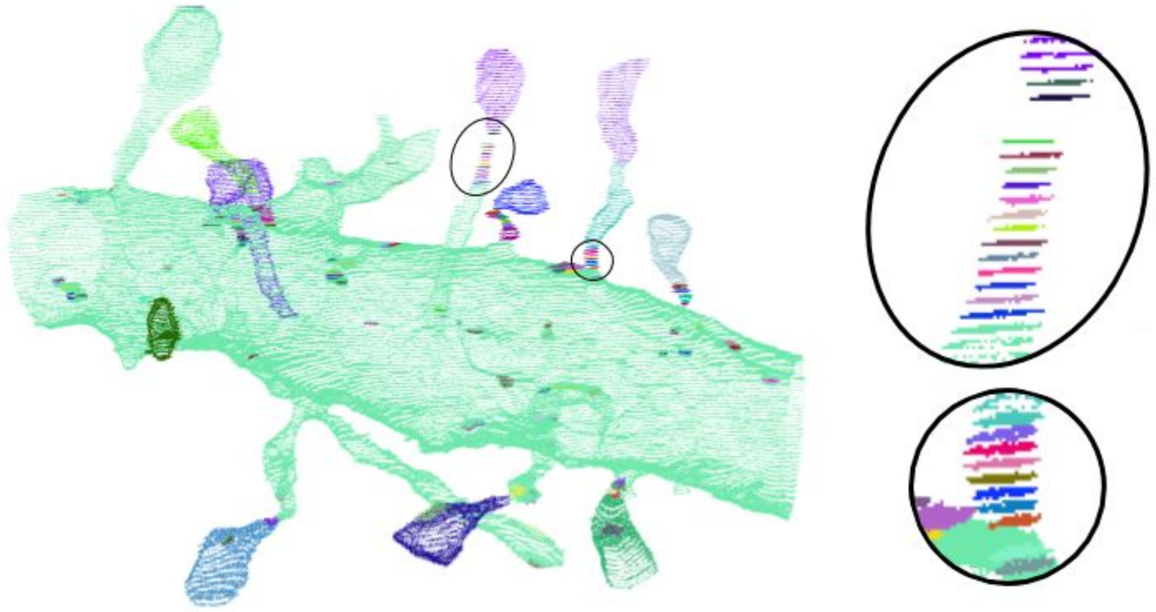

Merging Other Small Segments

Up to 80% of remaining segments are very small with little shape information

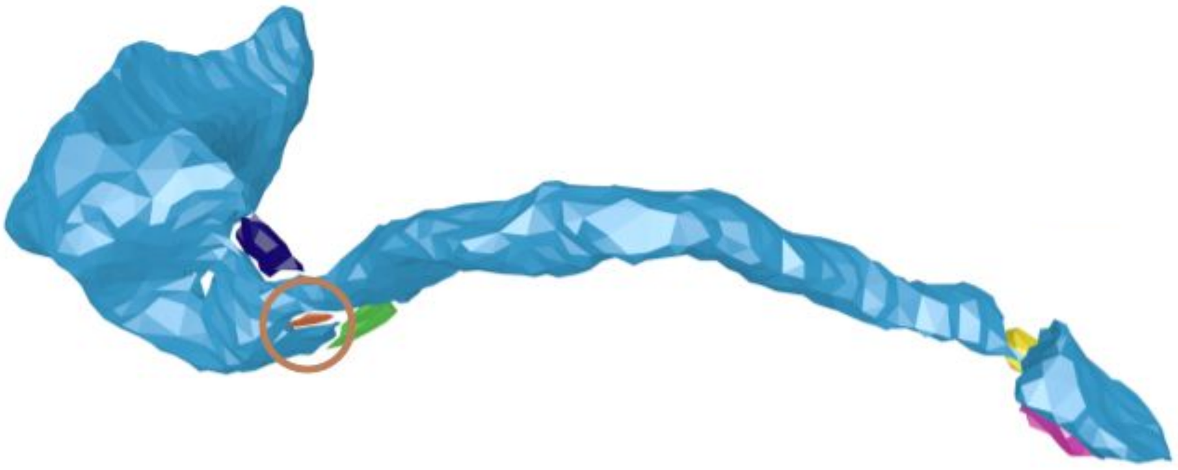

Merging Other Small Segments

Up to 80% of remaining segments are very small with little shape information

These small segments often occur at narrow locations with noisy affinities

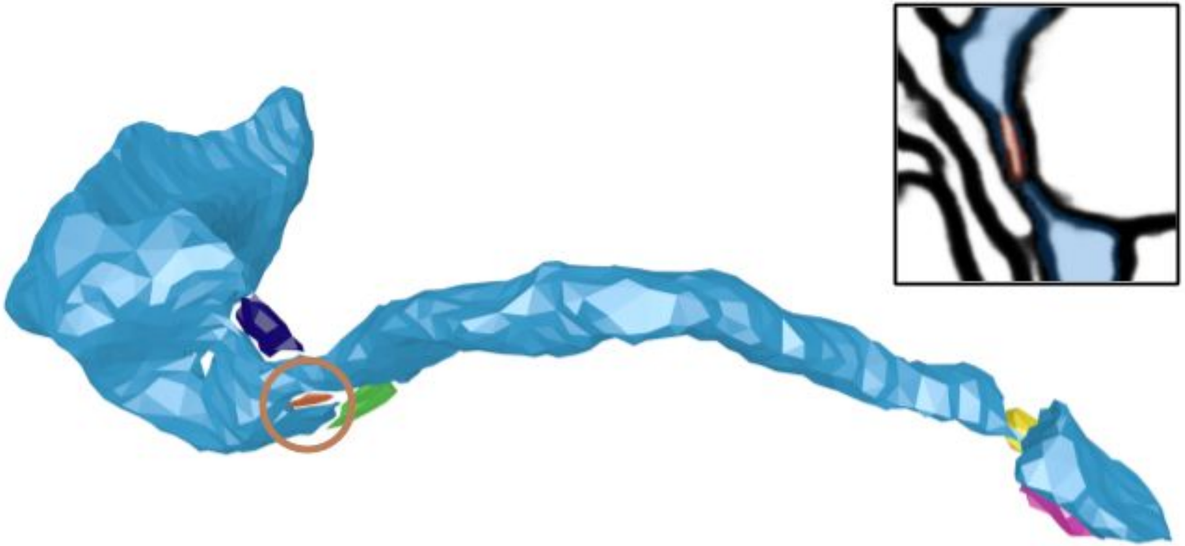

Small Segment Merging

Each small segment is merged with a nearby large segment

Small Segment Merging

Each small segment is merged with a nearby large segment

A 3D CNN predicts the most likely neighbor to belong to the same neuronal process

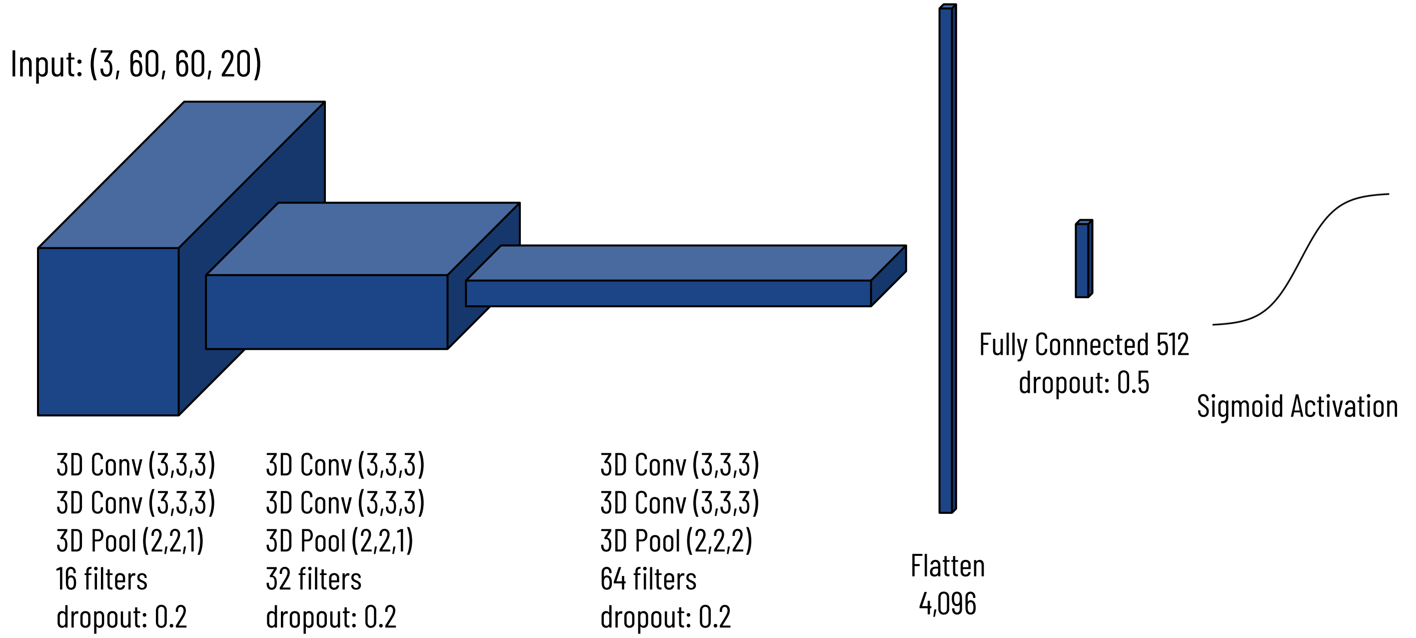

Edge Generation

Each segment has too many adjacent neighbors to use the adjacency matrix

Edge Generation

Each segment has too many adjacent neighbors to use the adjacency matrix

Typical Segment

Edge Generation

Each segment has too many adjacent neighbors to use the adjacency matrix

Typical Segment

103 Adjacent Neighbors

Handcrafted Geometric Constraints

Use directional information to identify potential split error locations

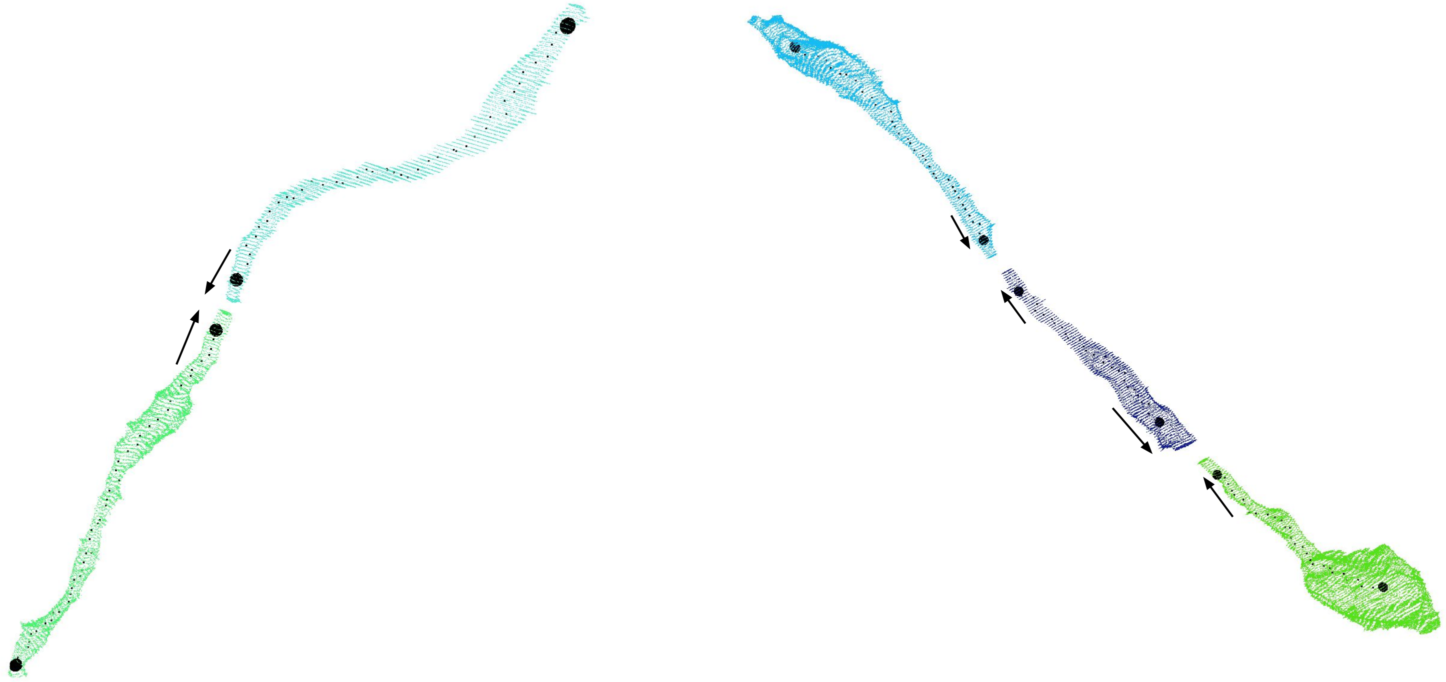

Skeleton Generation

Approximate volume shapes with 1D skeletons and identify potential split errors based on skeletal geometry around the endpoints

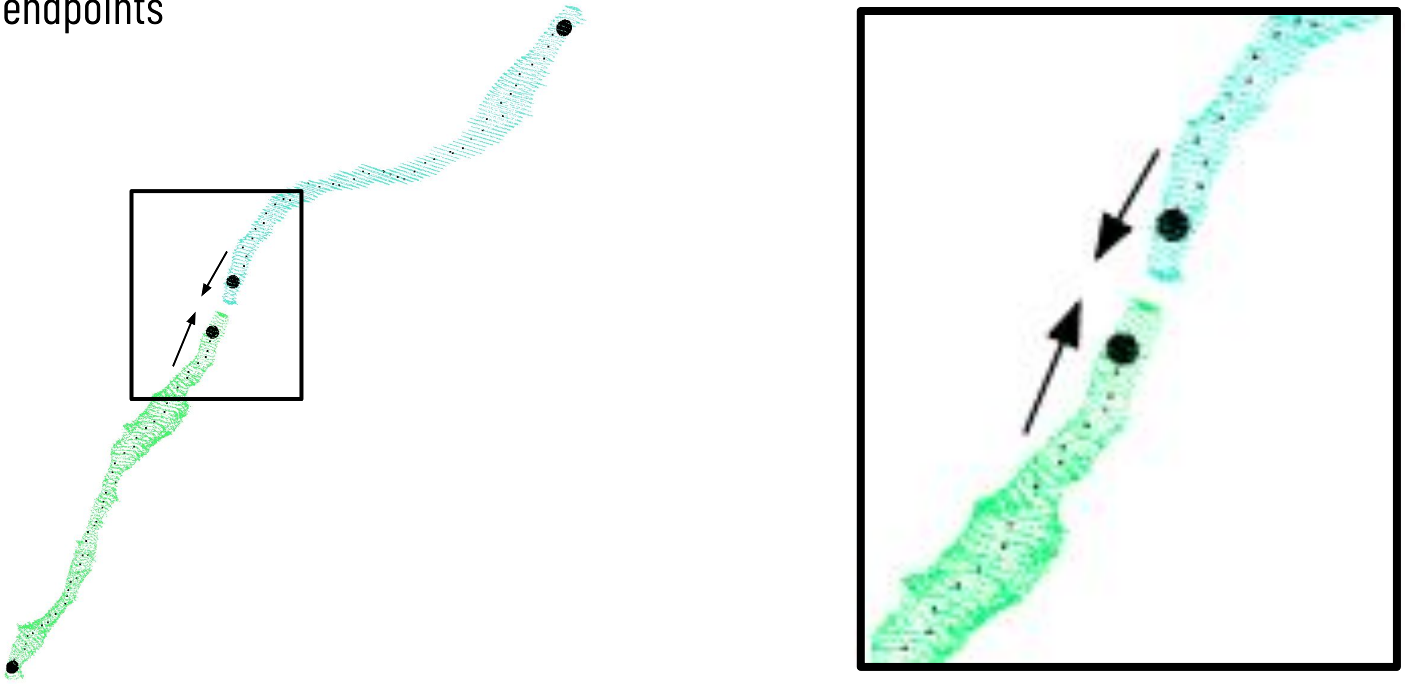

Skeleton Benchmark

Many skeleton generation strategies have been developed in the connectomics and volume processing communities

Skeleton Benchmark

Many skeleton generation strategies have been developed in the connectomics and volume processing communities

We create and publish a skeleton benchmark dataset to test existing and novel strategies on connectomics data

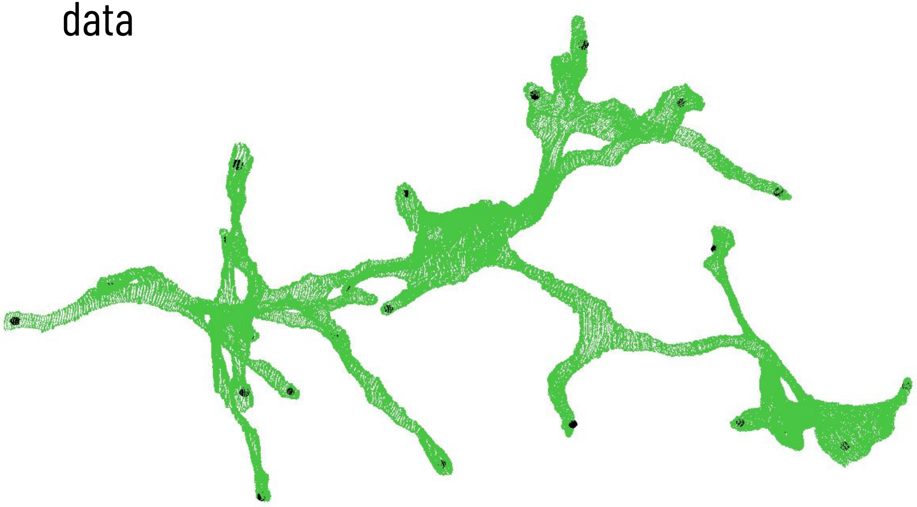

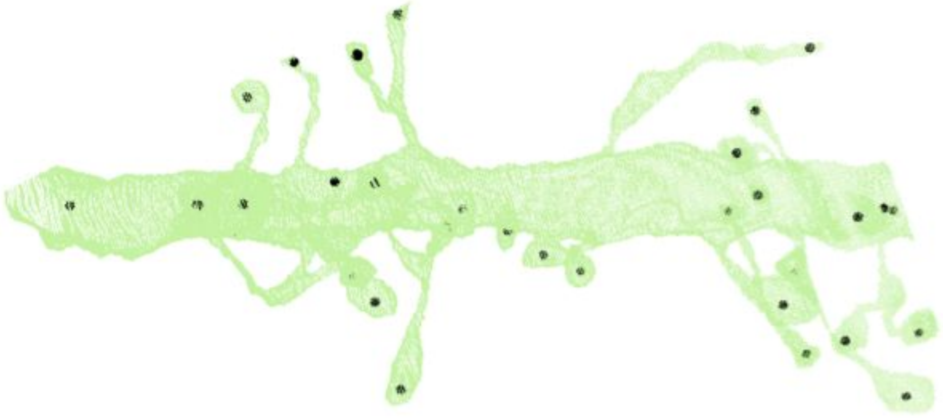

Skeleton Benchmark

Manually identified desired endpoint locations on 500 ground truth segments in the Kasthuri dataset

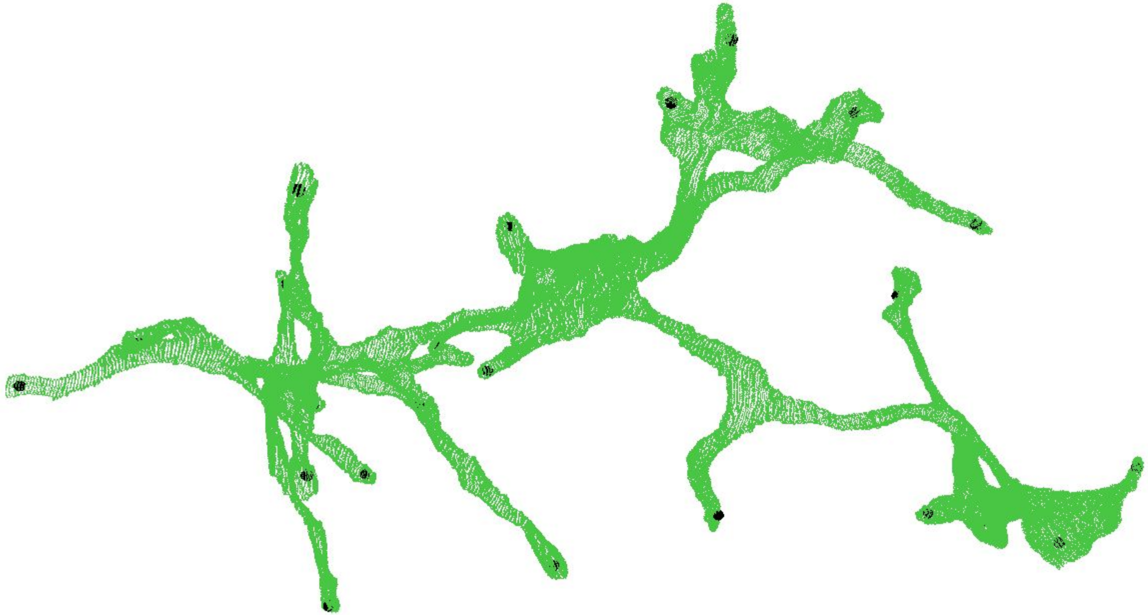

Skeleton Benchmark

Manually identified desired endpoint locations on 500 ground truth segments in the Kasthuri dataset

- Over 2000 endpoints manually located

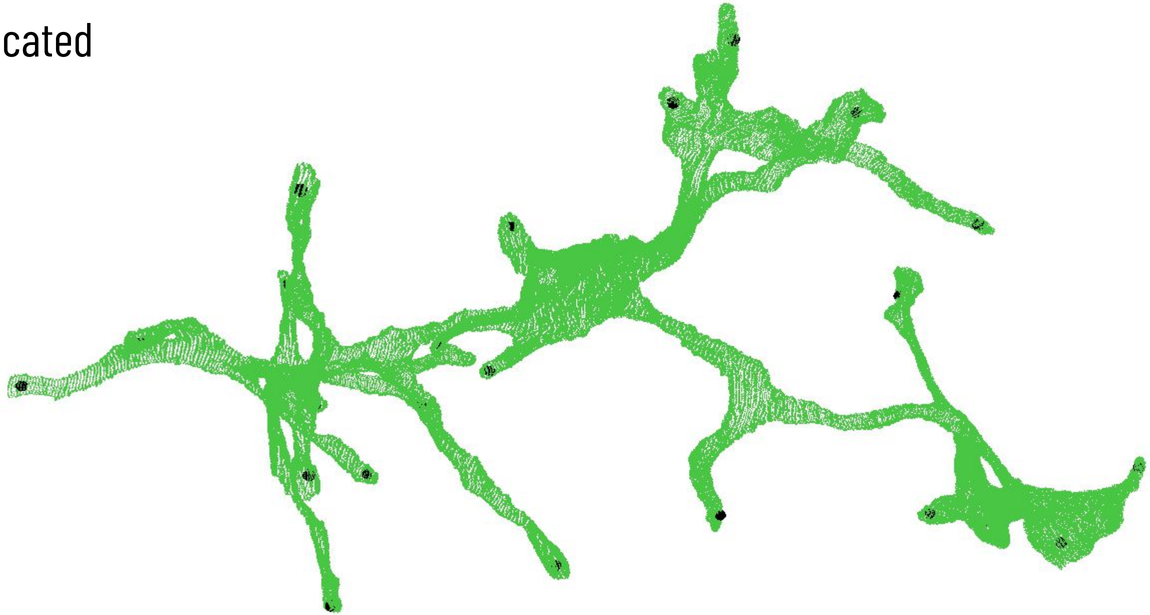

Skeleton Evaluation

We consider three different skeleton generation techniques: TEASER

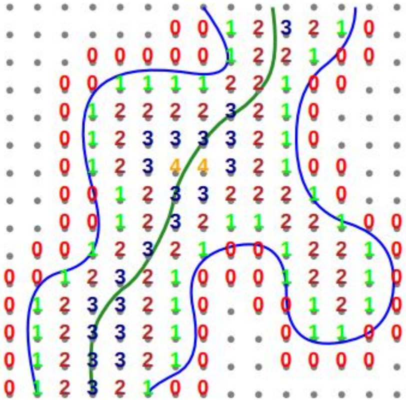

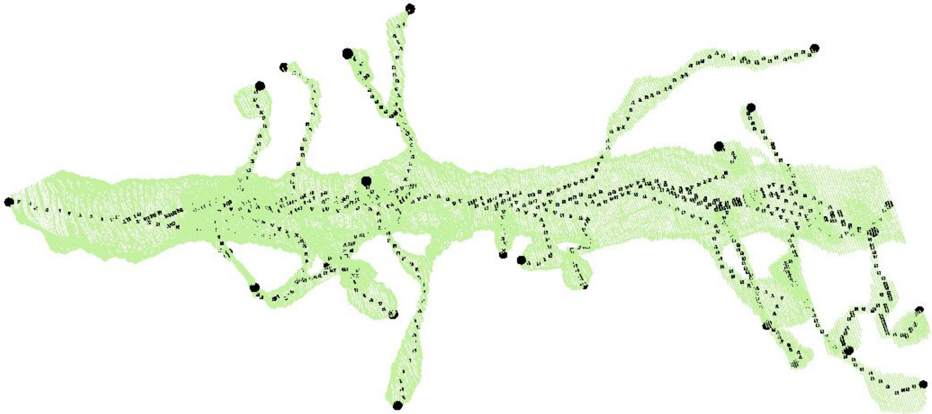

Sato et al.,TEASER: Tree-structure Extraction Algorithm for Accurate and Robust Skeletons,, Pacific Conference on Computer Graphics and Applications, 2000.

Skeleton Evaluation

We consider three different skeleton generation techniques: TEASER, Isthmus Thinning

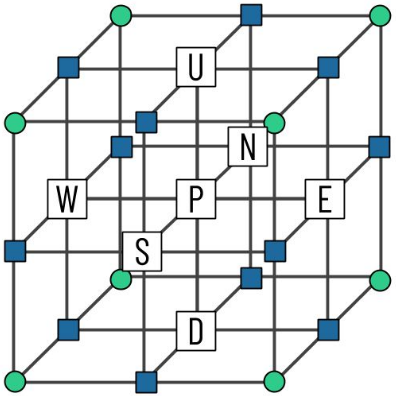

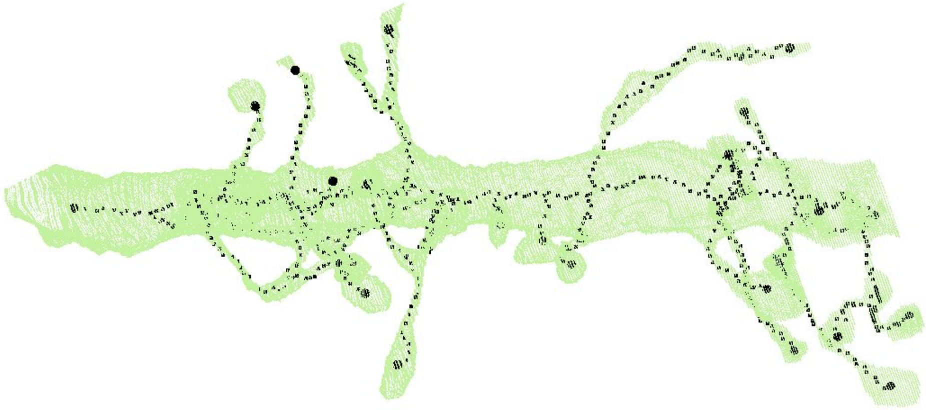

Palagyi, K., A Sequential 3D Curve-Thinning Algorithm Based on Isthmuses, International Symposium on Visual Computing, 2014.

Skeleton Evaluation

We consider three different skeleton generation techniques: TEASER, Isthmus Thinning, and Medial Surface Thinning

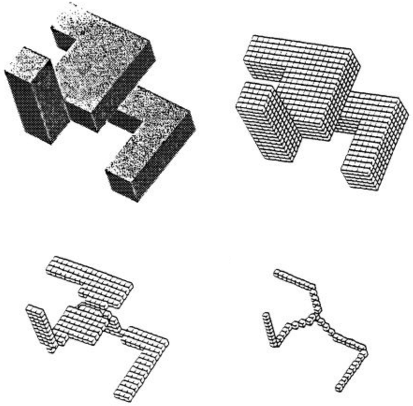

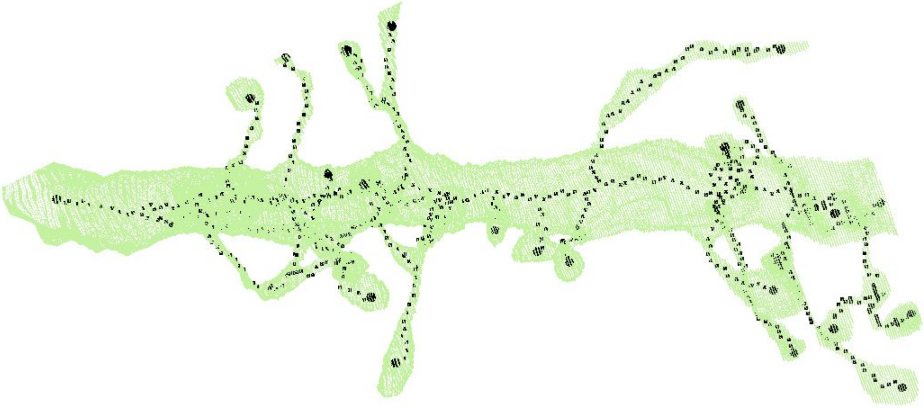

Lee et al., Building Skeleton Models via 3-D Medial Surface Axis Thinning Algorithms, CVGIP, 1994.

Skeleton Evaluation

We consider three different skeleton generation techniques: TEASER, Isthmus Thinning, and Medial Surface Thinning

We evaluated nearly 1000 different parameter configurations to find the best results on the benchmark

Skeleton Evaluation

We consider three different skeleton generation techniques: TEASER, Isthmus Thinning, and Medial Surface Thinning

We evaluated nearly 1000 different parameter configurations to find the best results on the benchmark

A positive result is a generated endpoint within 800 nanometers of our manually labeled endpoint

Skeleton Evaluation

Isthmus thinning outperformed the existing methods in both endpoint accuracy and speed

Precision: 94.7%
Recall: 86.7%
F-Score: 90.5%

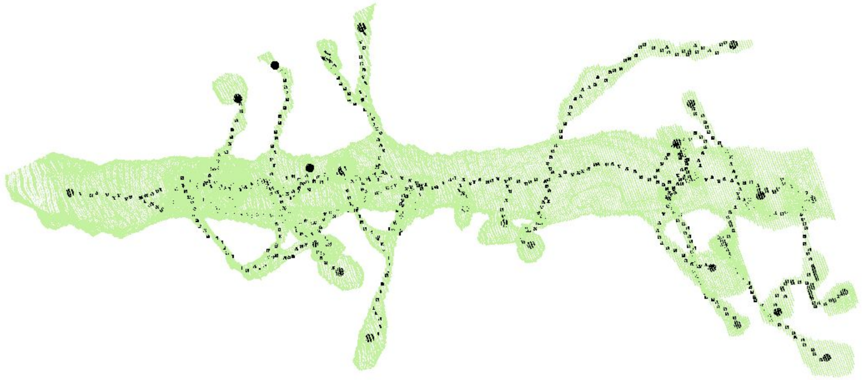

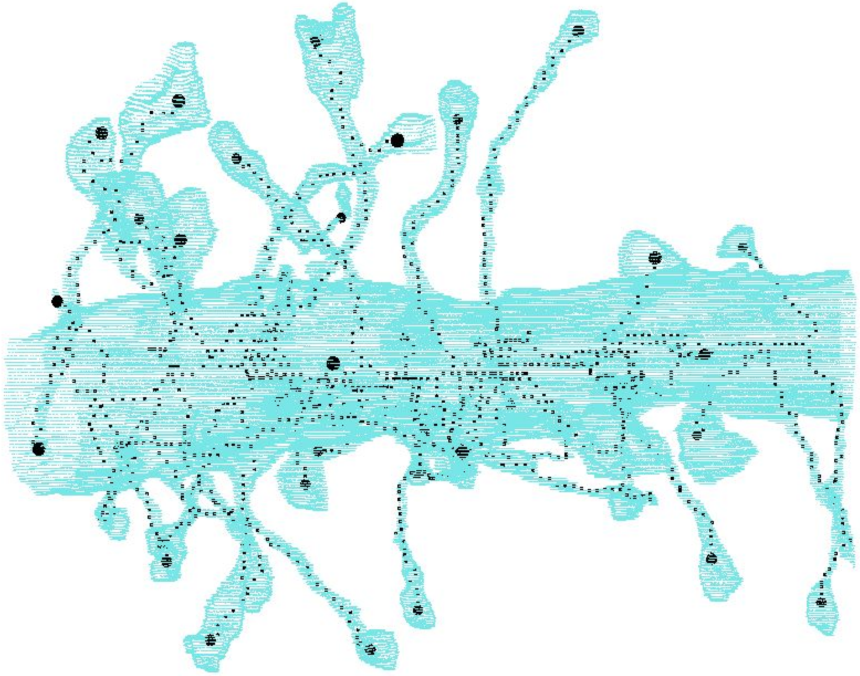

Edge Generation

Two nodes receive an edge in the graph if one of the corresponding skeletons has an endpoint vector towards the other segment

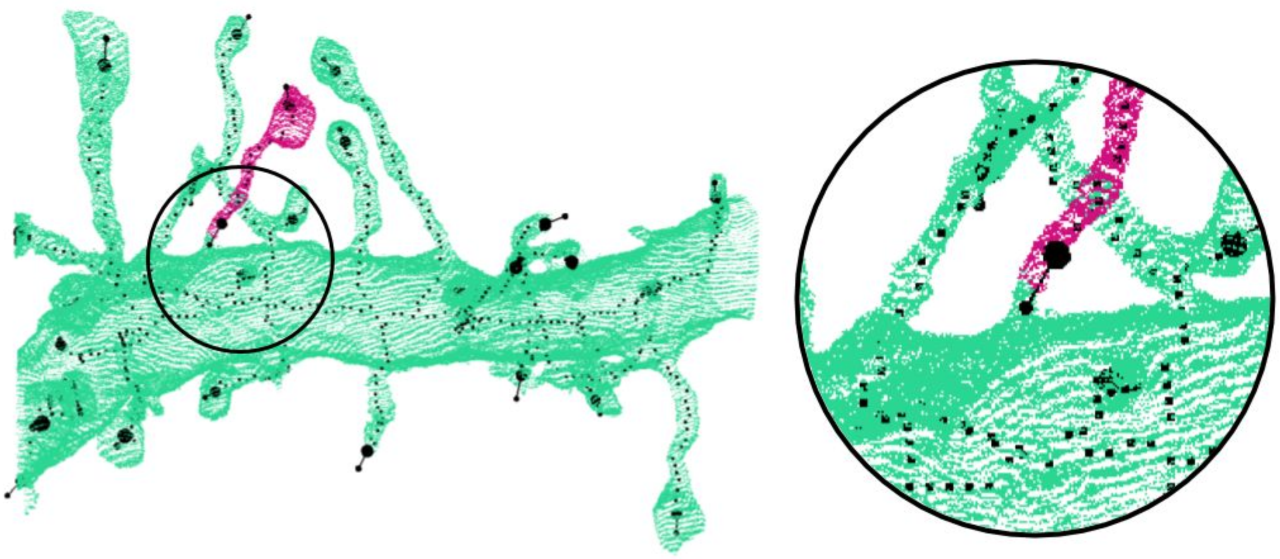

Edge Generation

Endpoint must be within 500 nanometers of the other segment

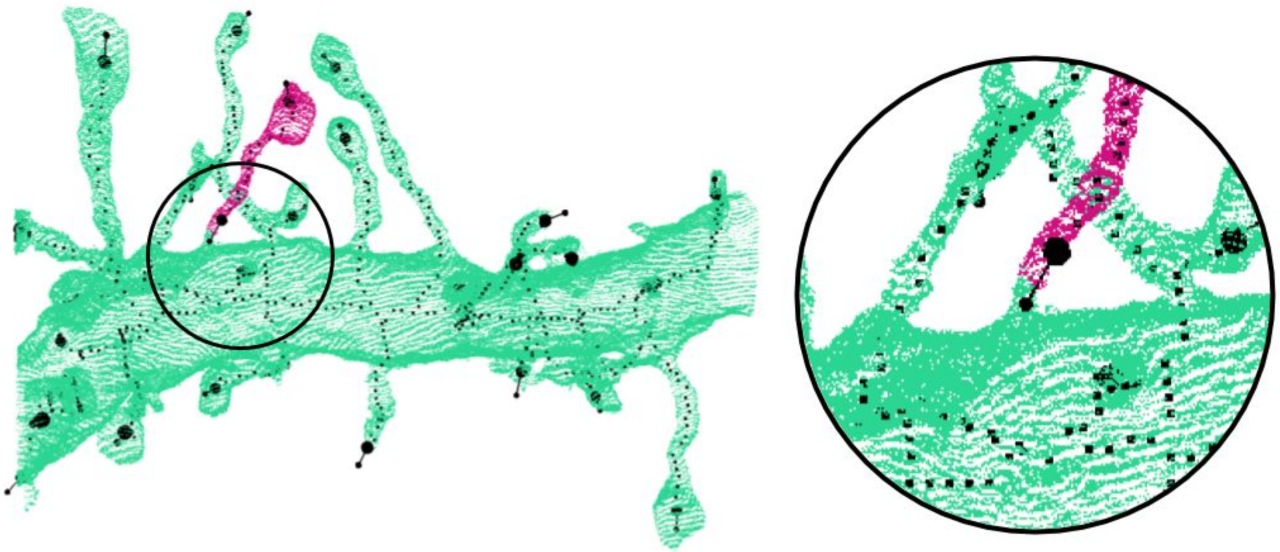

Machine-learned Morphologies

We train a VGG-style convolutional neural network to predict if two segments belong to the same neuronal process

Architecture and Training Parameters

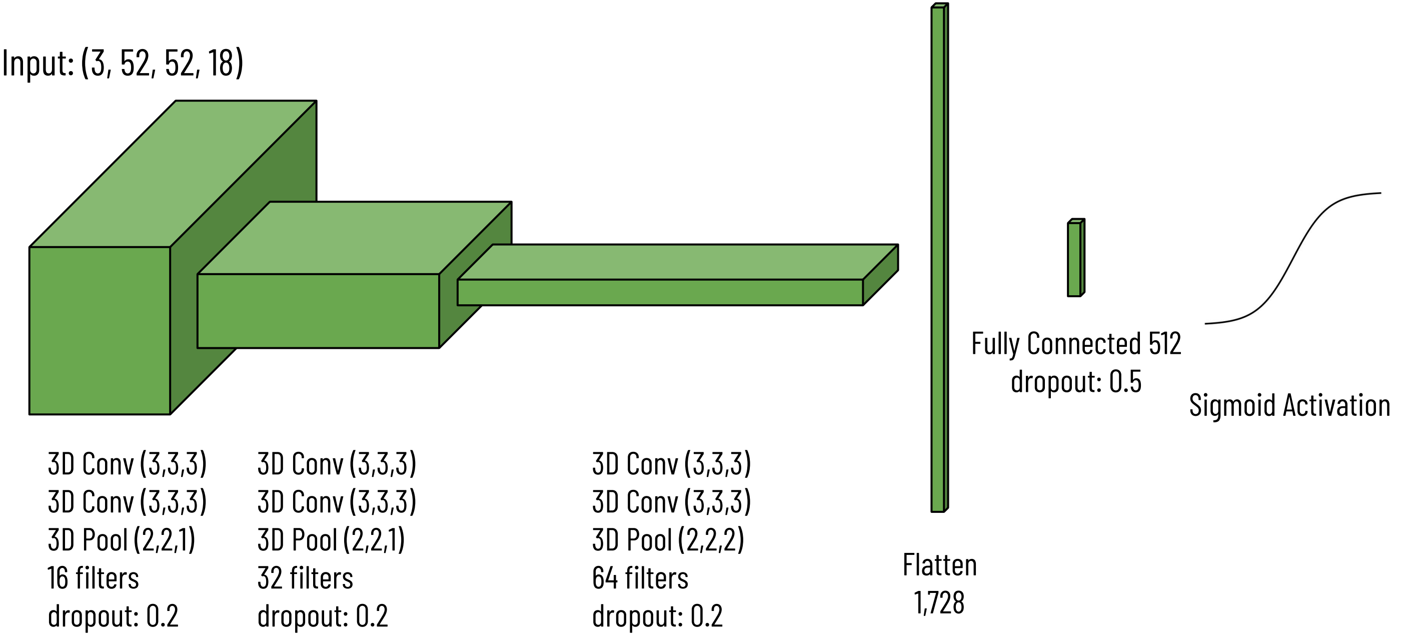

Architecture and Training Parameters

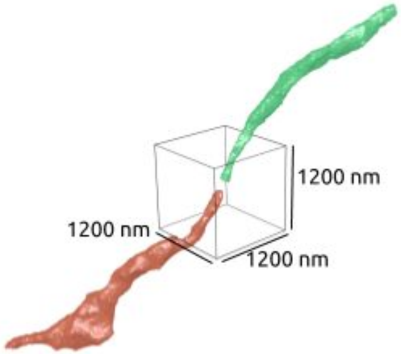

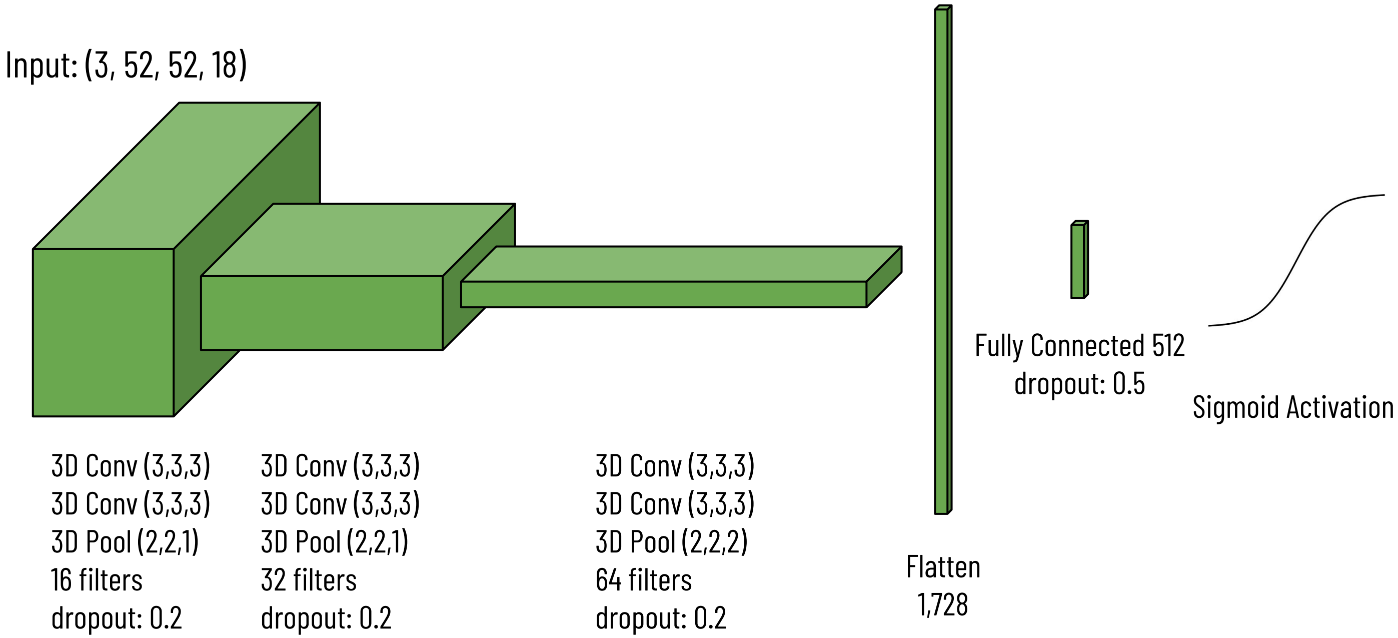

Architecture and Training Parameters

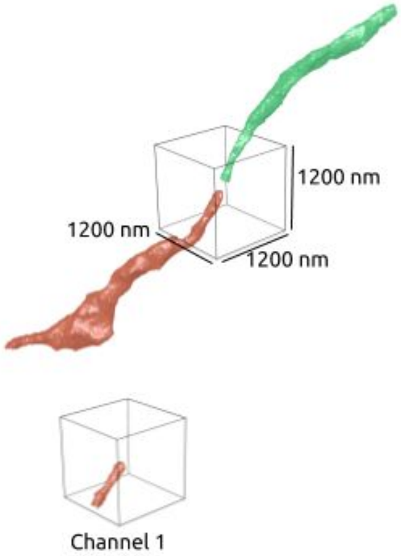

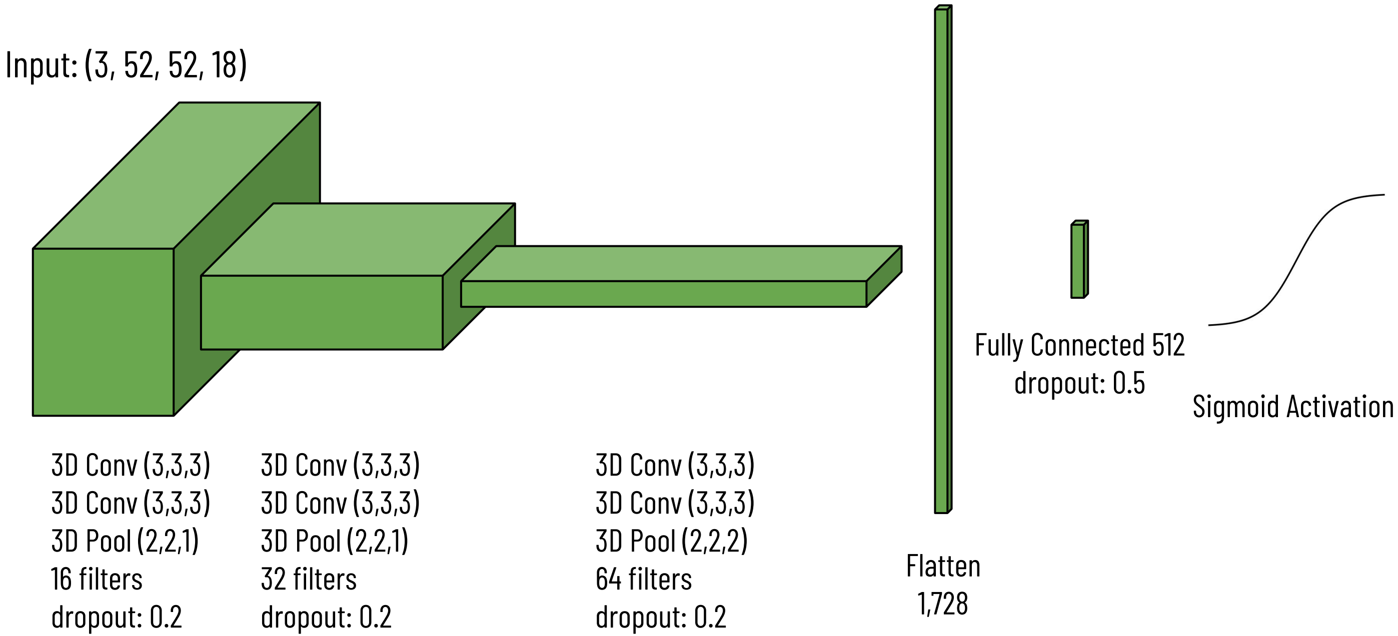

Architecture and Training Parameters

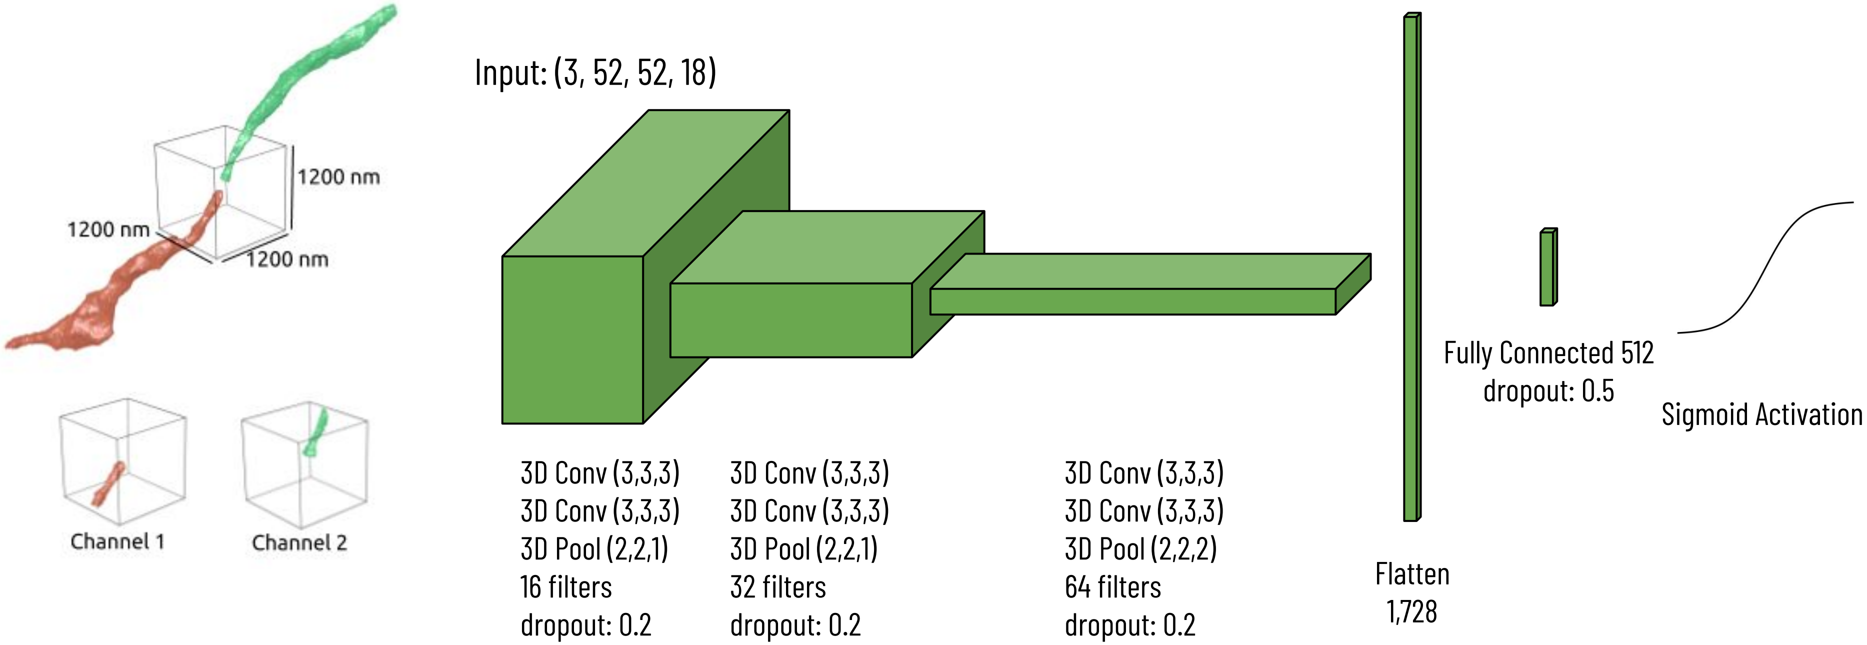

Architecture and Training Parameters

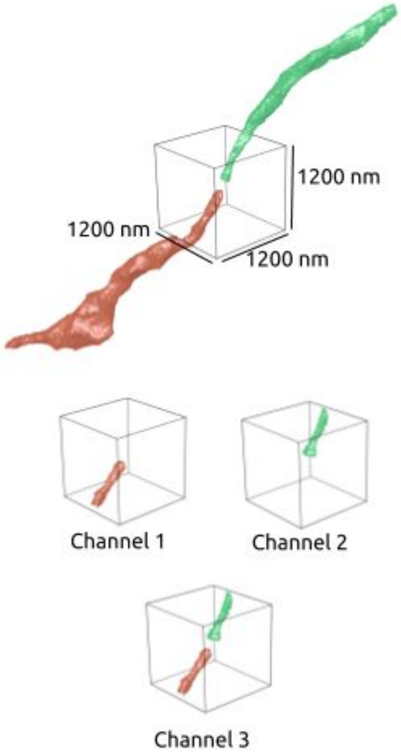

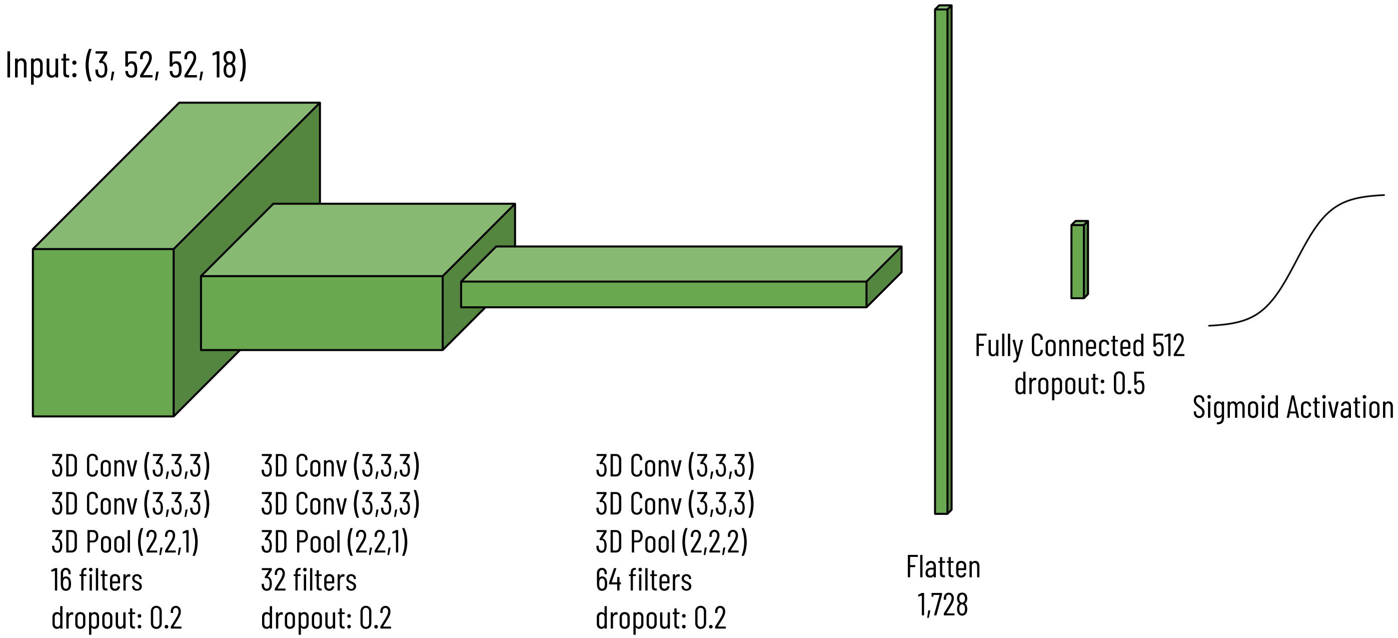

Regions of Interest

Too small and there is not enough local context

Regions of Interest

Too large and unnecessary detail inhibits learning

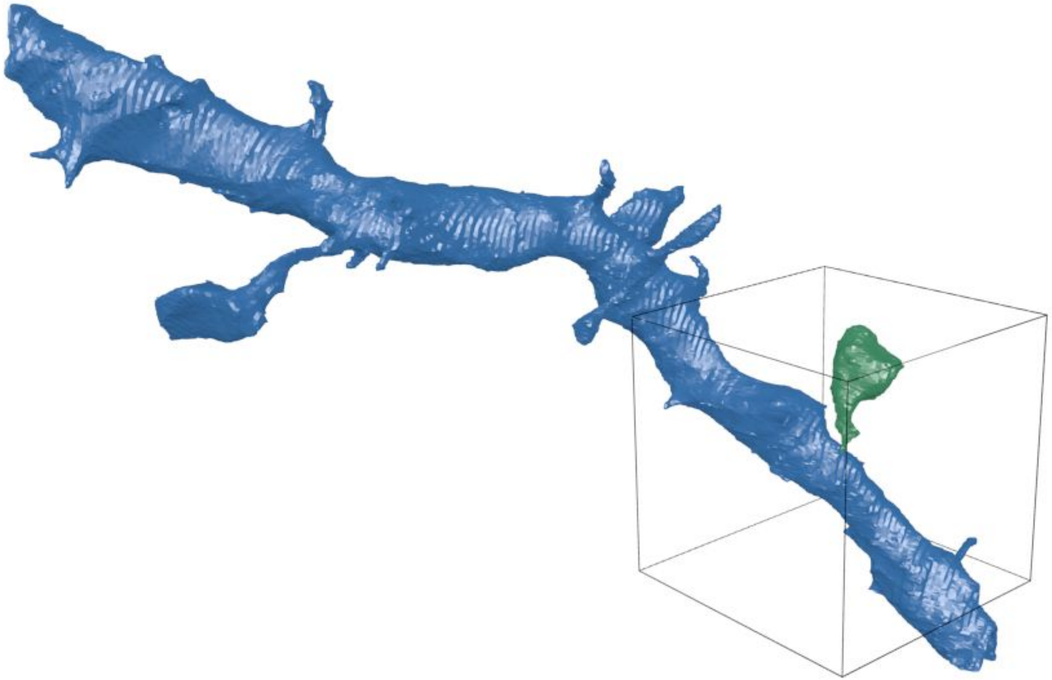

Regions of Interest

Found that cubes of size 1200 x 1200 x 1200 nm^3 work well

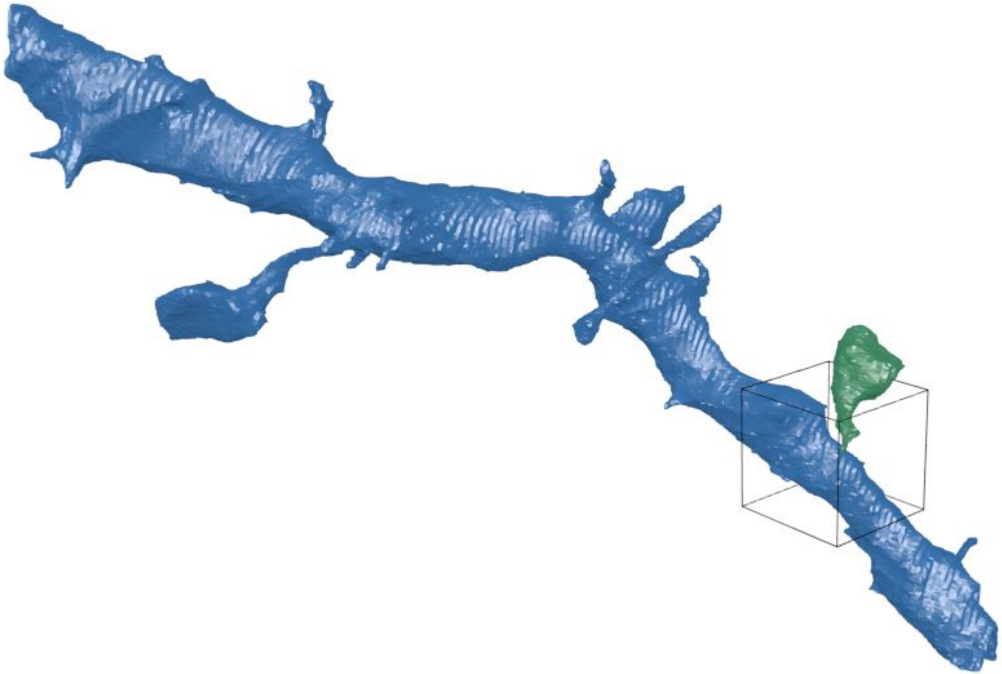

Input Examples

Should Merge

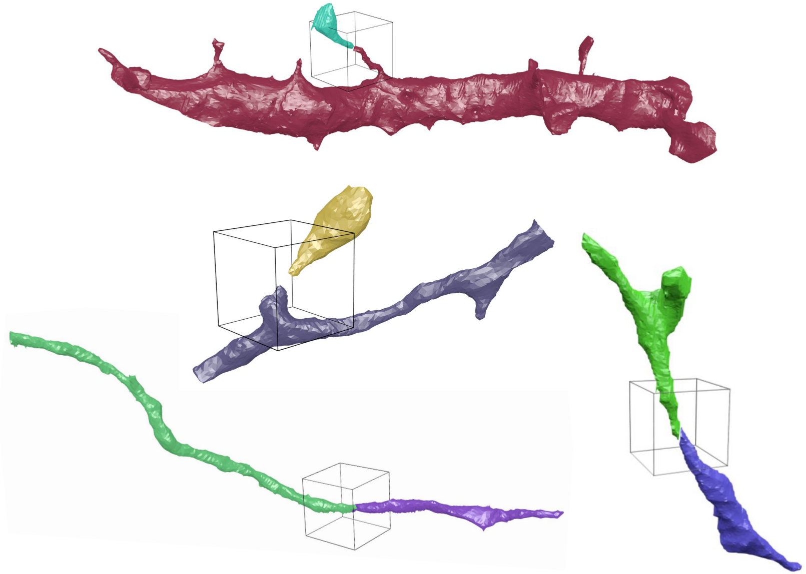

Input Examples

Should Merge

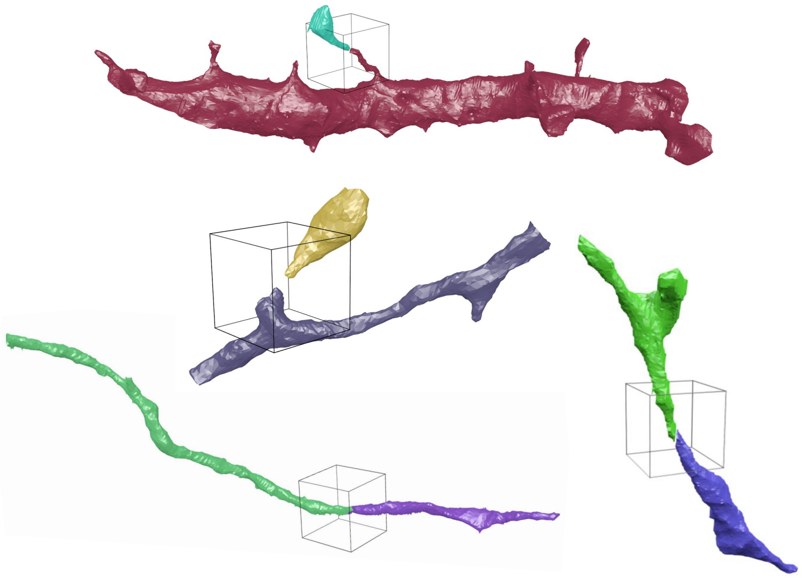

Should Not Merge

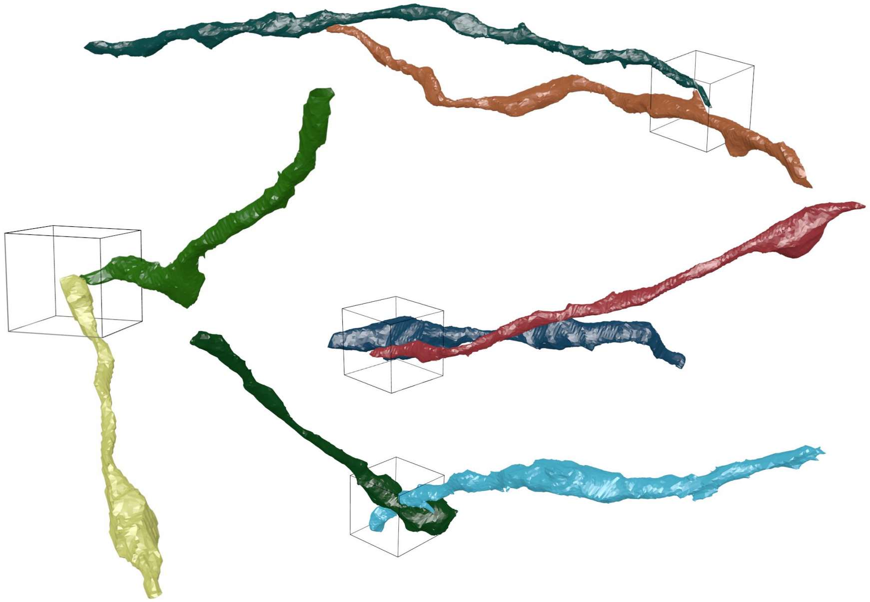

Goal: Partition graph into neuronal processes

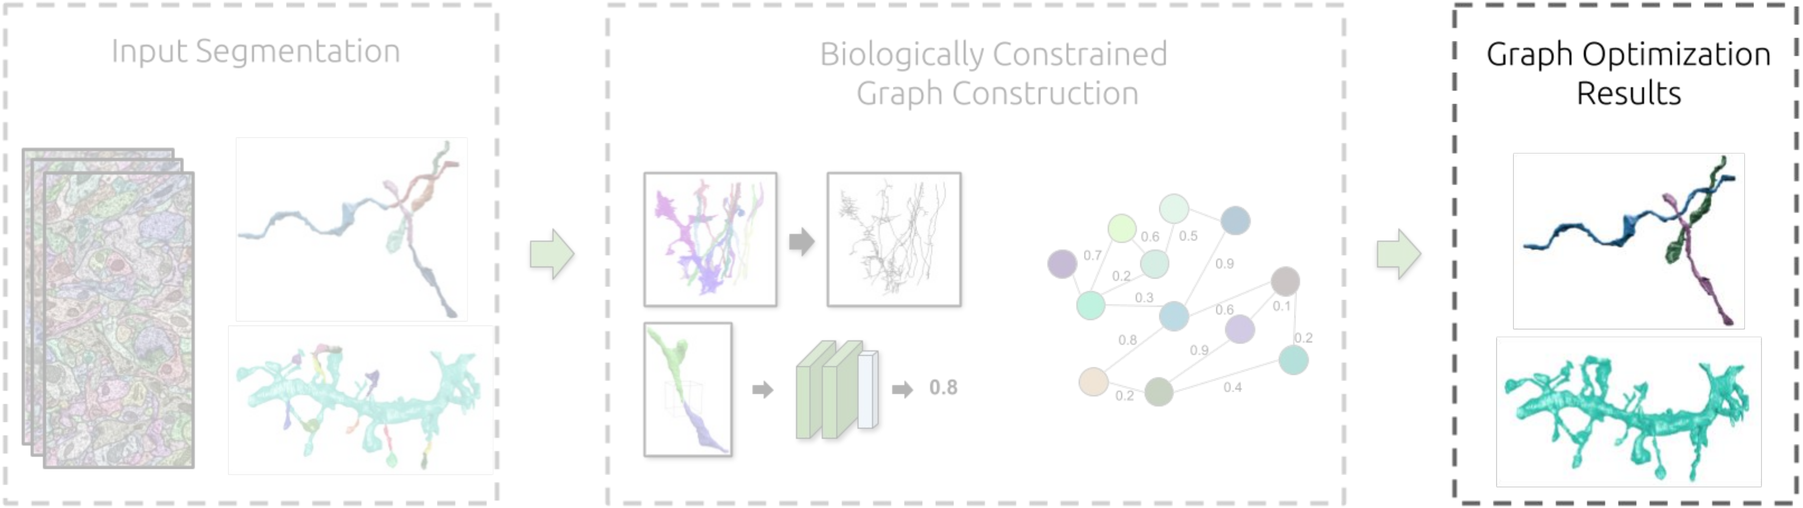

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

- The final number of segments is not predetermined

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

- The final number of segments is not predetermined
- Guarantees a globally consistent solution

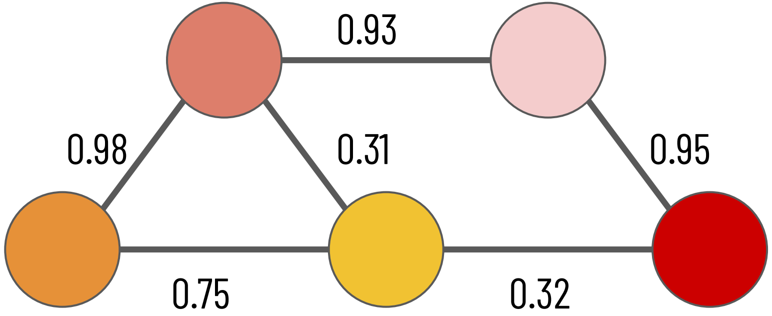

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

- The final number of segments is not predetermined
- Guarantees a globally consistent solution

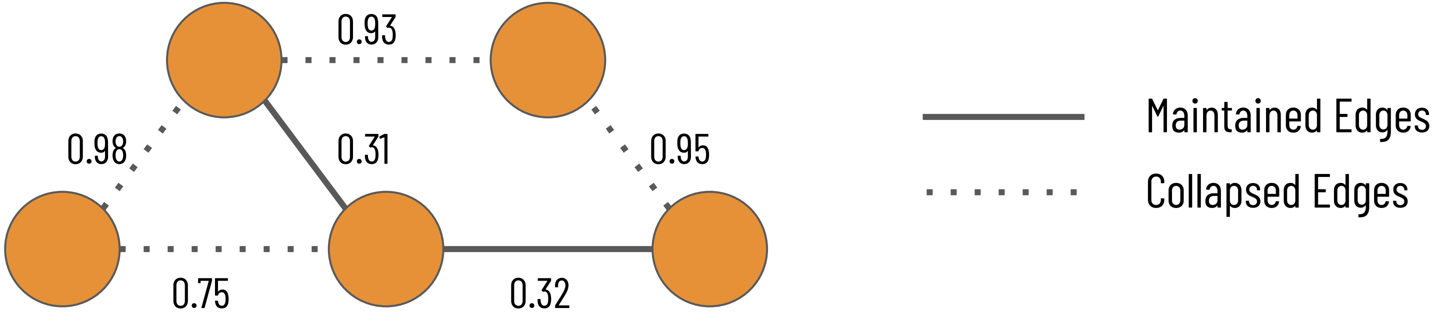

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

The final number of segments is not predetermined

Guarantees a globally consistent solution

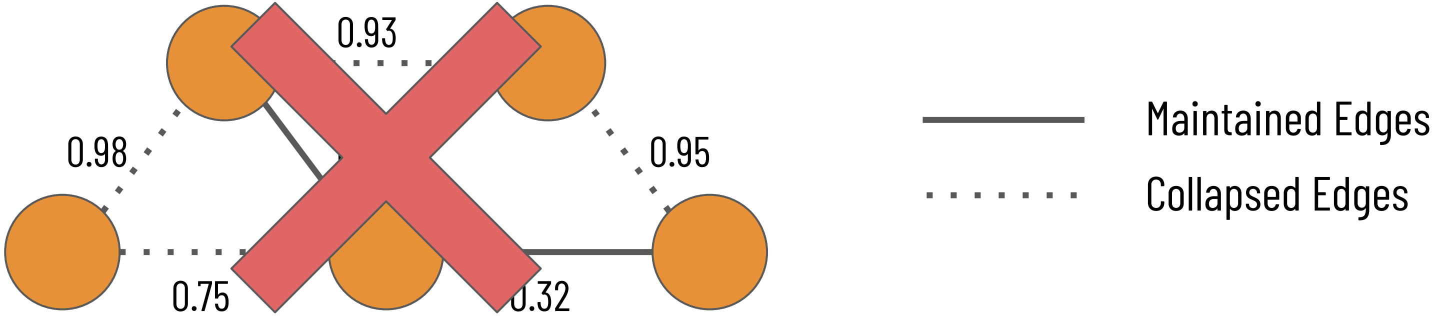

Lifted Edges

Represent long range probabilities of non-adjacent nodes belonging to the same neuron

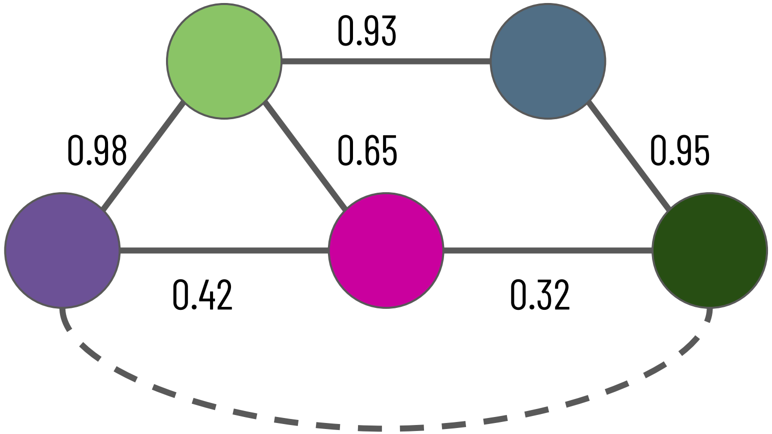

Lifted Edges

Represent long range probabilities of non-adjacent nodes belonging to the same neuron

Take value of highest probability path between two nodes

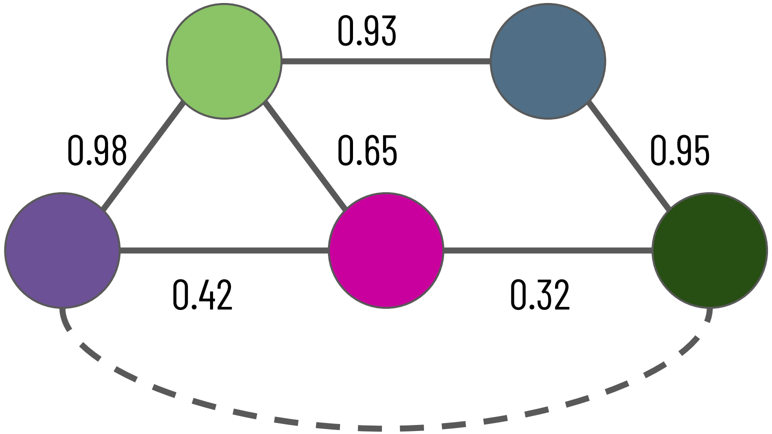

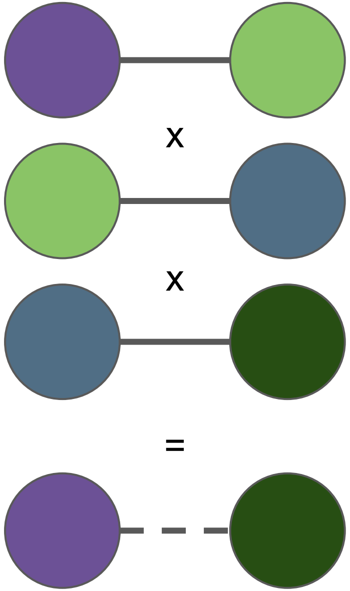

Lifted Edges

Represent long range probabilities of non-adjacent nodes belonging to the same neuron

Take value of highest probability path between two nodes

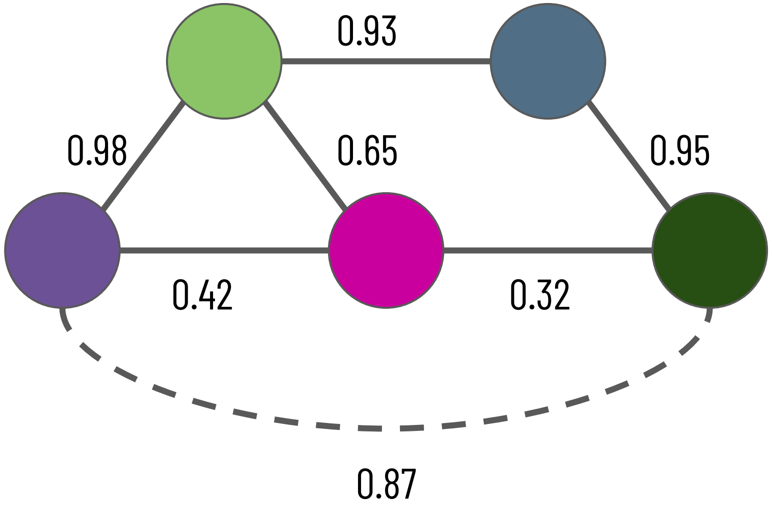

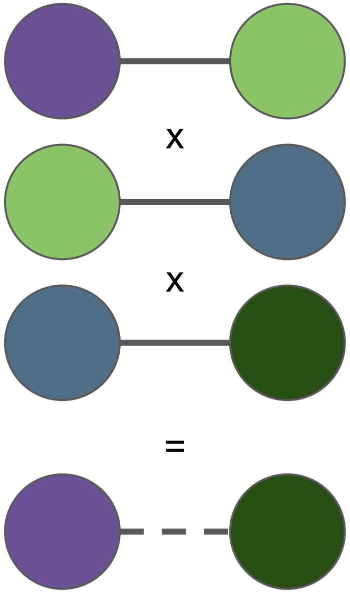

Datasets

Kasthuri

Princeton Neuroscience Institute

SNEMI3D

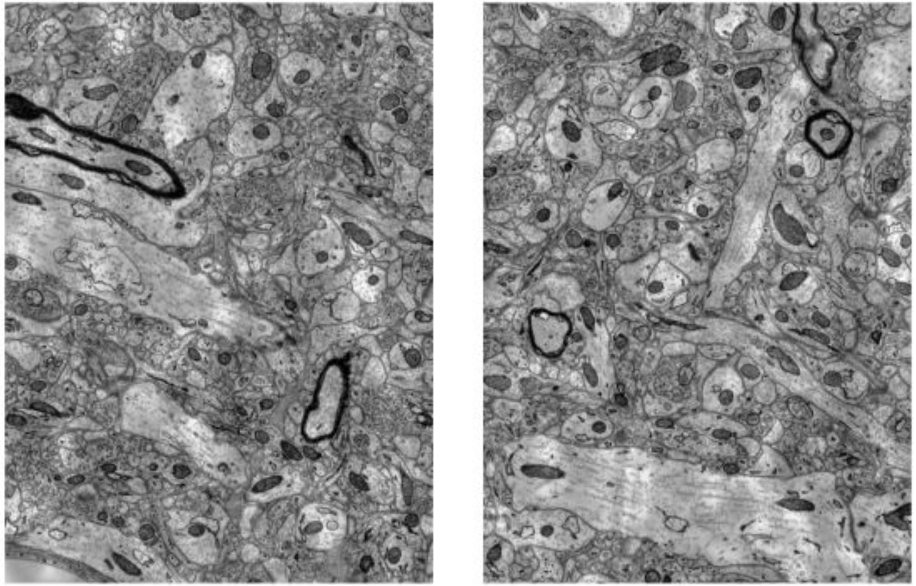

2 Volumes

6 x 6 x 30 nm^3 / vx

1335 x 1809 x 338 vx

8.01 x 10.85 x 10.14 μm^3

Datasets

Kasthuri

Princeton Neuroscience Institute

SNEMI3D

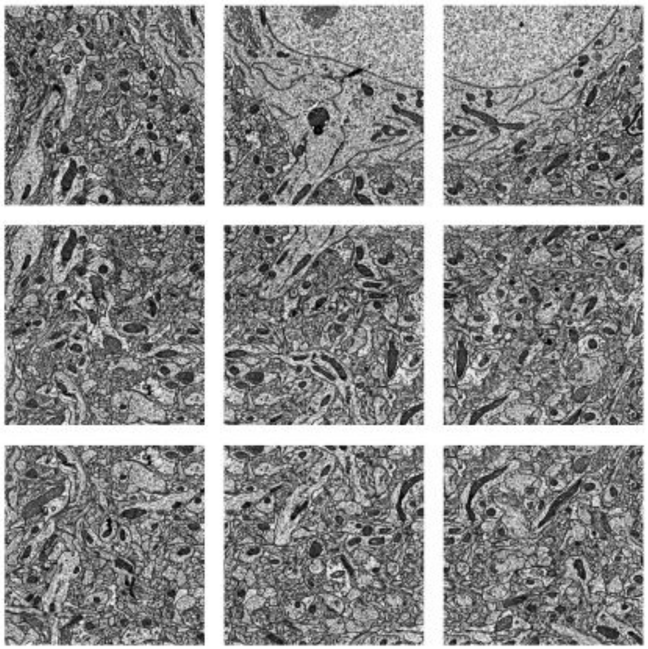

9 Volumes

$3.6 \times 3.6 \times 40$ nm^3 / vx

$2048 \times 2048 \times 256$ vx

$7.37 \times 7.37 \times 10.24$ μm^3

Datasets

Kasthuri

Princeton Neuroscience Institute

SNEMI3D

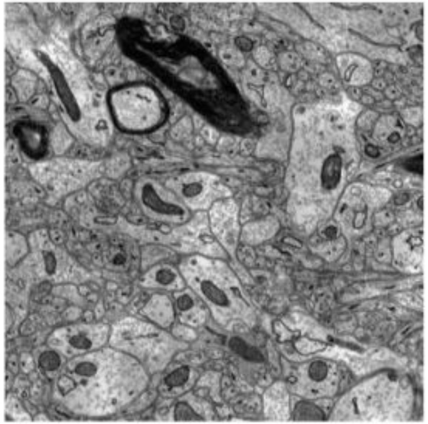

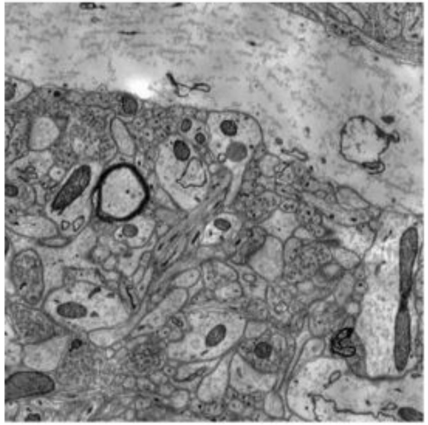

2 Volumes

$3 \times 3 \times 30\ nm^3$ / vx

1024 x 1024 x 100 vx

$3.07 \times 3.07 \times 3\ \mu m^3$

Input Segmentations

For the two PNI Test datasets, we use zwatershed and mean agglomeration

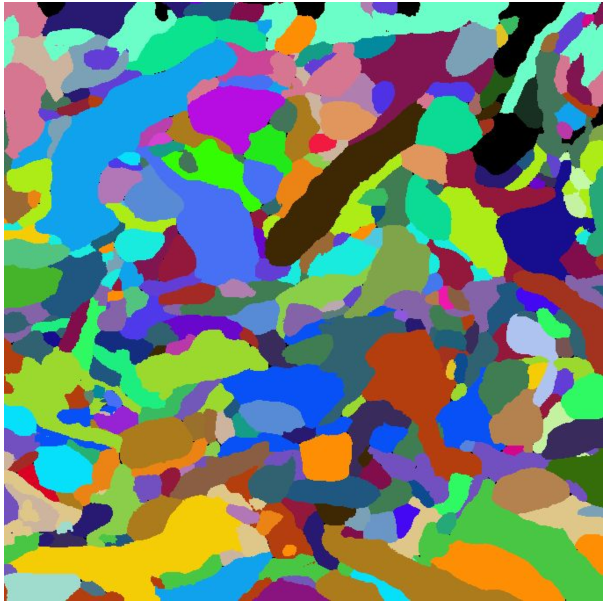

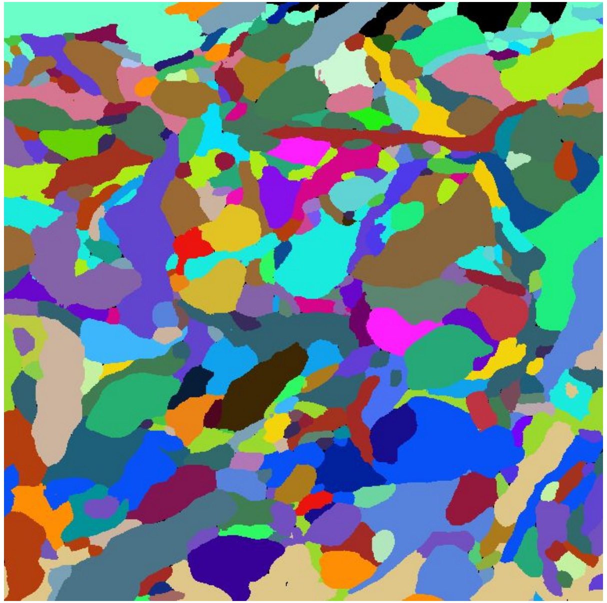

Input Segmentations

For the Kasthuri and SNEMI3D datasets, we use the waterz agglomeration strategy

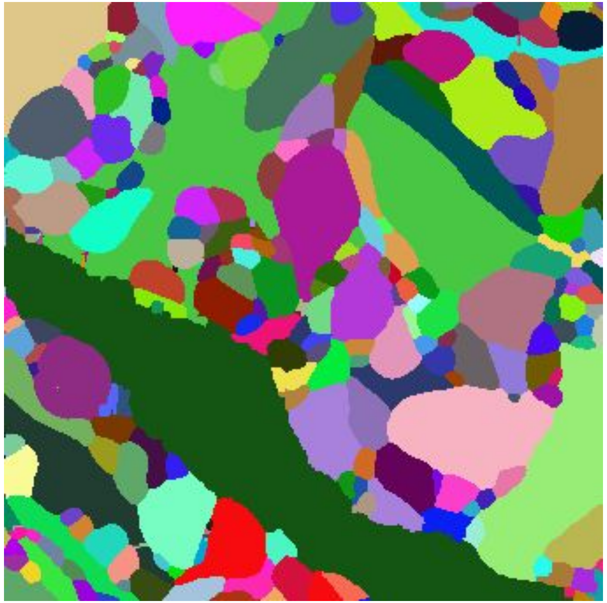

Split Variation of Information

Measure of entropy between segmentation and ground truth

Split Variation of Information

Measure of entropy between segmentation and ground truth

VI Split: Increases if two voxels from the same neuron have different labels

Split Variation of Information

Measure of entropy between segmentation and ground truth

VI Split: Increases if two voxels from the same neuron have different labels

VI Merge: Increases if two voxels from different neurons have the same label

Split Variation of Information

Measure of entropy between segmentation and ground truth

VI Split: Increases if two voxels from the same neuron have different labels

VI Merge: Increases if two voxels from different neurons have the same label

Total Variation of Information = VI Split + VI Merge

Variation of Information

Dataset	Baseline	Proposed	Decrease
PNI Test One	0.491	0.388	-20.9%
PNI Test Two	0.416	0.297	-28.7%
Kasthuri Test	0.965	0.815	-15.6%
SNEMI3D	0.807	0.647	-19.8%

Qualitative Results

Qualitative Results

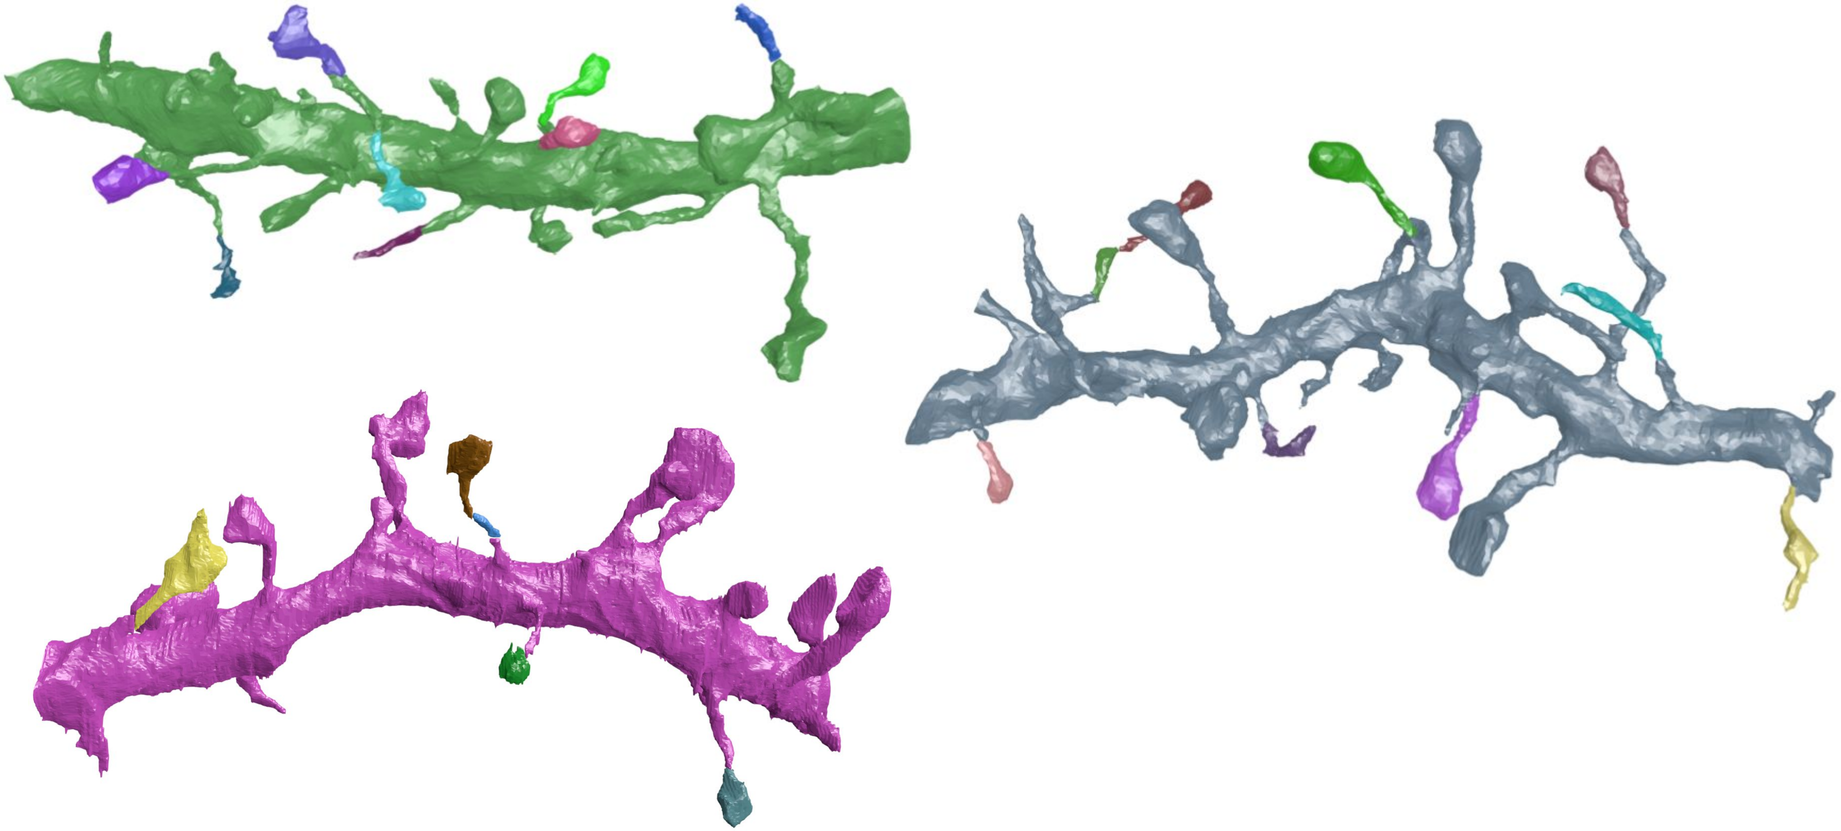

Failure Cases

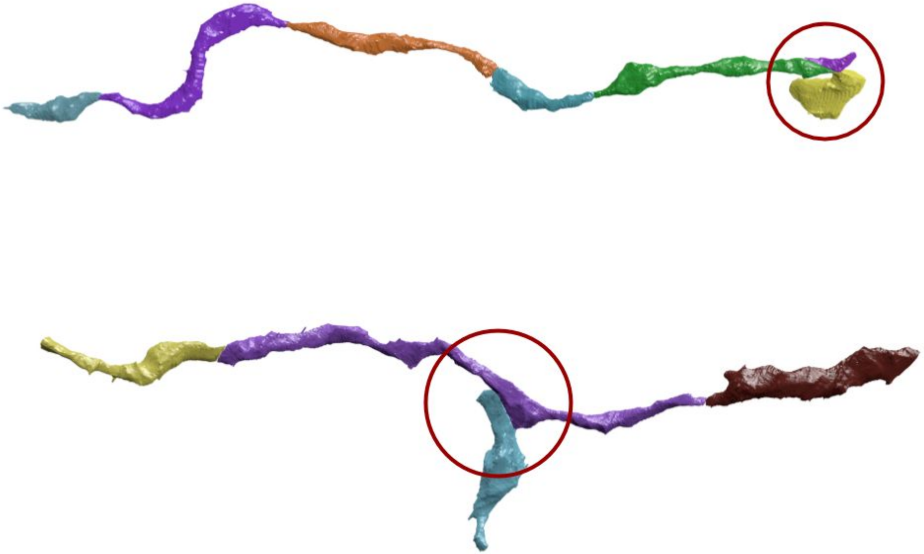

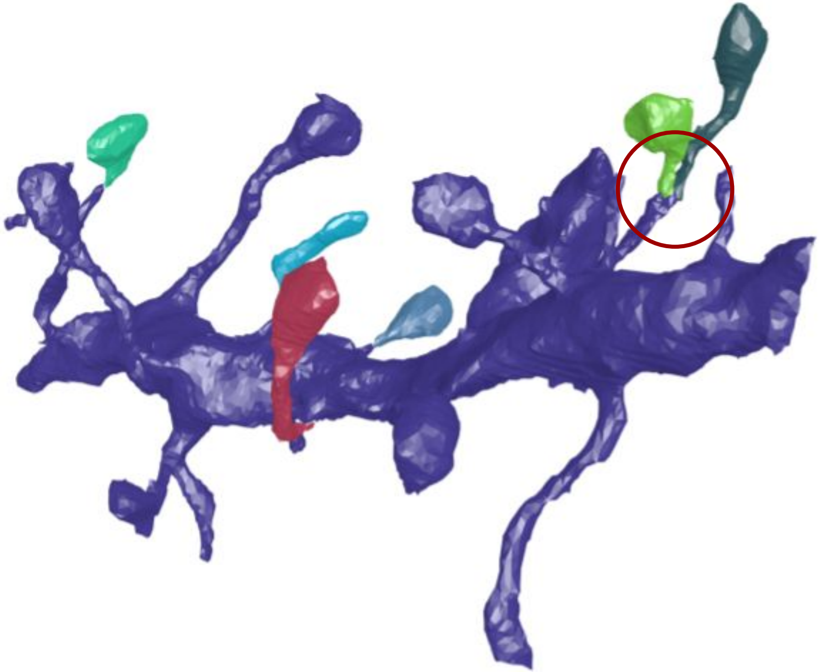

Failure Cases

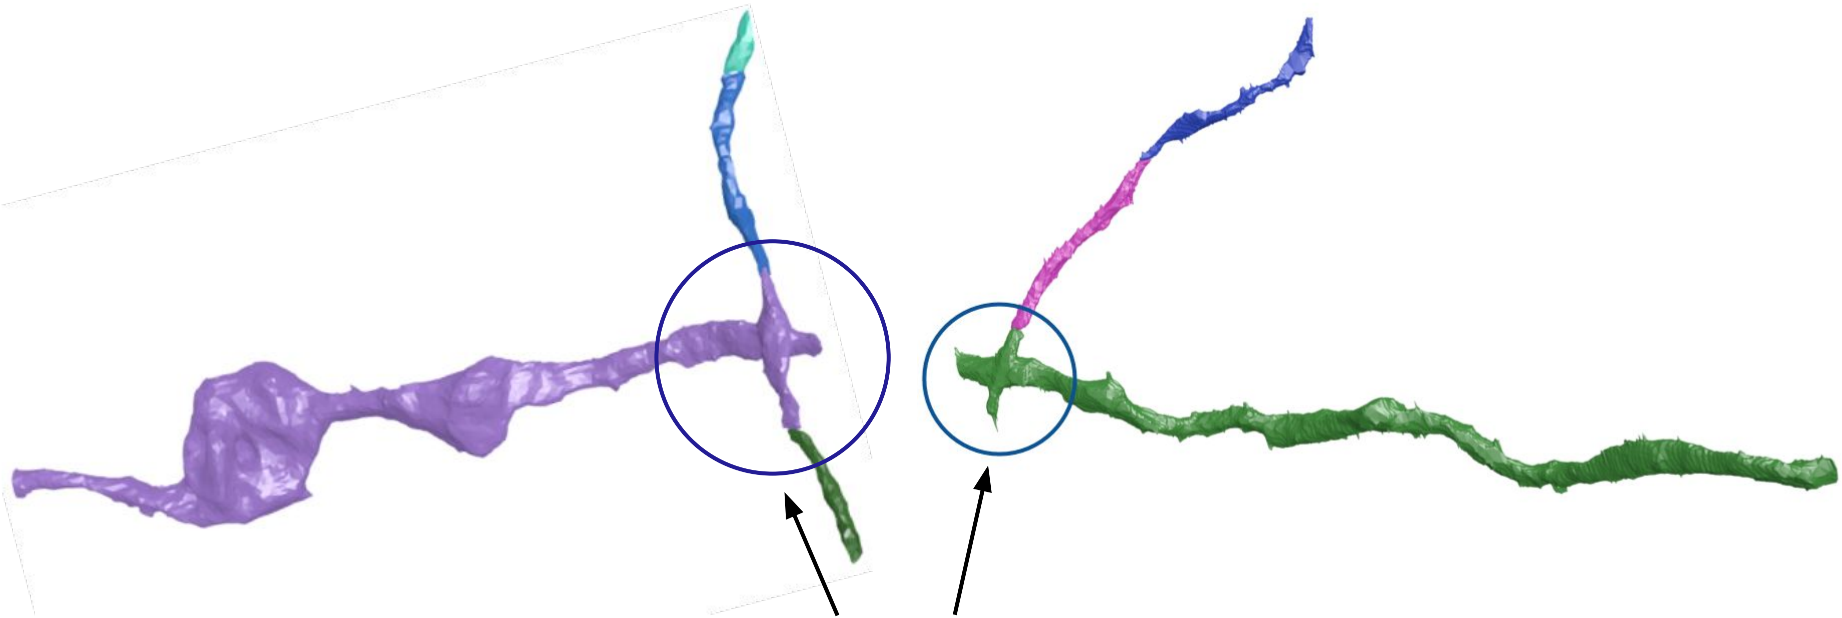

Errors in Input Segmentation

Ablation Studies: Node Generation

Goal: Merge all small segments with a nearby larger segment from the same neuronal process

Ablation Studies: Node Generation

Goal: Merge all small segments with a nearby larger segment from the same neuronal process

Baseline: How many small segments belong to the same neuron as the high affinity large neighbor?

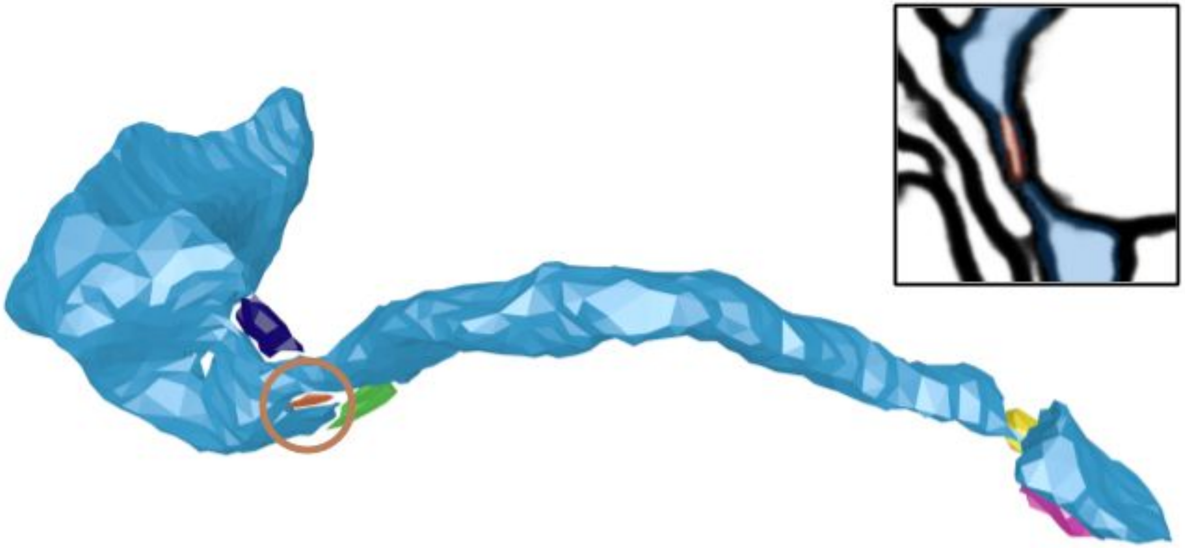

Ablation Studies: Node Generation

Goal: Merge all small segments with a nearby larger segment from the same neuronal process

Baseline: How many small segments belong to the same neuron as the high affinity large neighbor?

Dataset	**Baseline** (↑)	**Proposed** (↑)
PNI Test One	305/521 (36.9%)	686/129 (80.2%)
PNI Test Two	185/281 (39.7%)	444/75 (85.5%)
Kasthuri Test	4,514/8,604 (52.5%)	6,623/2,020 (76.6%)

The number of correctly merged small segments versus the number of incorrectly merged segments

Ablation Studies: Edge Generation

Goal: Identify all split errors while minimizing the number of total edges in the graph

Ablation Studies: Edge Generation

Goal: Identify all split errors while minimizing the number of total edges in the graph

Baseline: How many total edges are there in the adjacency graph?

Dataset	**Baseline**	**Proposed**	**Edge Recall** (↑/↓)
PNI Test One	528 / 25,619	417 / 10,074	79.0% / 39.3%
PNI Test Two	460 / 30,388	370 / 11,869	80.4% / 39.1%
Kasthuri Test	1,193 / 43,951	936 / 18,168	78.5% / 41.3%

The number of edges in the graph that correspond to split errors, the total number of edges, and the recall

Ablation Studies: Edge Generation

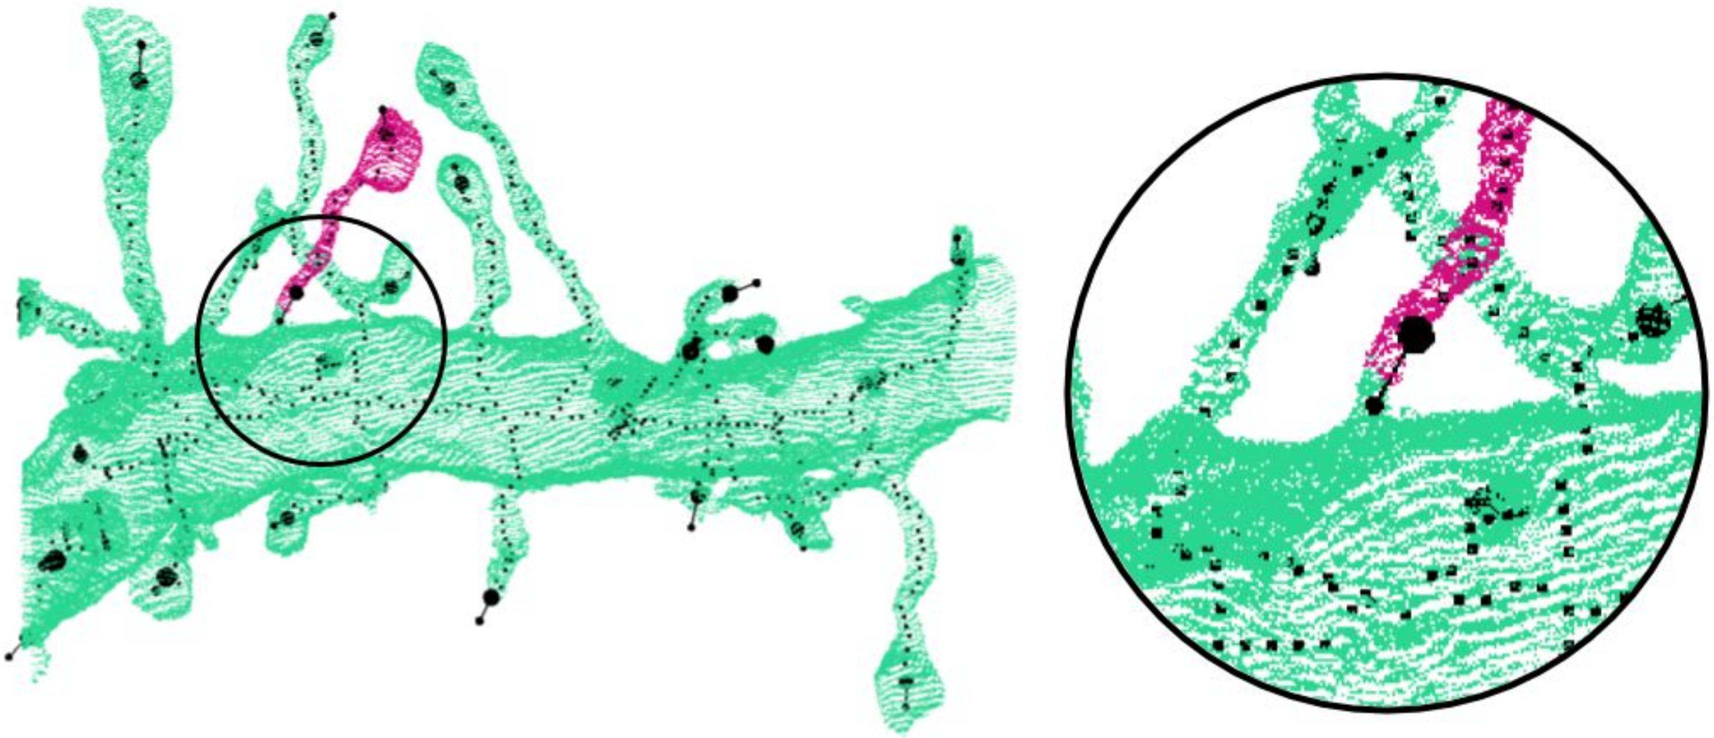

Ablation Studies: Edge Generation

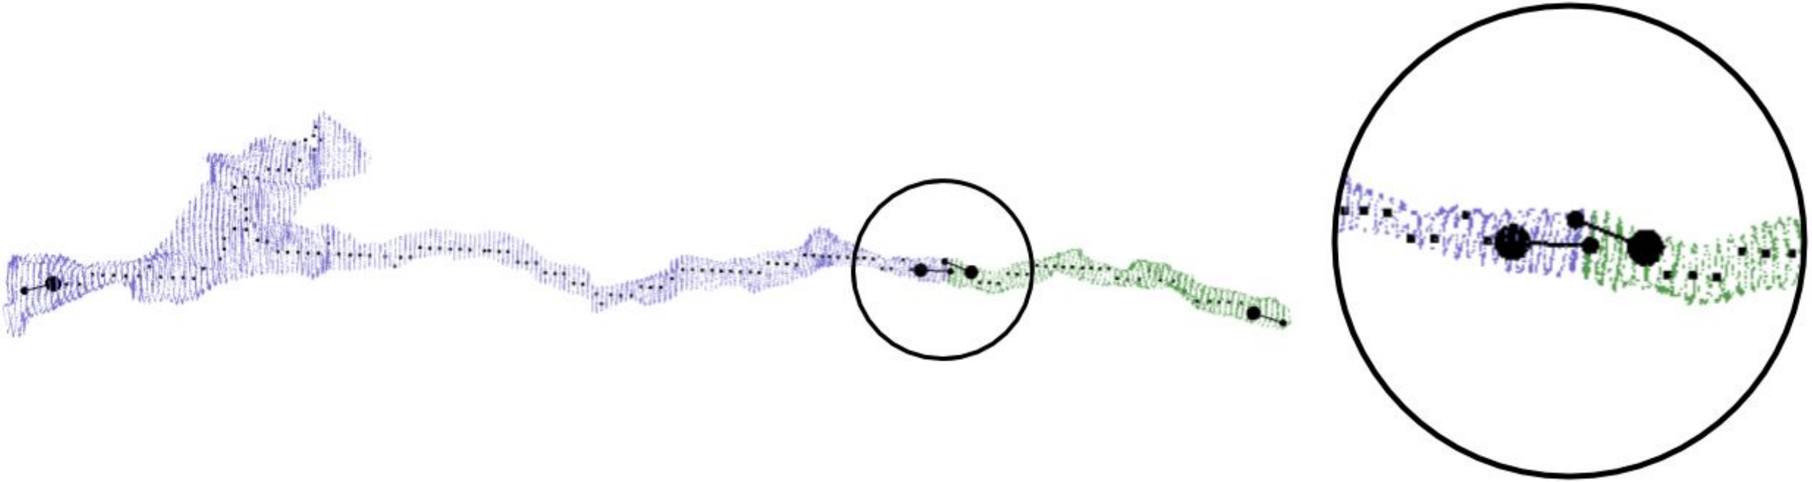

Edge Generation Failure Cases

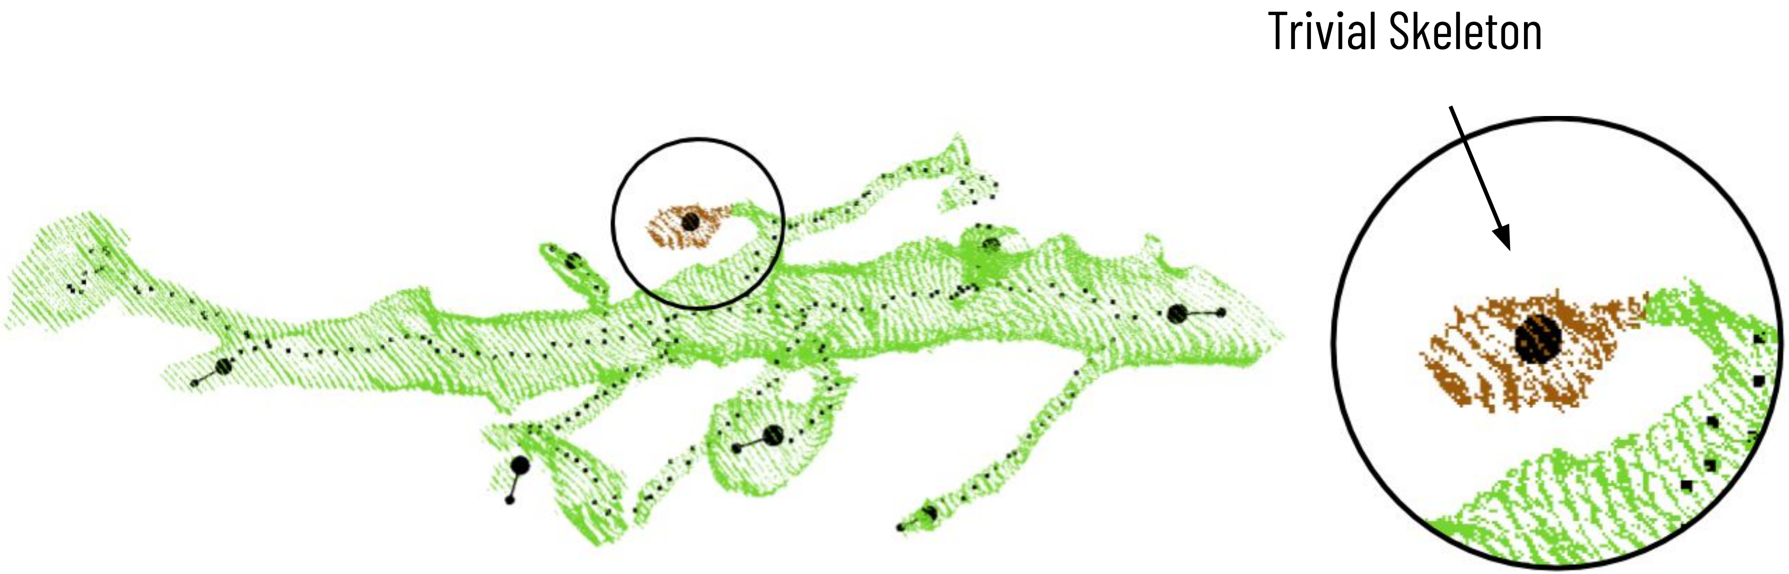

Edge Generation Failure Cases

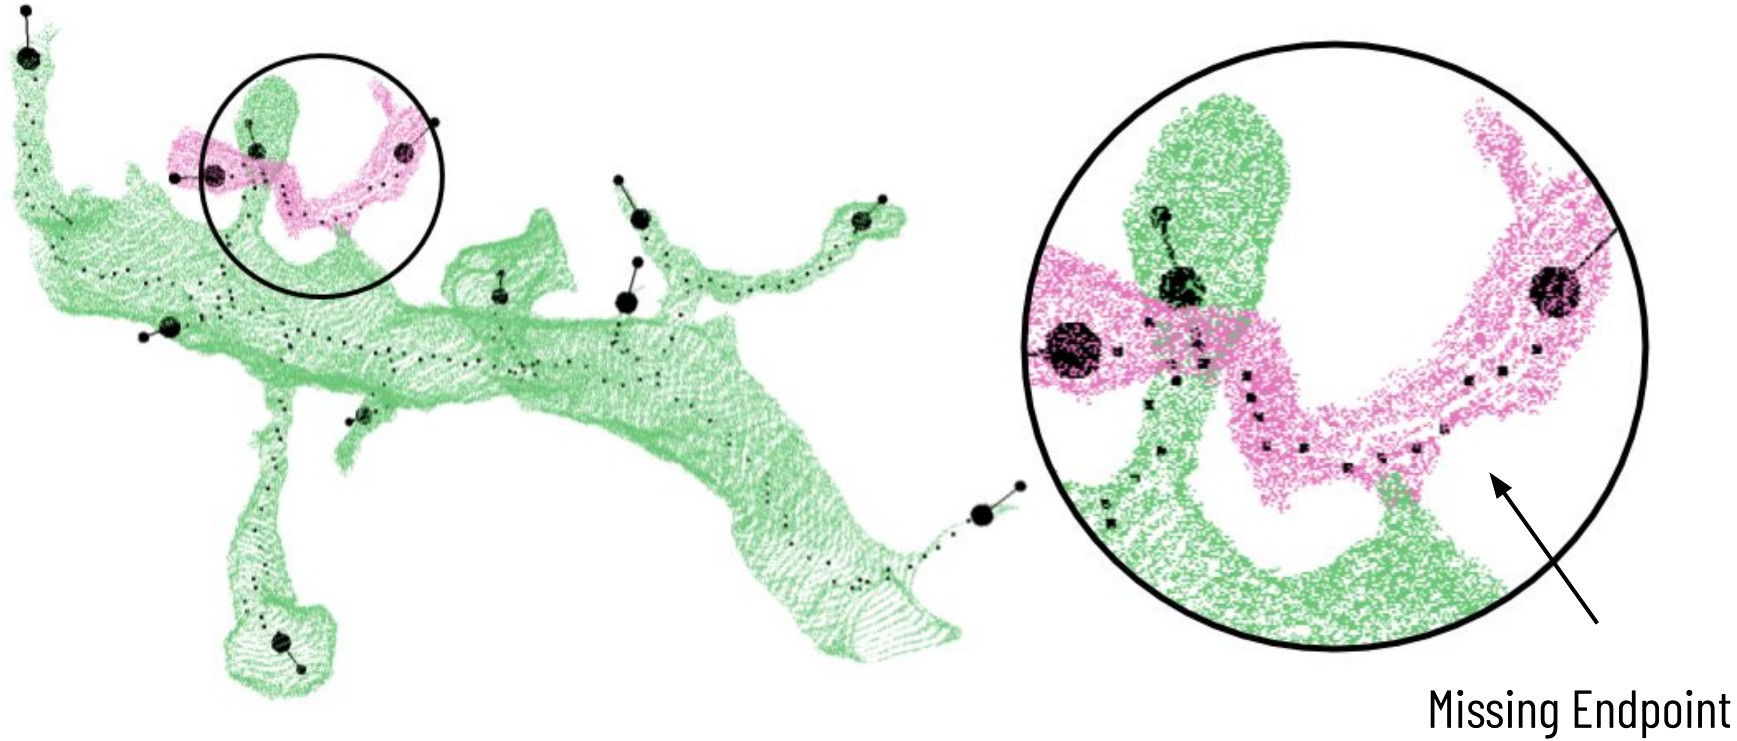

Ablation Studies: Edge Weight Assignment

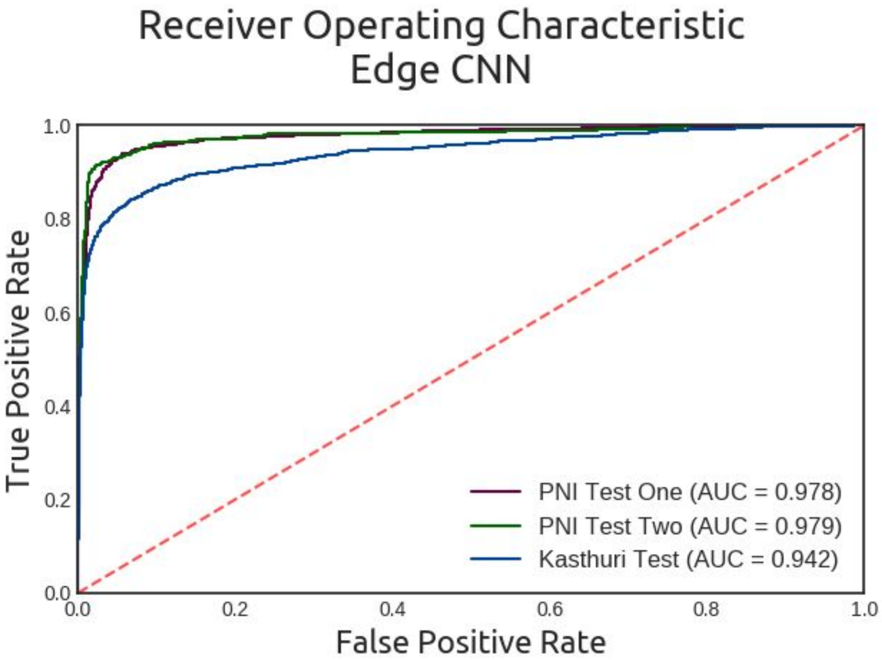

Accuracies:

PNI Test One: 96.4%
PNI Test Two: 97.2%
Kasthuri: 93.4%

Running Times

Time to process a gigavoxel dataset

Step	Running Time
Node Generation	281 seconds
Edge Generation	351 seconds
Lifted Multicut	13 seconds
Total	10.75 minutes

Efficient Correction for EM Connectomics with Skeletal Representation

Konstantin Dmitriev[1], Toufiq Parag[2], Brian Matejek[2], Arie Kaufman[1], Hanspeter Pfister[2]

[1]Stony Brook University

[2]Harvard University

Harvard John A. Paulson
School of Engineering and
Applied Sciences

Visual
Computing
Group

Addressing Merge Errors

Hard to efficiently correct since the search space of possible errors is very large

Addressing Merge Errors

Hard to efficiently correct since the search space of possible errors is very large

However, still occur in all segmentation strategies

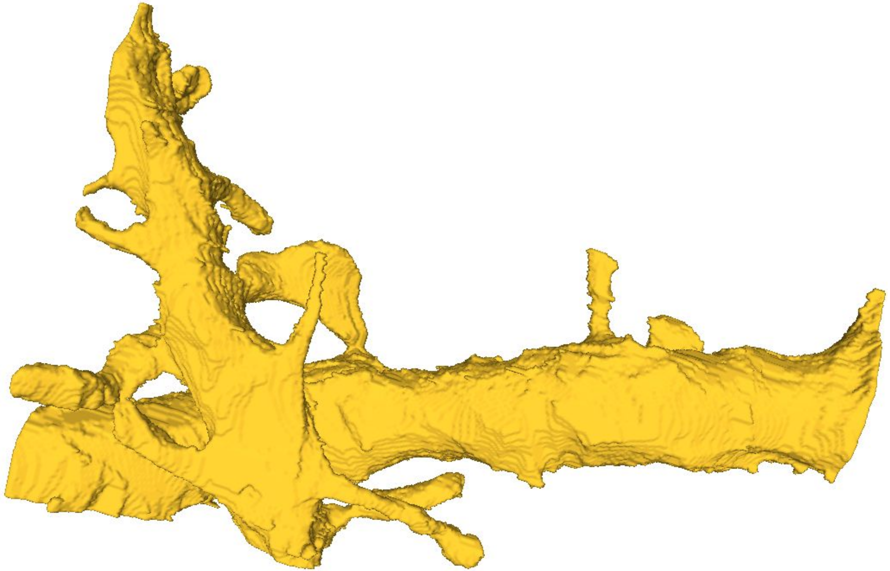

Addressing Merge Errors

Hard to efficiently correct since the search space of possible errors is very large

However, still occur in all segmentation strategies

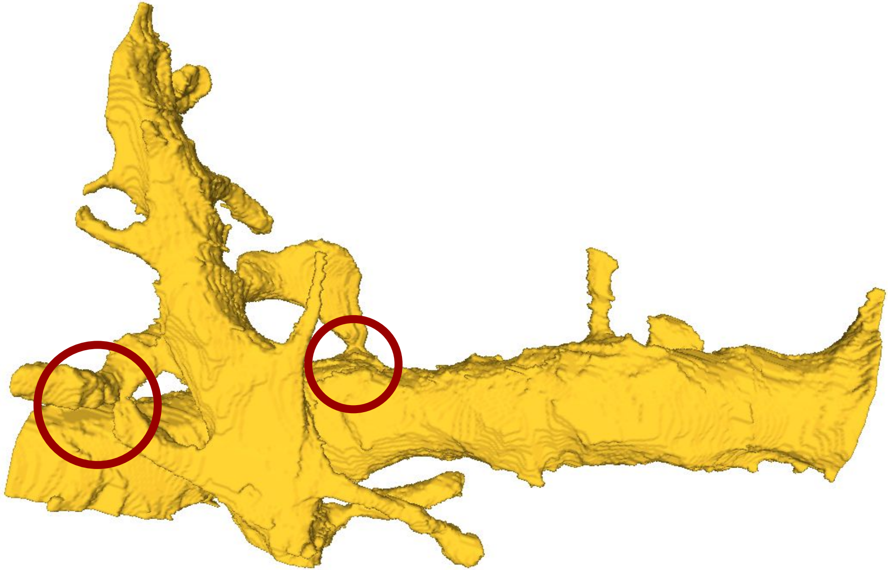

Addressing Merge Errors

Hard to efficiently correct since the search space of possible errors is very large

However, still occur in all segmentation strategies

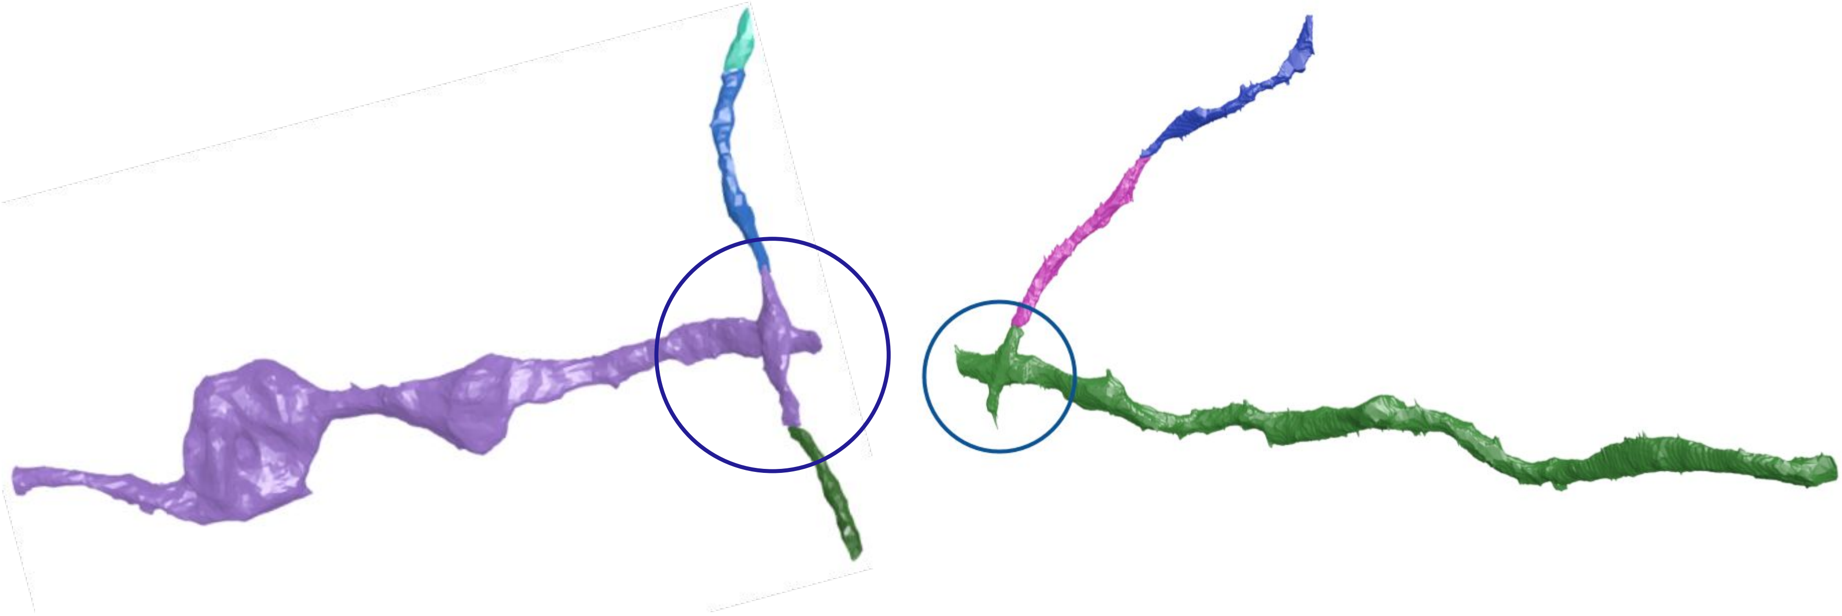

Skeleton Generation

Use NeuTu to generate skeletons for all segments

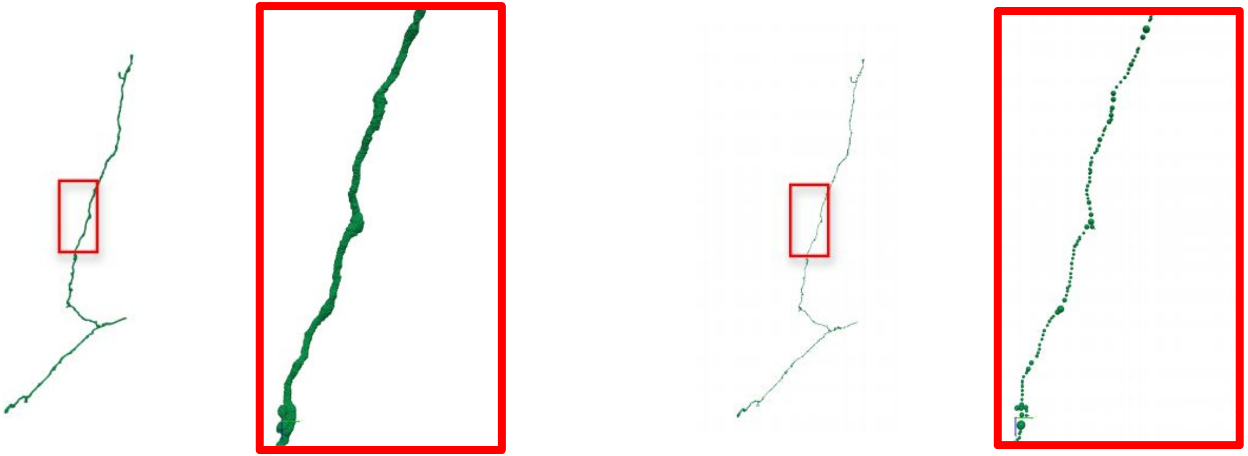

Zhao, T. et al., NeuTu: Software for Collaborative, Large Scale, Segmentation-Based Connectome Reconstruction, Frontiers in Neural Circuits, 2018.

Skeleton Generation

Use NeuTu to generate skeletons for all segments

Reduce search locations to skeleton points

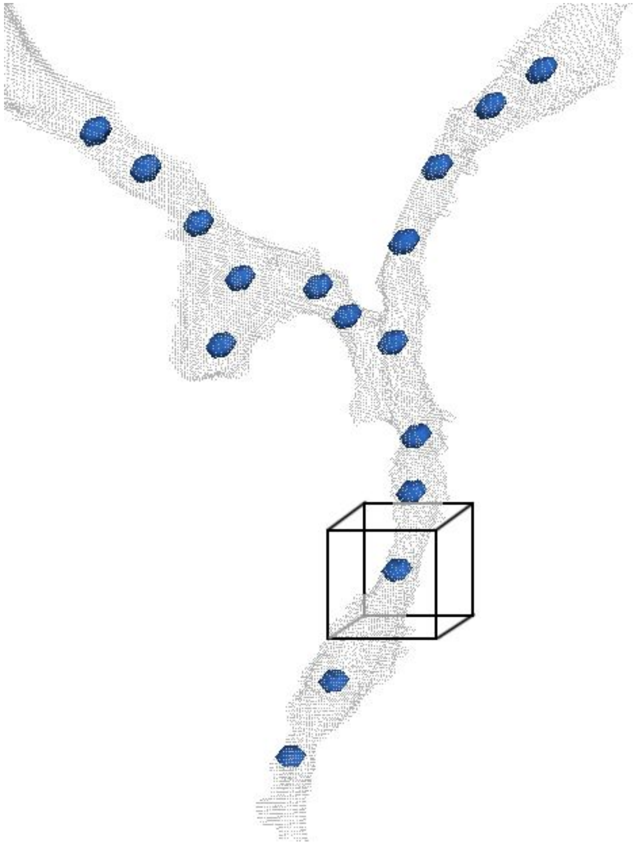

Skeleton Generation

Use NeuTu to generate skeletons for all segments

Reduce search locations to skeleton points

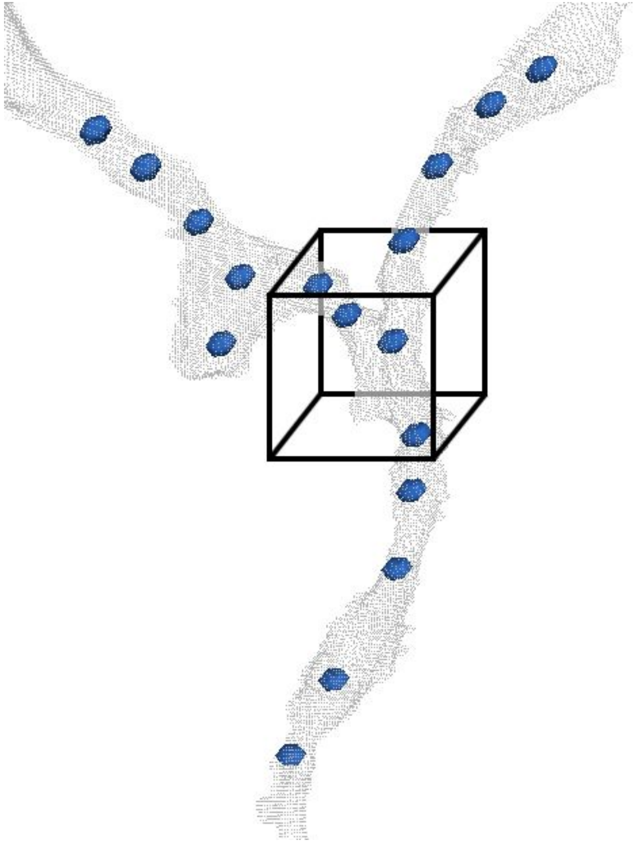

Identifying Merge Errors

We use a CNN that takes the segmentation and image data as input and produces a split candidate

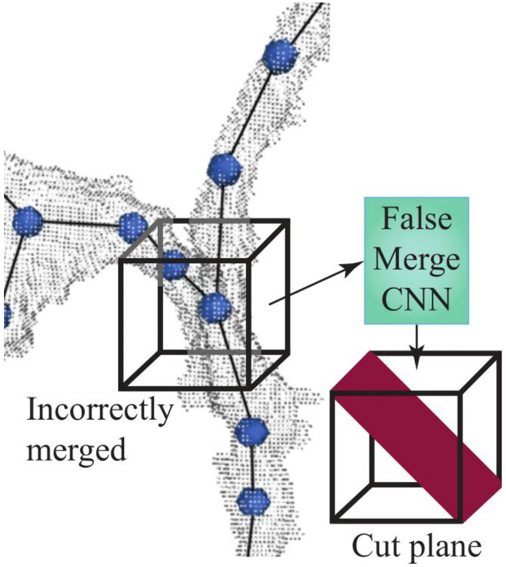

Identifying Merge Errors

We use a CNN that takes the segmentation and image data as input and produces a split candidate

Correctly segmented regions return no

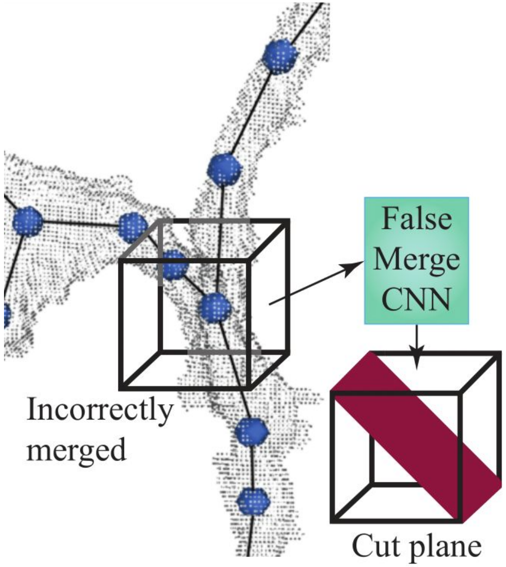

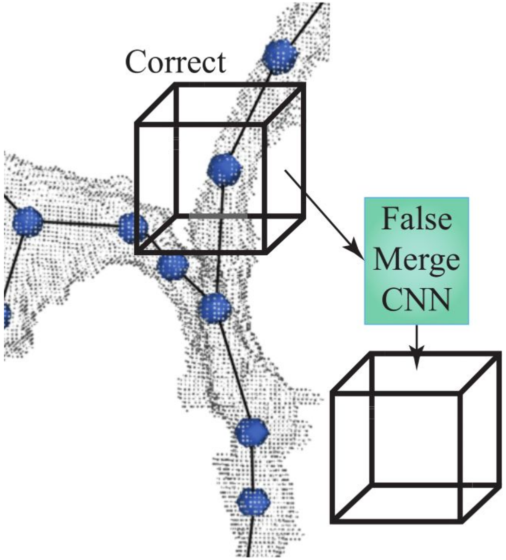

Merge Correction Example

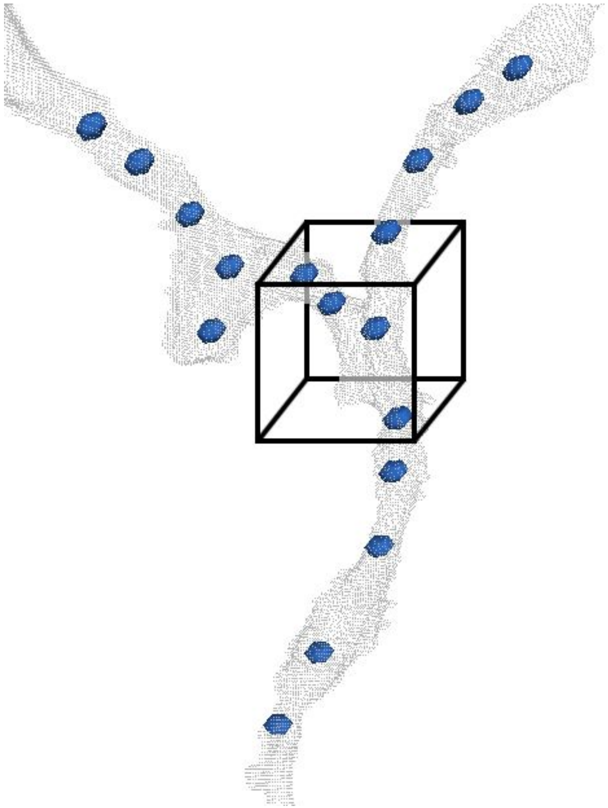

Merge Correction Example

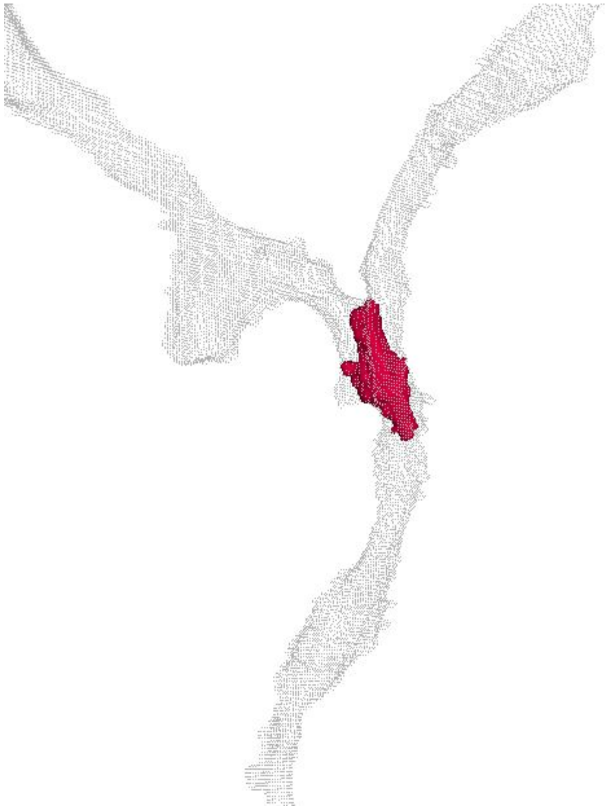

Merge Correction Example

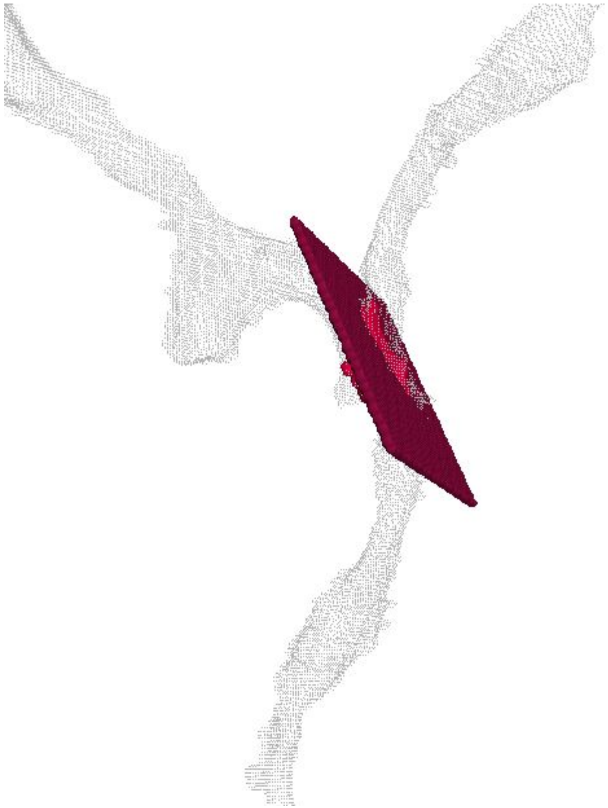

Merge Correction Example

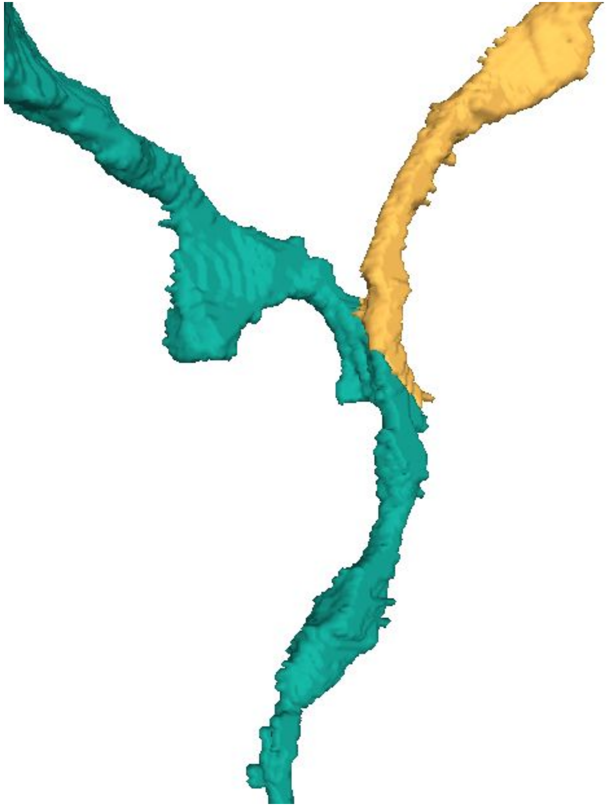

Network Architecture

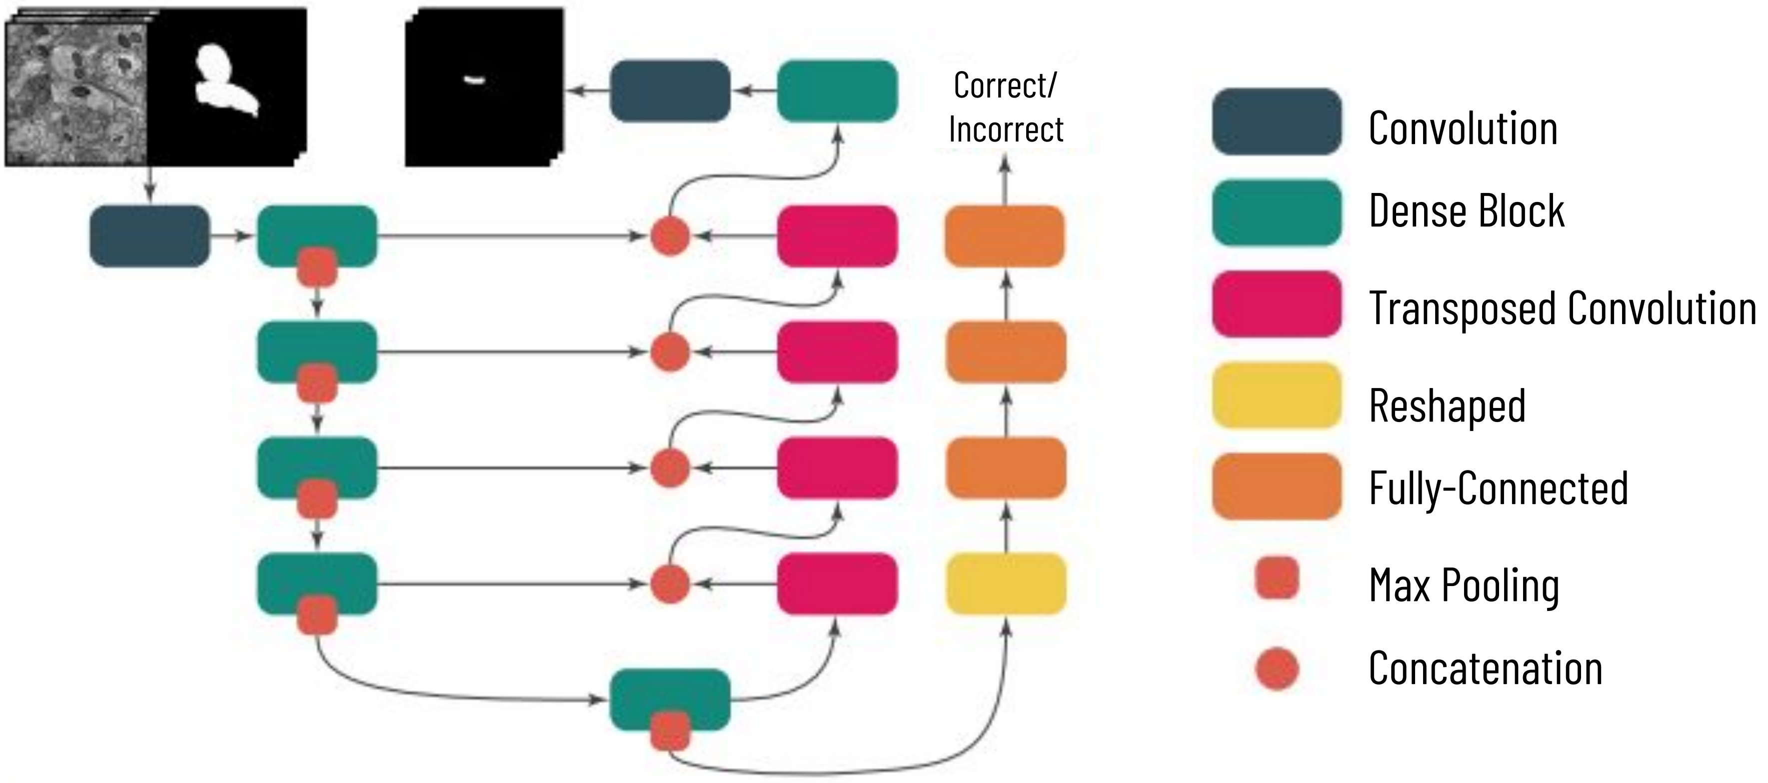

Datasets

Kasthuri

Princeton Neuroscience Institute

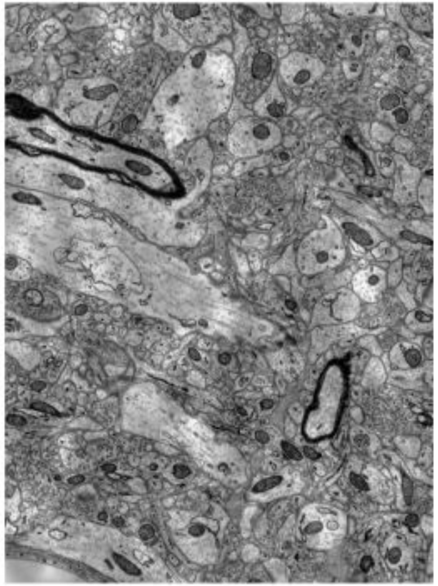 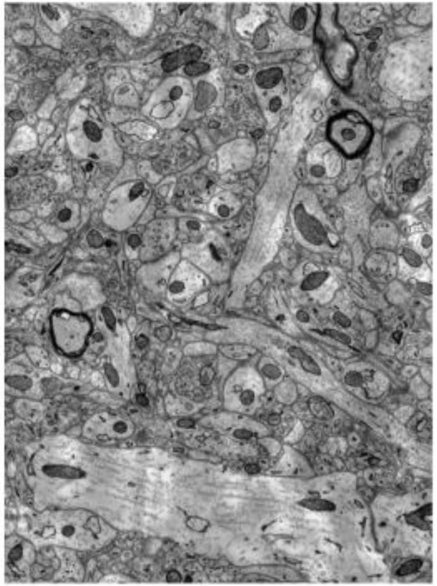

2 Volumes

6 x 6 x 30 nm^3 / vx

1335 x 1809 x 338 vx

8.01 x 10.85 x 10.14 μm^3

Datasets

Kasthuri

Princeton Neuroscience Institute

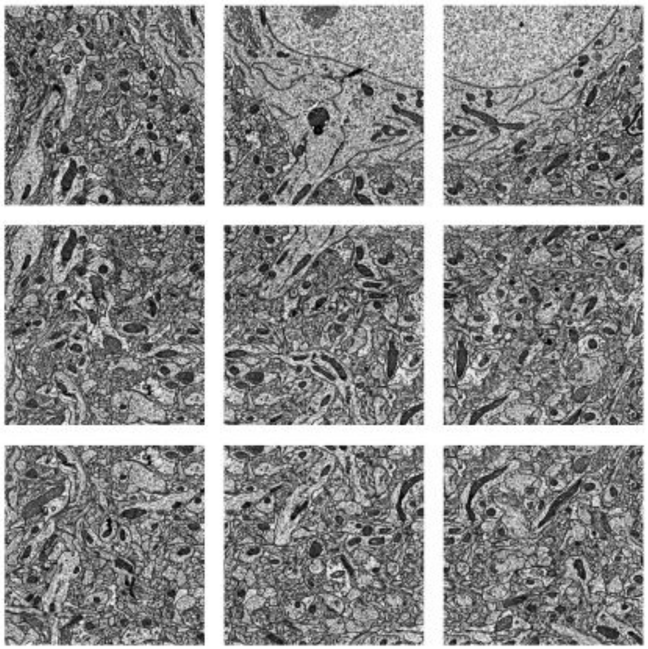

9 Volumes

$3.6 \times 3.6 \times 40\ nm^3$ / vx

2048 x 2048 x 256 vx

$7.37 \times 7.37 \times 10.24\ \mu m^3$

Results

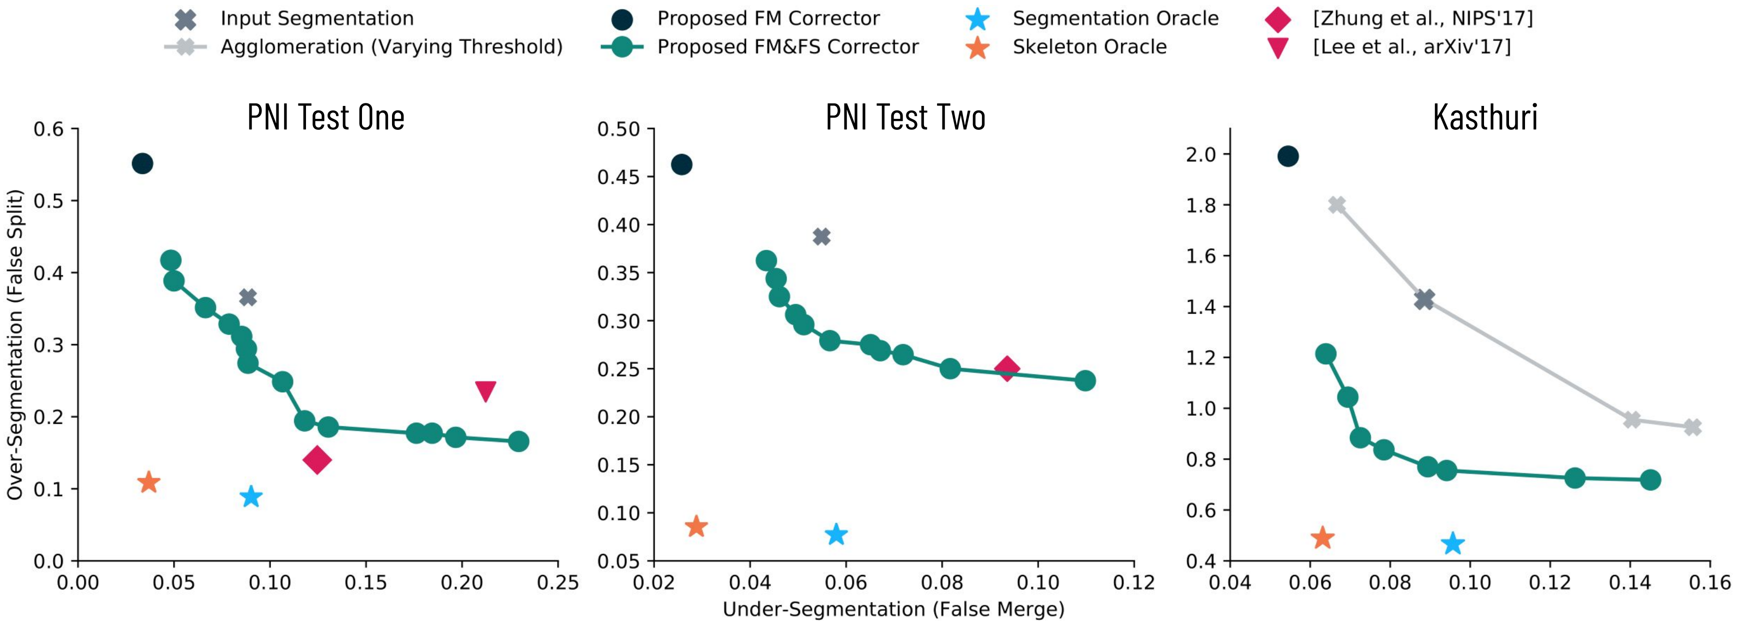
Input Segmentation
Agglomeration (Varying Threshold)
Proposed FM Corrector
Proposed FM&FS Corrector
Segmentation Oracle
Skeleton Oracle
[Zhung et al., NIPS'17]
[Lee et al., arXiv'17]
PNI Test One
PNI Test Two
Kasthuri
Over-Segmentation (False Split)
Under-Segmentation (False Merge)

Search Space Reduction

Size	PNI Test One	PNI Test Two	Kasthuri
Volume Size	1.074×10^9 vx	1.074×10^9 vx	0.816×10^9 vx
Query Points	40,621 pts	41,513 pts	62,815 pts
Search Reduction	26,433x	25,865x	12,994x

Research Sponsors

This research was supported in part by NSF grants IIS-1447344, IIS-1607800, IIS-1527200, IIS-1607800, NRT-1633299, and CNS-1650499.

This research was supported in part by IARPA contract DI6PC00002.

Visual Computing Group

Hanspeter Pfister

Thank you!

Questions?