Segmentation of Electron Microscopy Images in Connectomics

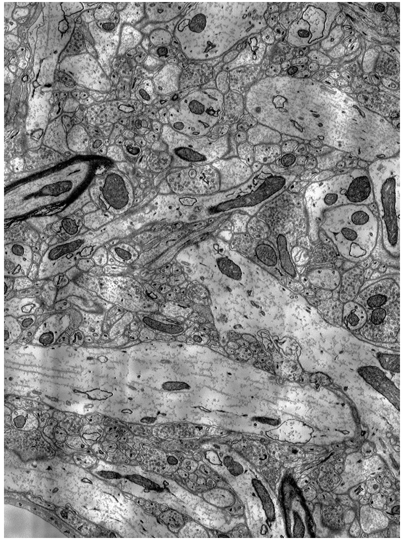

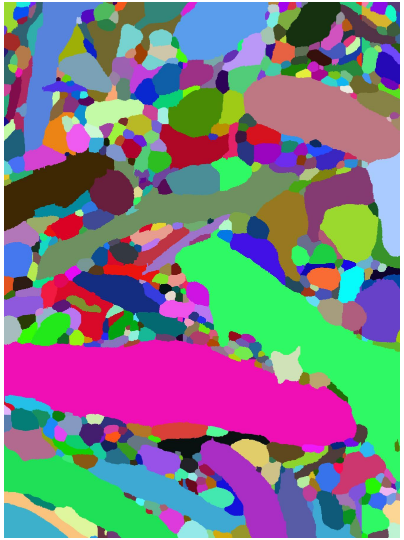

Brian Matejek
Advisor: Hanspeter Pfister

Connectomics

Goal: Extract the wiring diagram from a brain

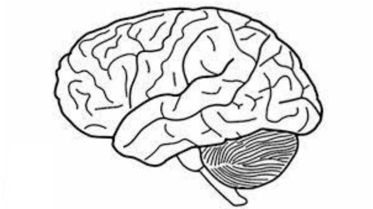

Connectomics

Goal: Extract the wiring diagram from a brain

Nano-resolution Imaging

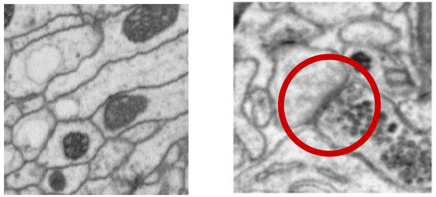

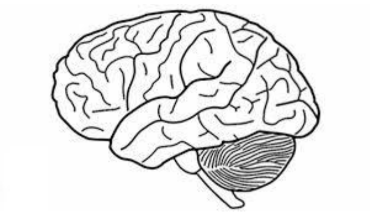

Connectomics

Goal: Extract the wiring diagram from a brain

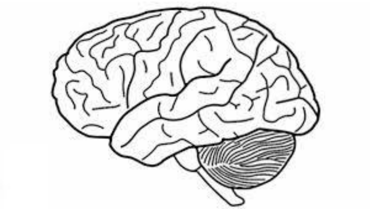

Nano-resolution Imaging

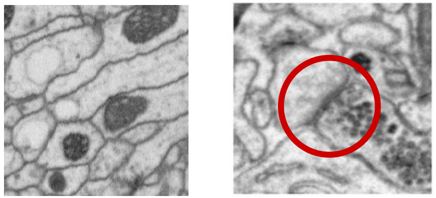

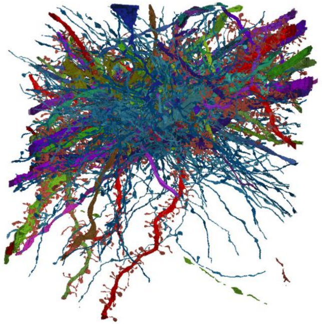

Circuit
Reconstruction

Connectomics

Goal: Extract the wiring diagram from a brain

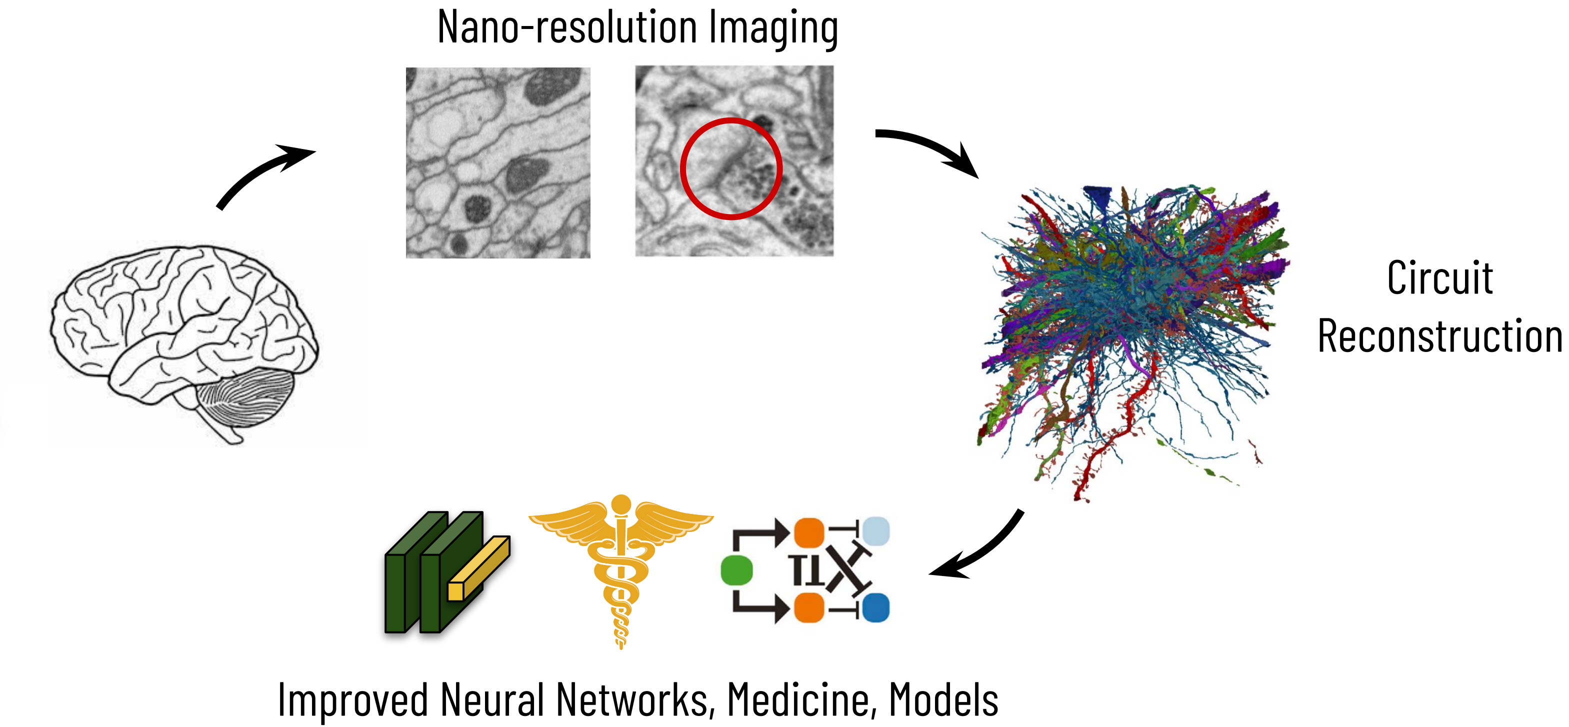

Connectomics

Goal: Extract the wiring diagram from a brain

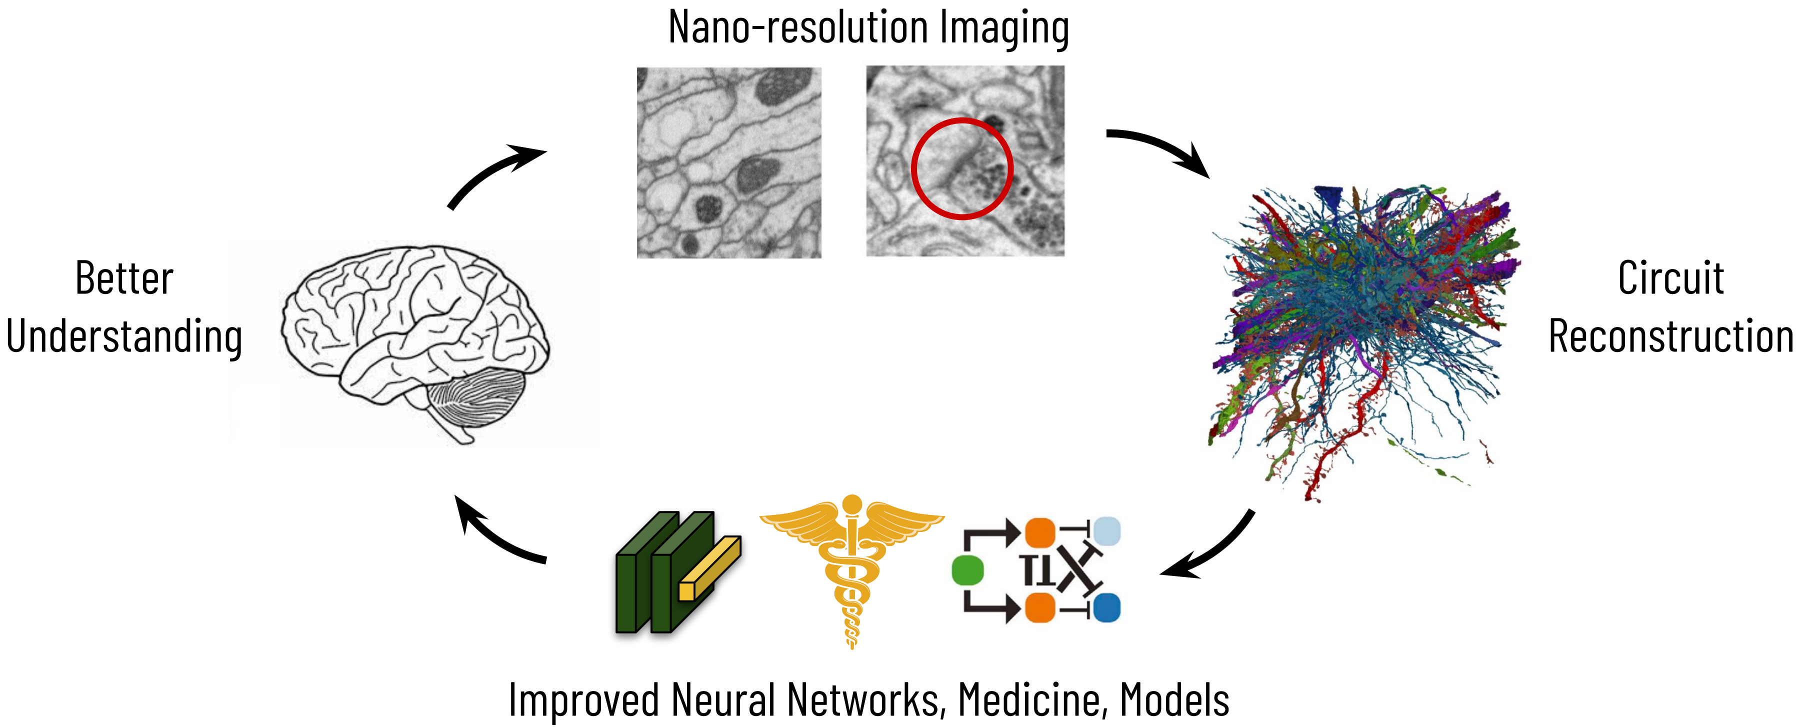

Connectomics

Goal: Extract the wiring diagram from a brain

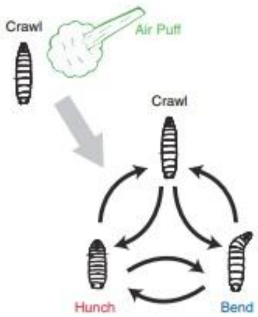

Behavior

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

Connectomics

Goal: Extract the wiring diagram from a brain

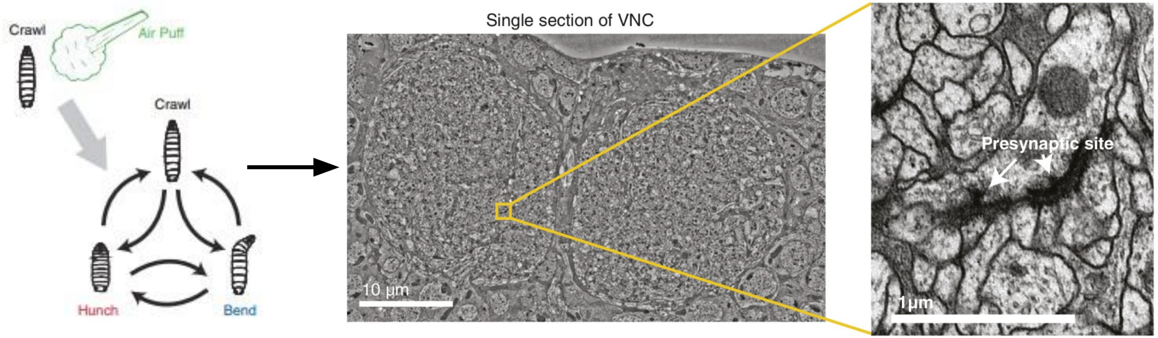

Behavior

Structure

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

Connectomics

Goal: Extract the wiring diagram from a brain

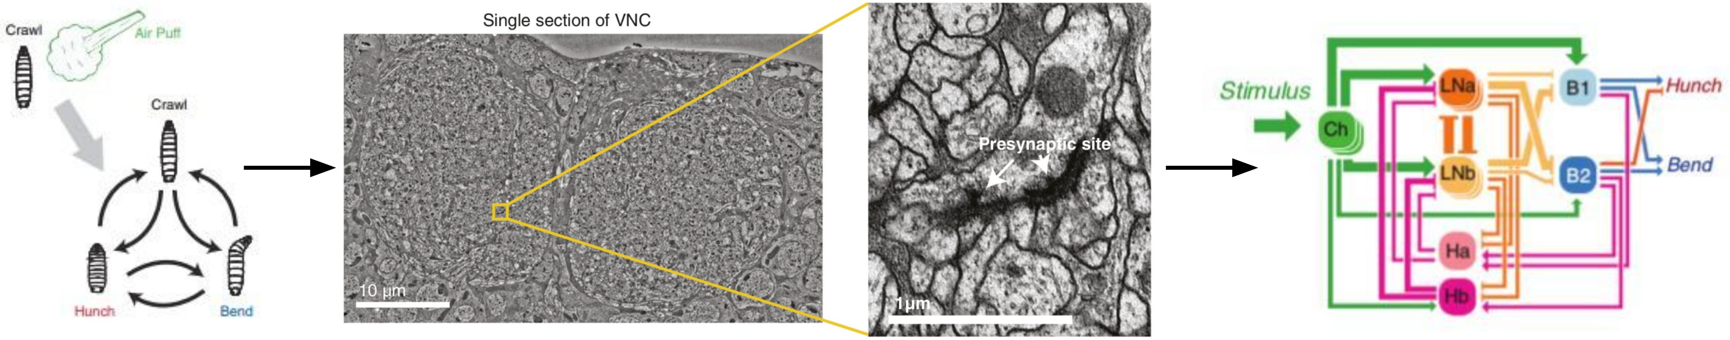

Behavior

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

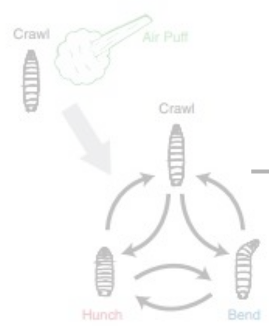

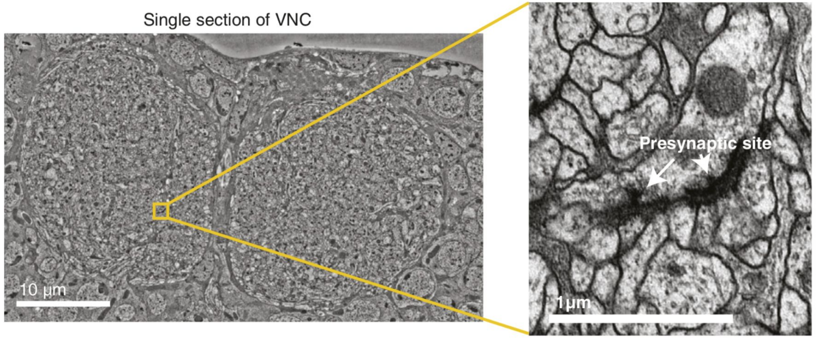

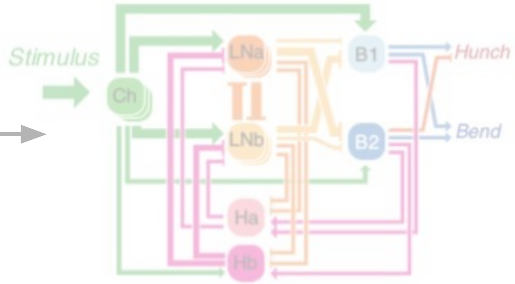

Behavior

Structure

Function

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

Connectomics Pipeline

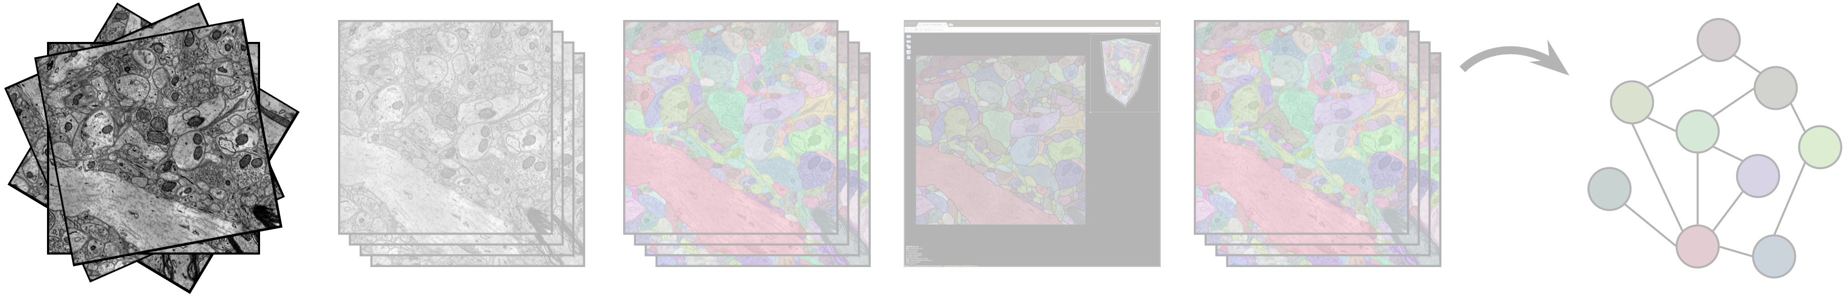

Acquisition Registration Segmentation Proofreading Network Analysis

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

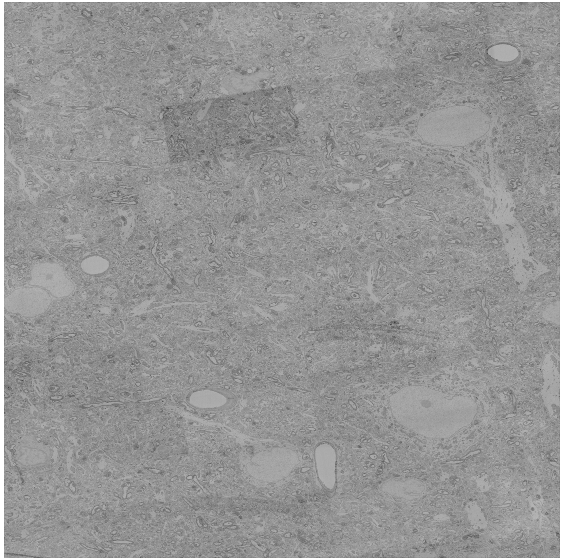

100 µm

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

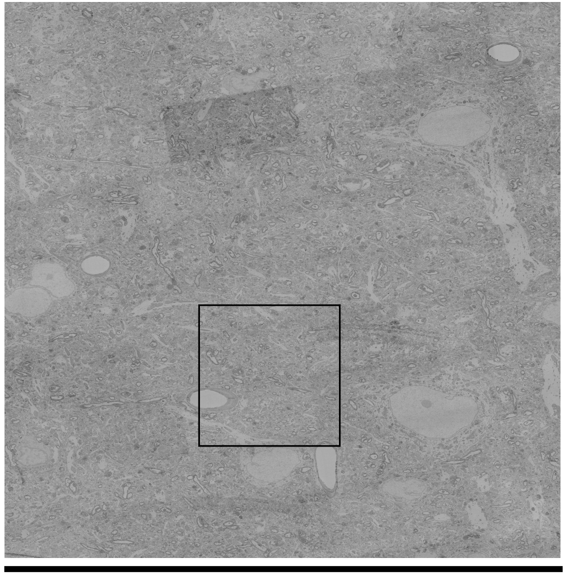

100 μm

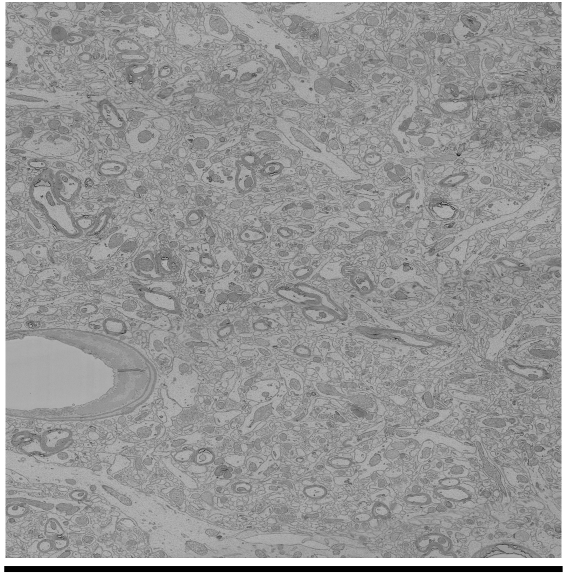

25 μm

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

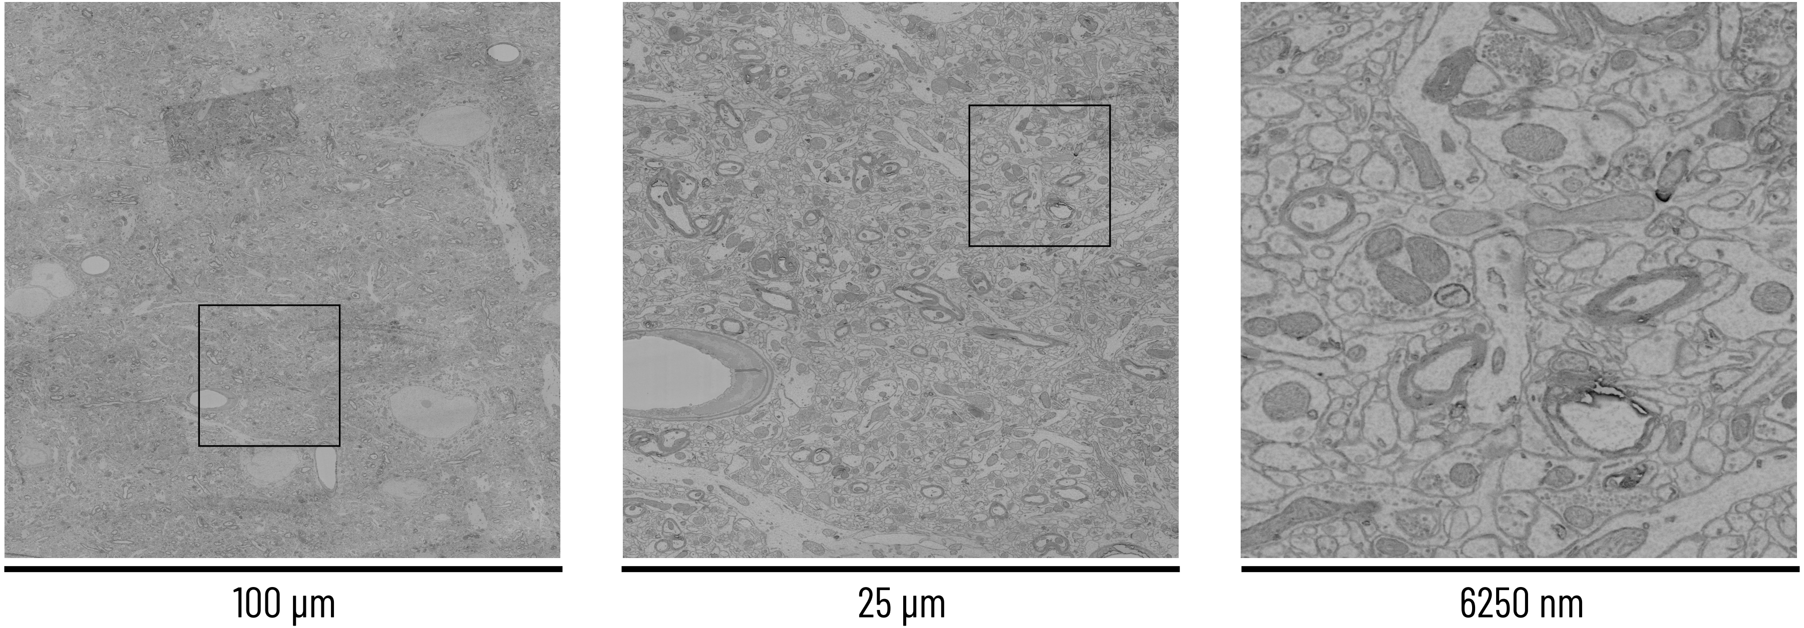

Connectomics Pipeline

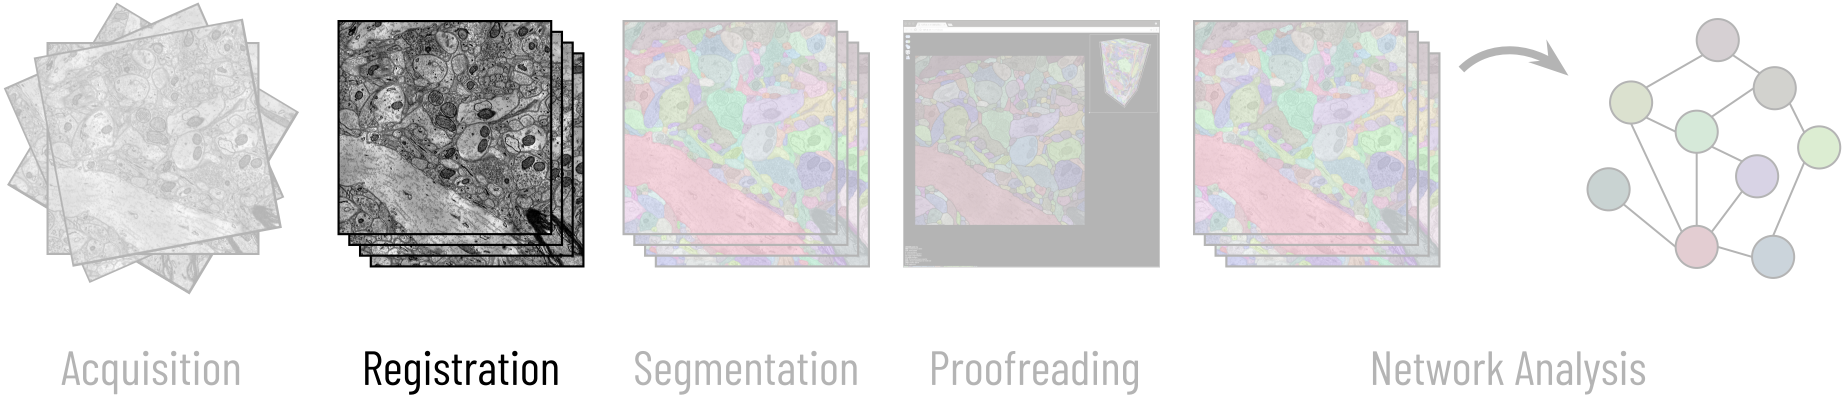

Saalfield et al., Elastic Volume Reconstruction from Series of Ultra-thin Microscopy Sections, Nature 2012
Khairy et al., Joint Deformable Registration of Large EM Image Volumes: A Matrix Solver Approach, 2018

Registration

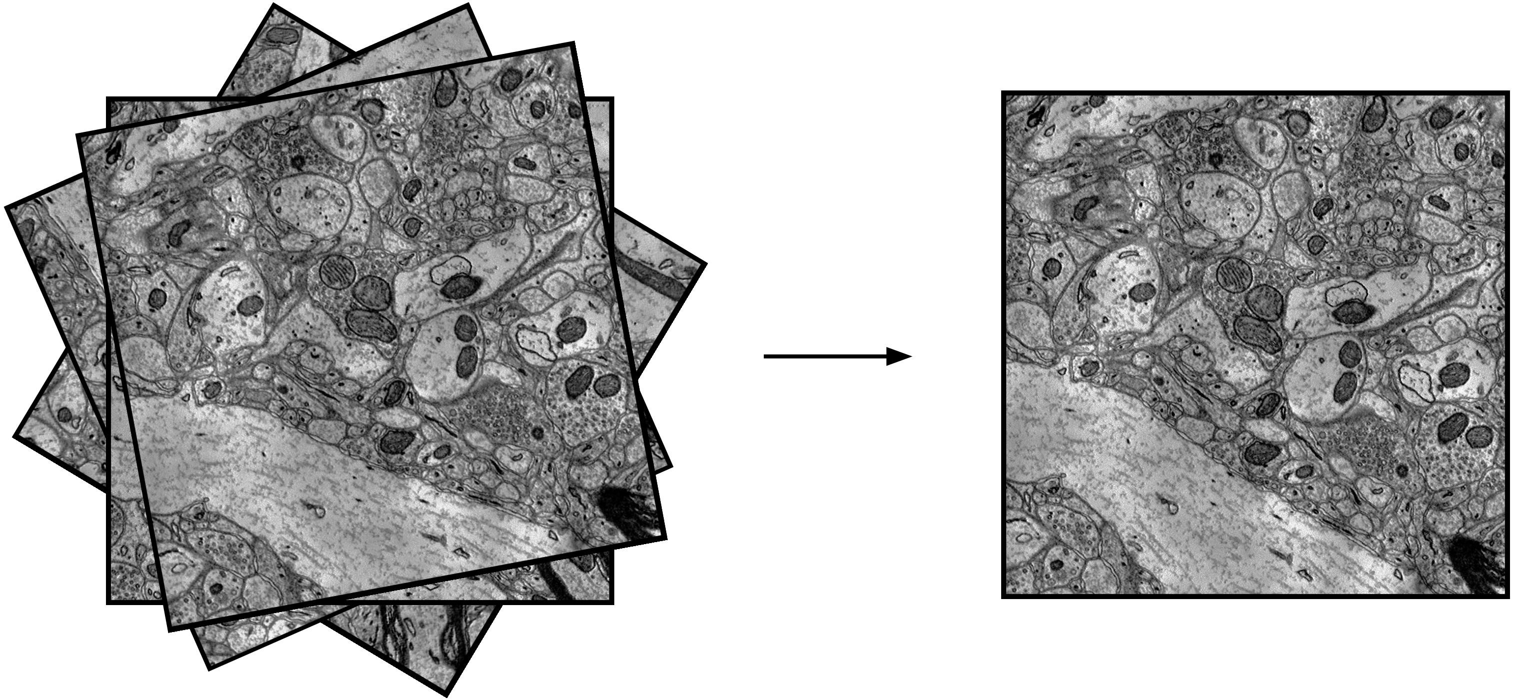

Connectomics Pipeline

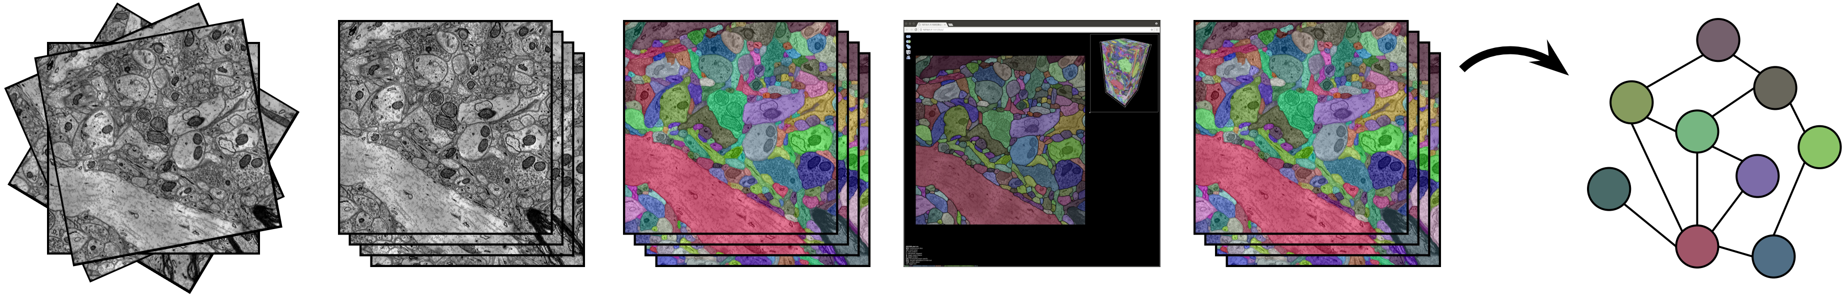

Acquisition Registration Segmentation Proofreading Network Analysis

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

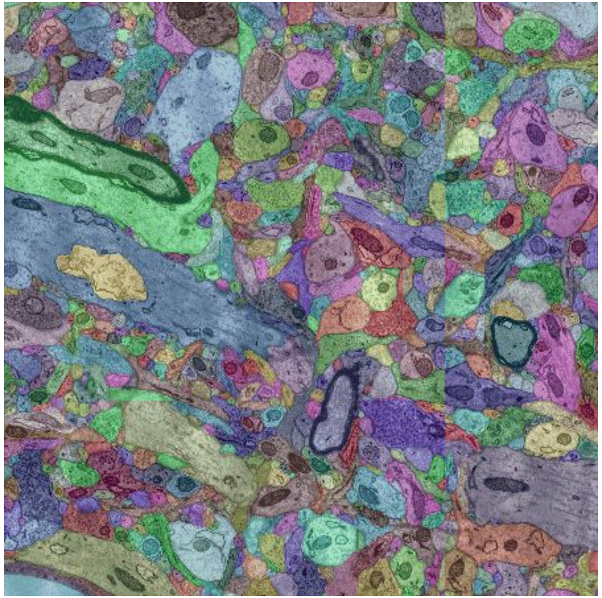

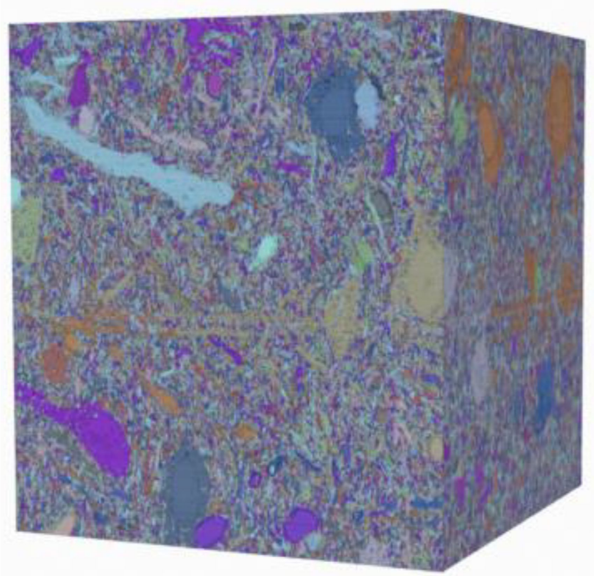

Label Volumes

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

Typically use 64 bits per voxel to label each segment uniquely

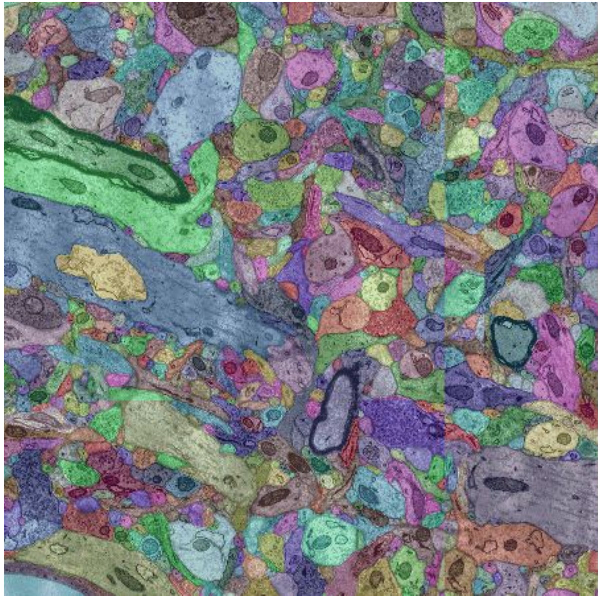

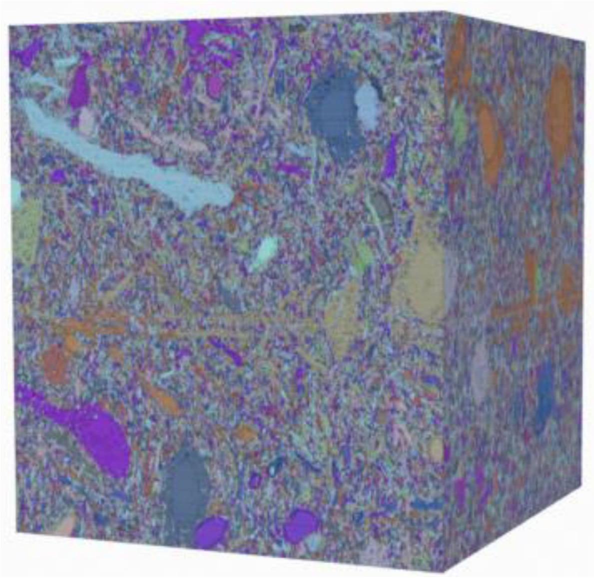

Connectomics Pipeline

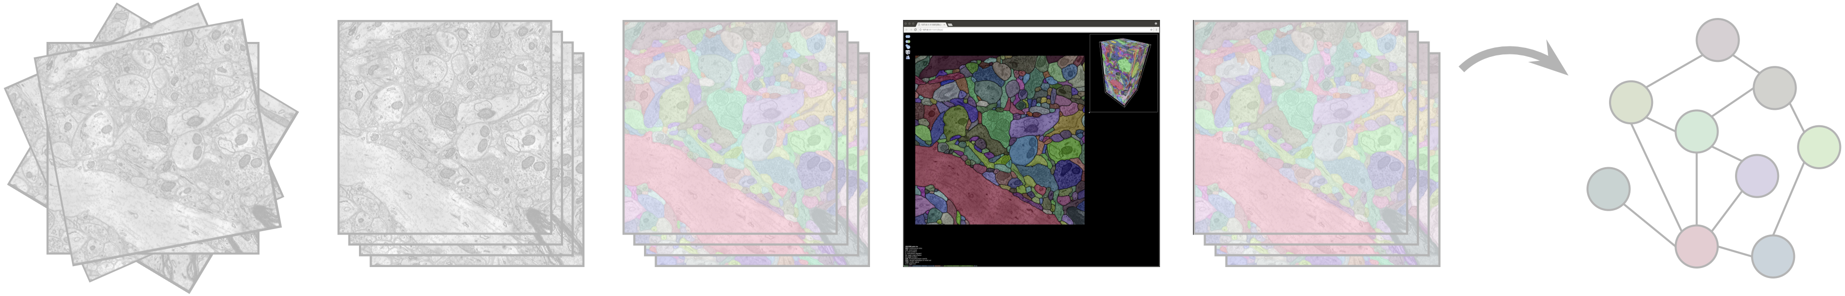

Acquisition Registration Segmentation **Proofreading** Network Analysis

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

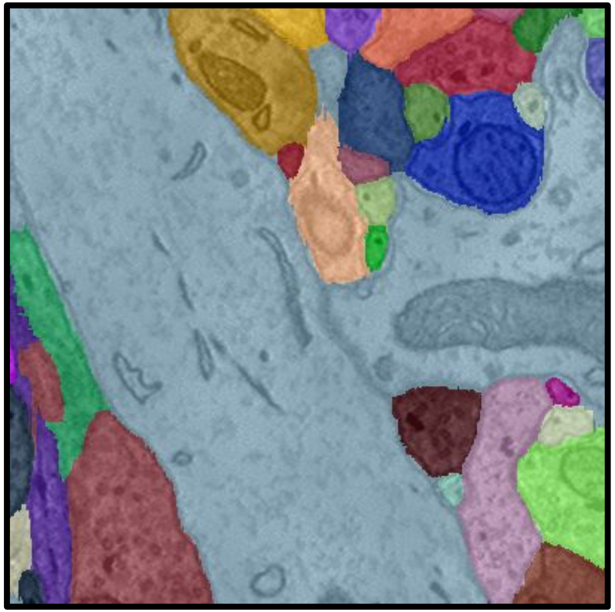

Ground Truth

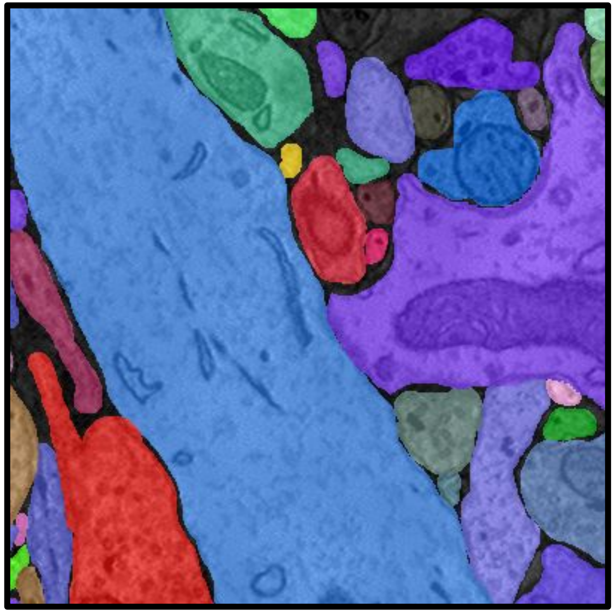

Merge Errors

Automatic Segmentation

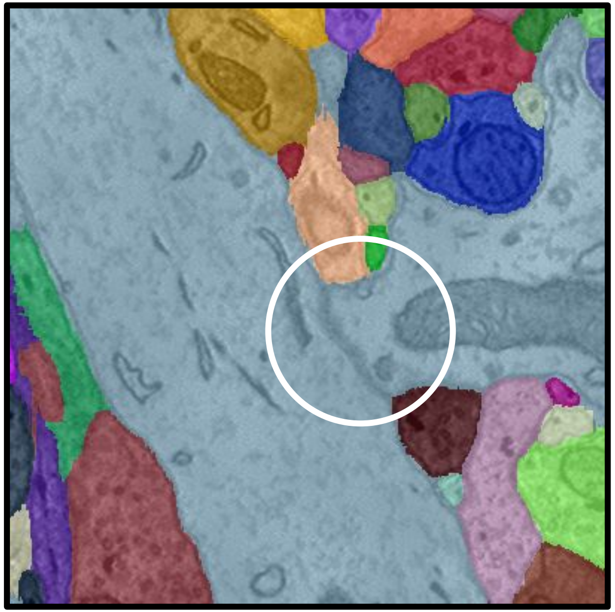

Ground Truth

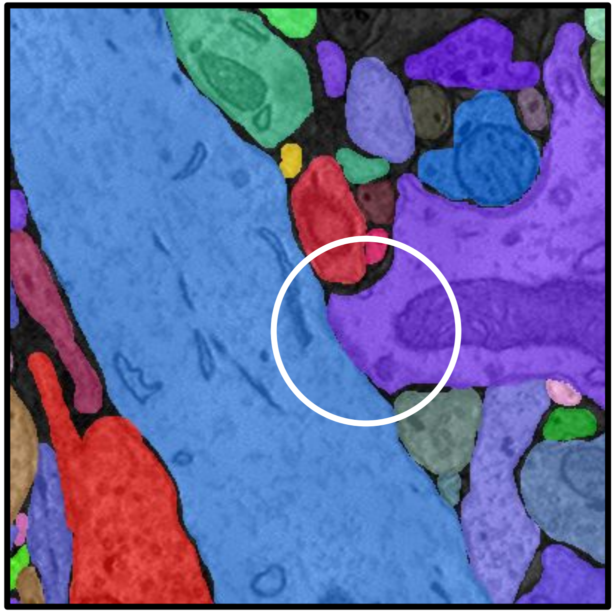

Split Errors

Automatic Segmentation

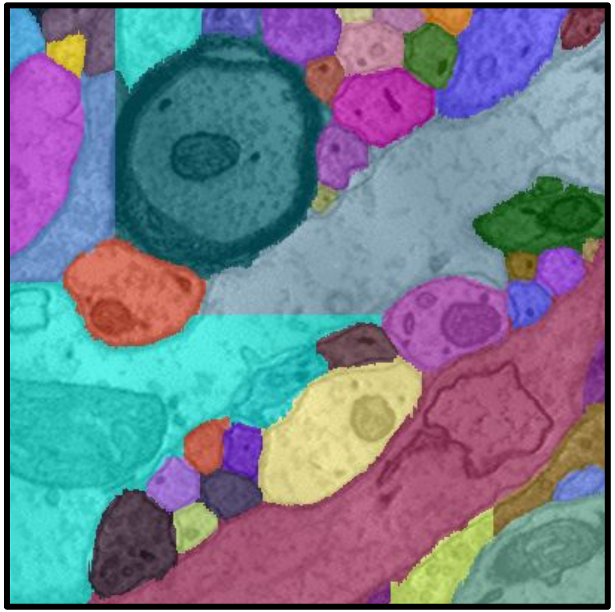

Ground Truth

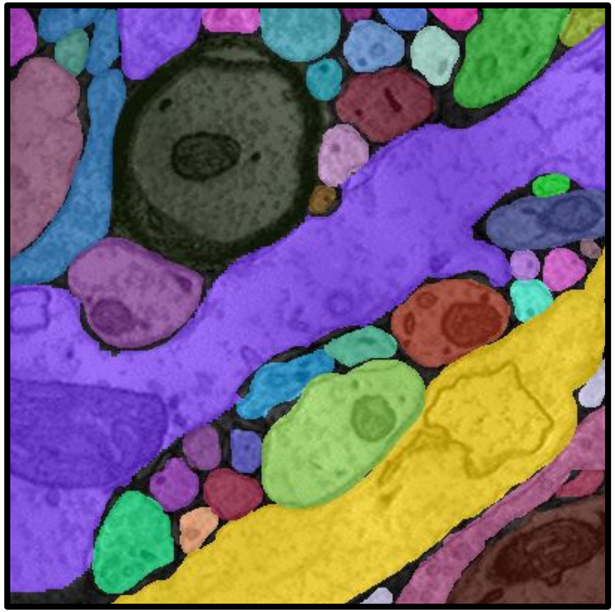

Split Errors

Automatic Segmentation

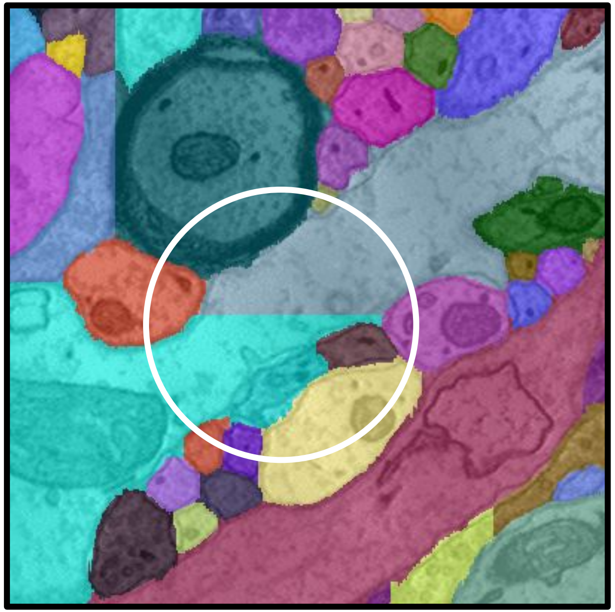

Ground Truth

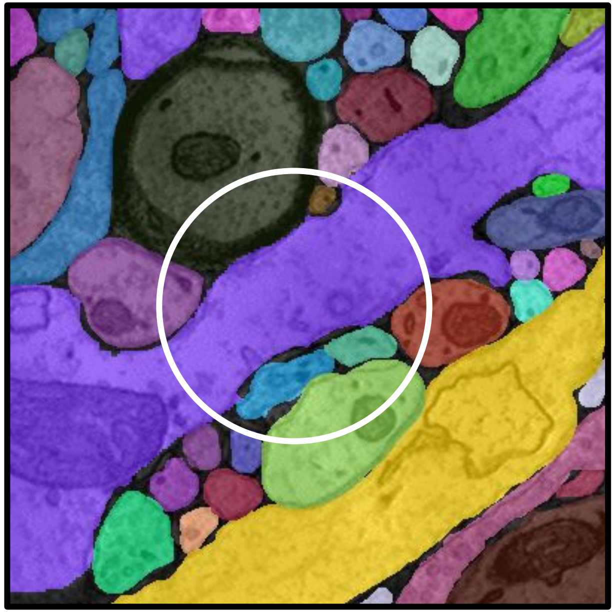

Connectomics Pipeline

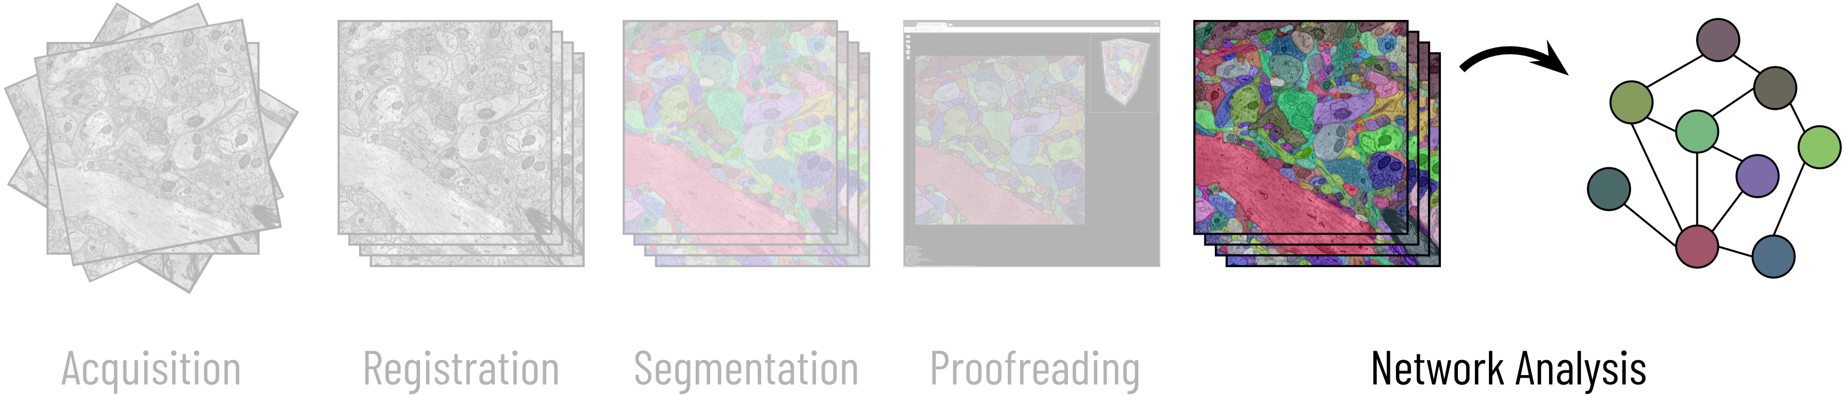

Sorger et al., neuroMAP - Interactive Graph-Visualization of the Fruit Fly's Neural Circuit, BioVIS 2013
Al-Awami et al., NeuroLines: A Subway Map Metaphor for Visualizing Nanoscale Neuronal Connectivity, IEEE VIS 2014
Haehn et al., Scalable Interactive Visualization for Connectomics, MDPI Informatics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

Network Analysis

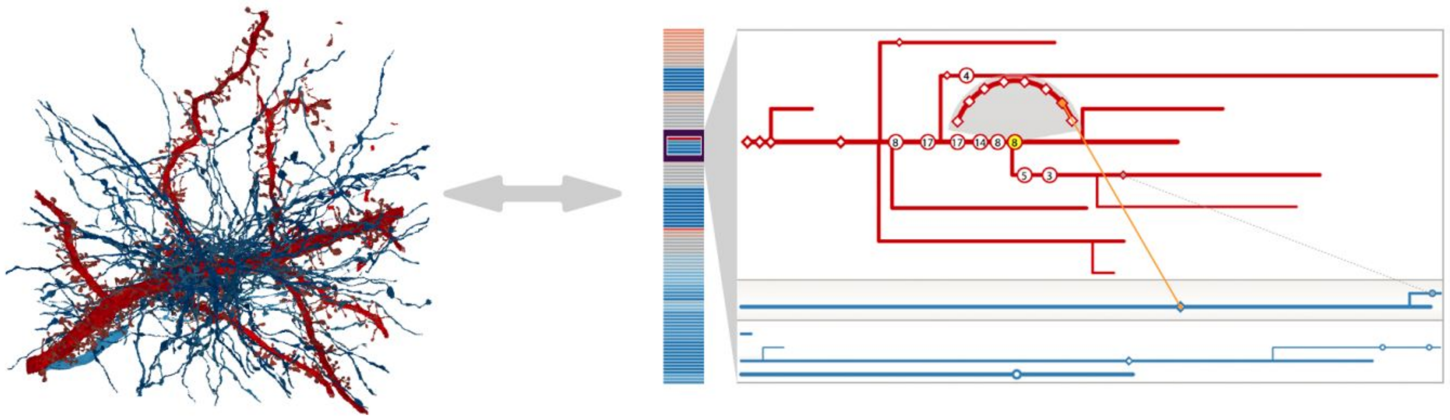

Connectomics Pipeline

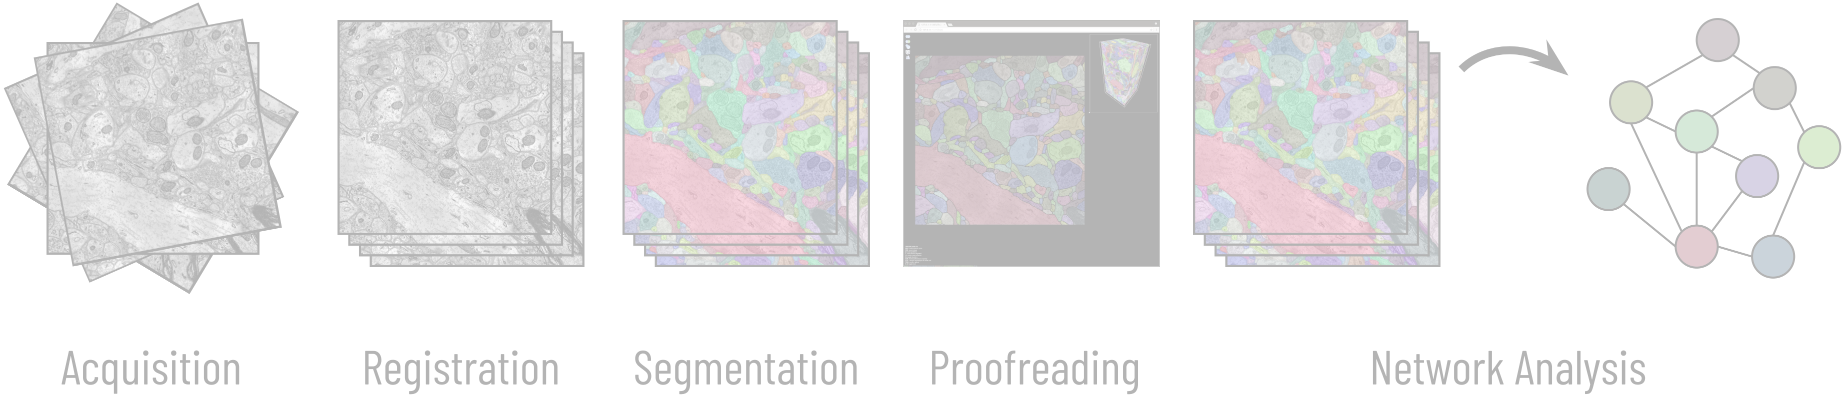

Connectomics Pipeline

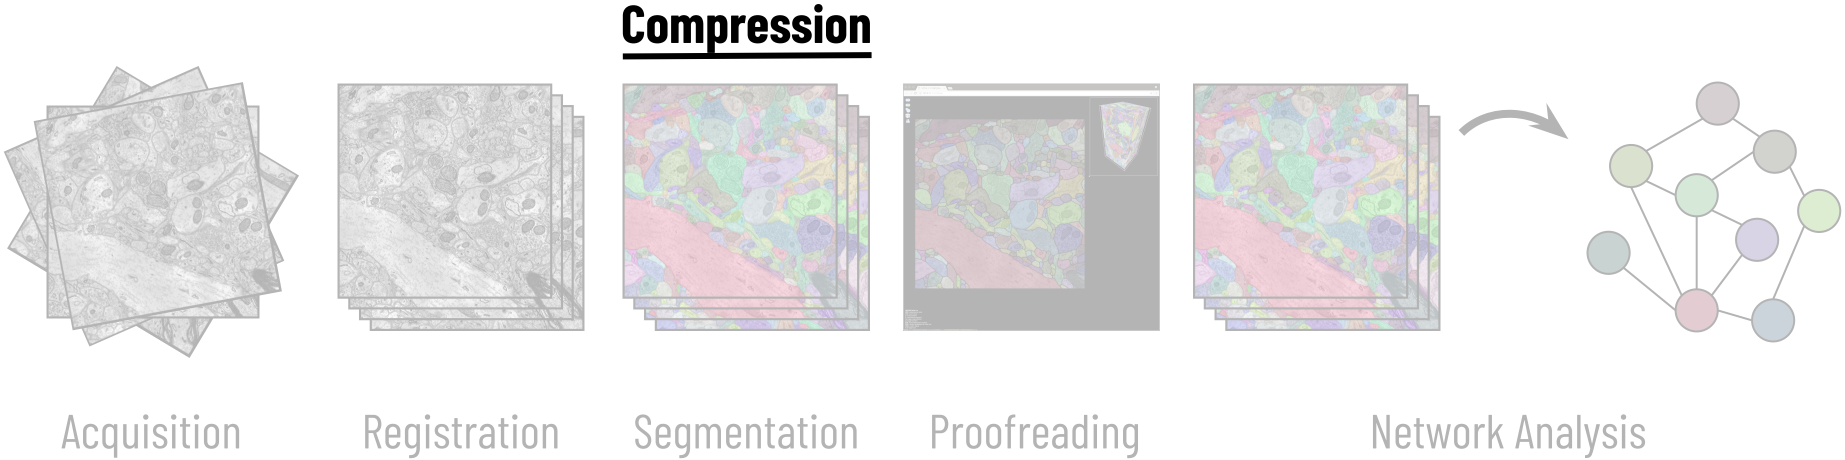

Connectomics Pipeline

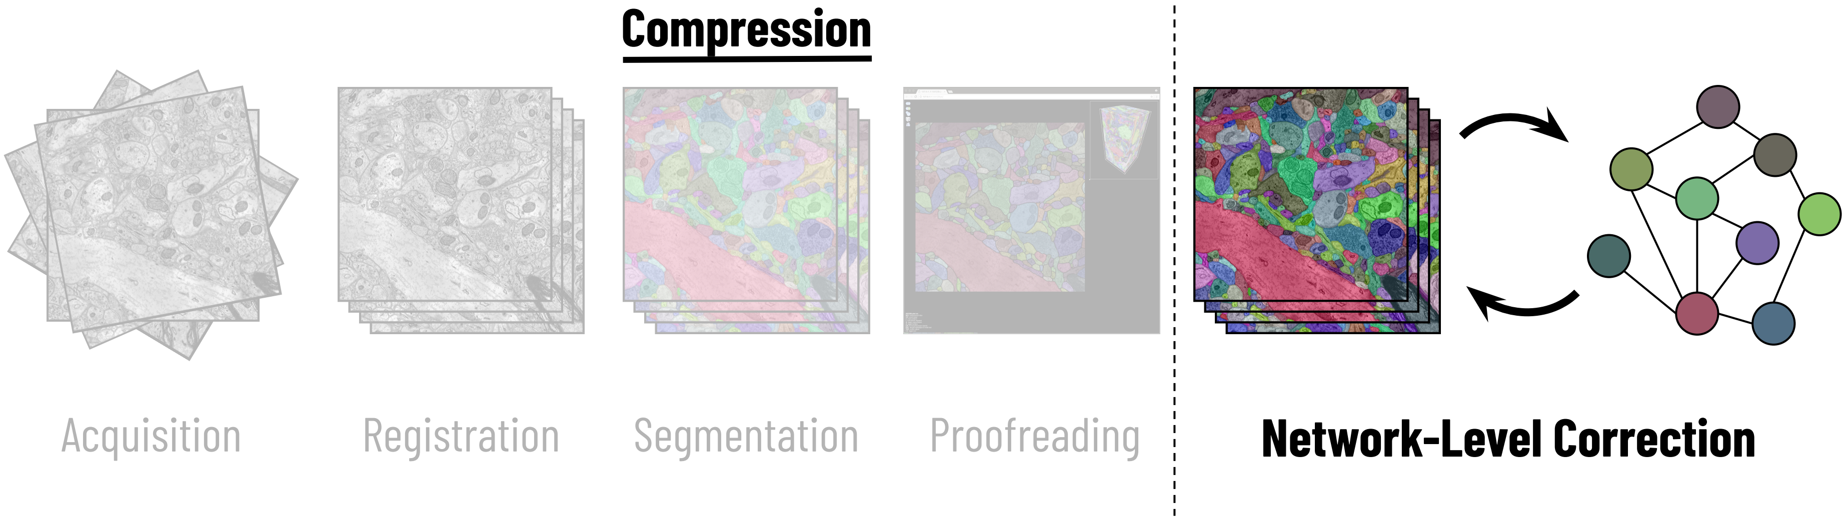

Connectomics Pipeline

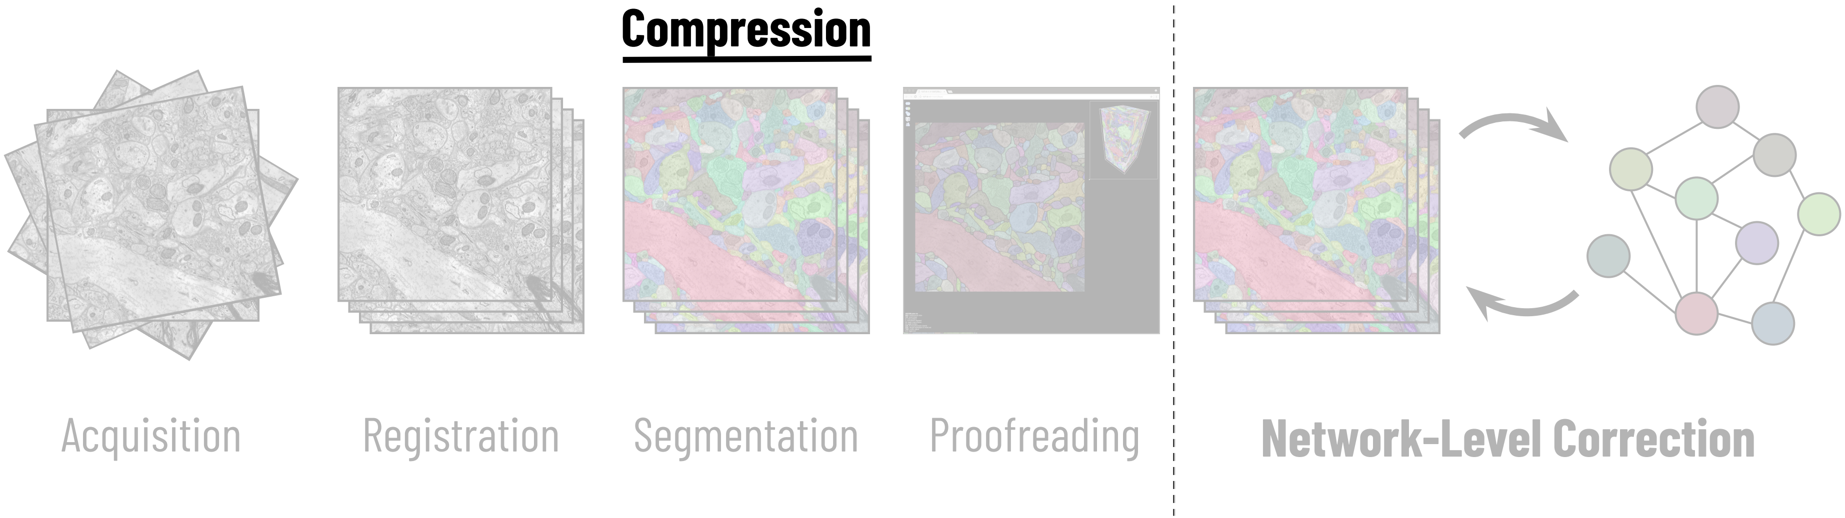

Compresso: Efficient Compression of Segmentation Data for Connectomics

Brian Matejek, Daniel Haehn, Fritz Lekschas, Michael Mitzenmacher, Hanspeter Pfister

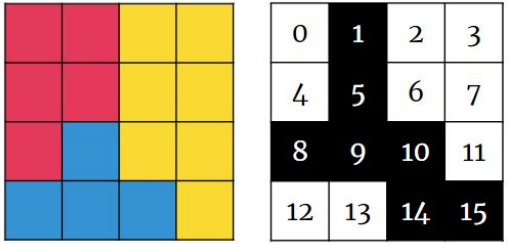

$$\sum_{i=0}^{n-1} \mathbb{I}(i)2^i$$

$$2^1 + 2^5 + 2^8 + 2^9 + 2^{10} + 2^{14} + 2^{15} = 50978$$

Published MICCAI 2017

Increasing Scales of Challenge Datasets

SNEMI

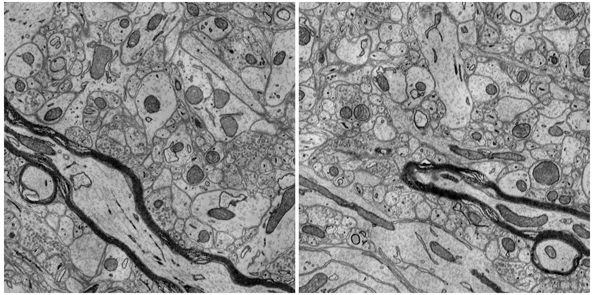

210 MB

2013

Increasing Scales of Challenge Datasets

SNEMI

CREMI

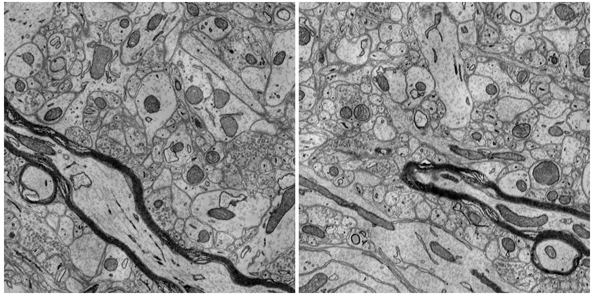

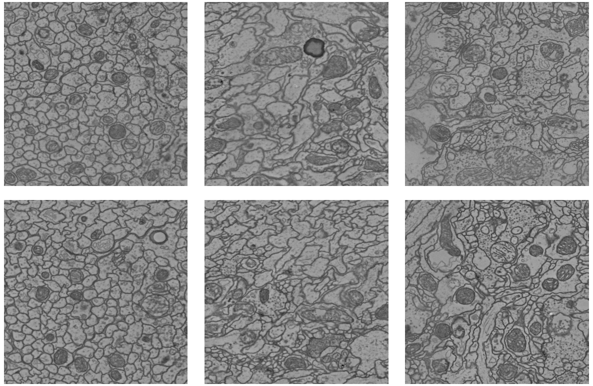

210 MB

1.2 GB

2013

2016

Increasing Scales of Challenge Datasets

SNEMI

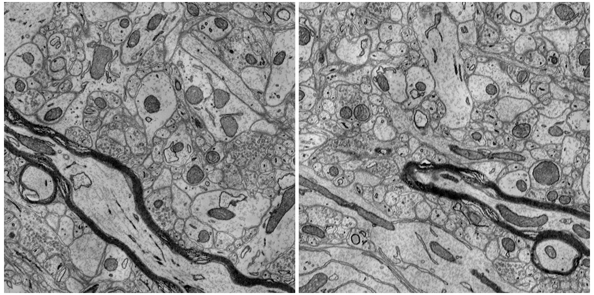

210 MB

2013

CREMI

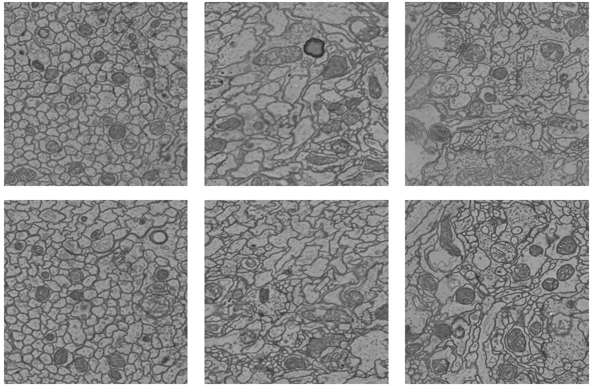

1.2 GB

2016

FIB-25

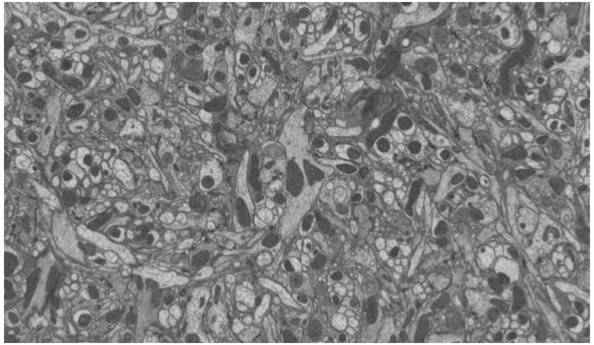

15.7 GB

2017

Connectomics Label Volumes

Large invariant regions without natural relationships between labels

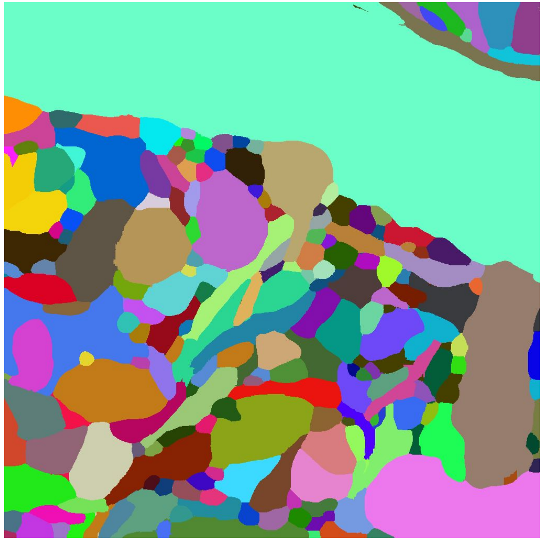

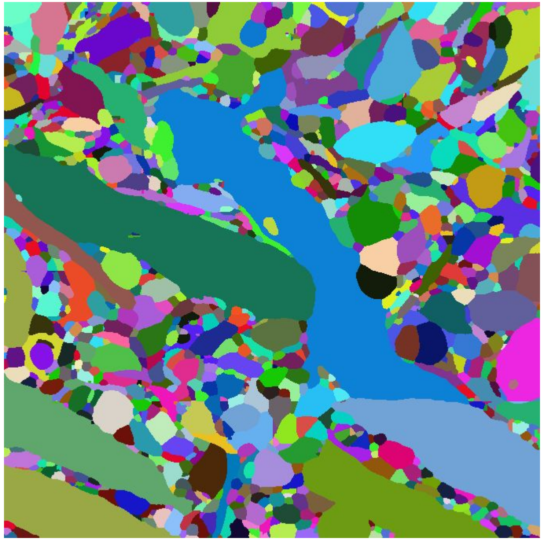

Existing Compression Schemes

- General-purpose compression schemes
  - BZ2, GZIP, LZMA, LZW, ZLIB, etc.
  - Not optimized for these unique label volumes

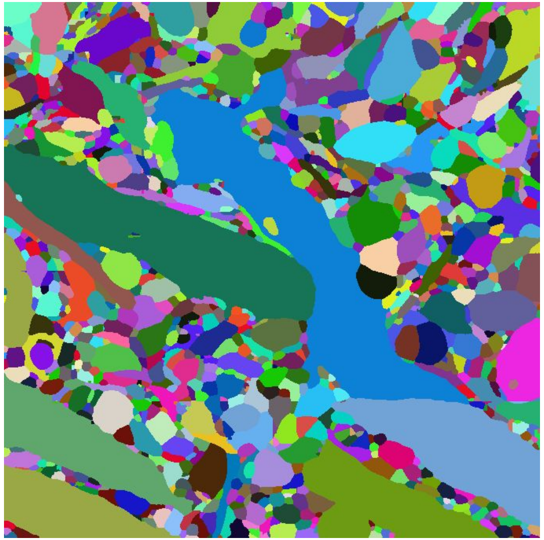

Existing Compression Schemes

- General-purpose compression schemes
  - BZ2, GZIP, LZMA, LZW, ZLIB, etc.
  - Not optimized for these unique label volumes
- Image compression schemes
  - JPEG, JPEG2000, PNG, etc.
  - Rely on frequency reduction and value prediction
    - Not useful with large invariant regions

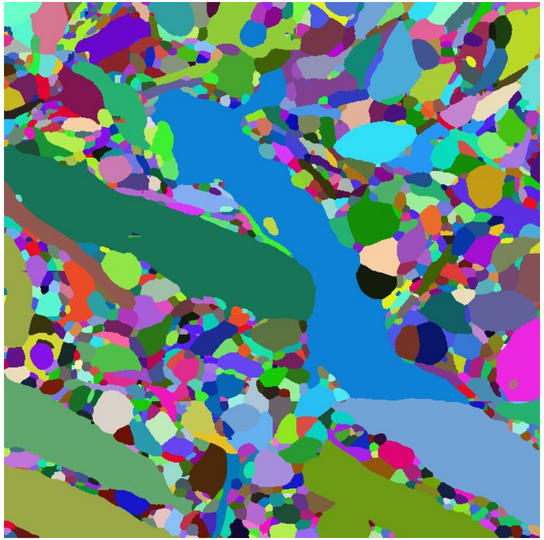

Existing Compression Schemes

- General-purpose compression schemes
  - BZ2, GZIP, LZMA, LZW, ZLIB, etc.
  - Not optimized for these unique label volumes
- Image compression schemes
  - JPEG, JPEG2000, PNG, etc.
  - Rely on frequency reduction and value prediction
    - Not useful with large invariant regions
- Video compression schemes
  - H.264, H.265, MPEG, etc.
  - Color space optimizations do not translate

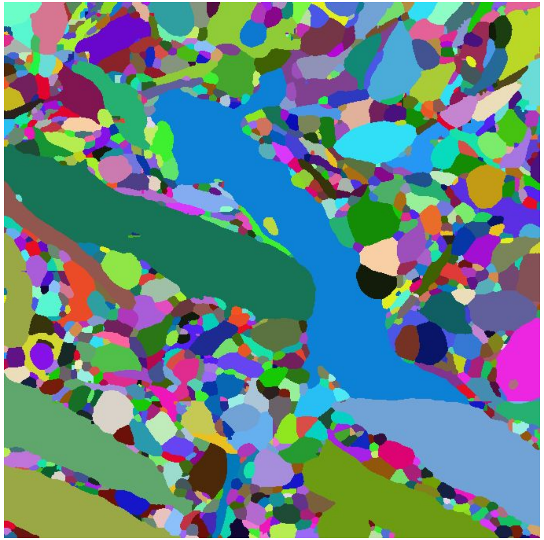

Existing Compression Schemes

- General-purpose compression schemes
  - BZ2, GZIP, LZMA, LZW, ZLIB, etc.
  - Not optimized for these unique label volumes
- Image compression schemes
  - JPEG, JPEG2000, PNG, etc.
  - Rely on frequency reduction and value prediction
    - Not useful with large invariant regions
- Video compression schemes
  - H.264, H.265, MPEG, etc.
  - Color space optimizations do not translate
- Neuroglancer compression scheme
  - Specifically designed for label volumes

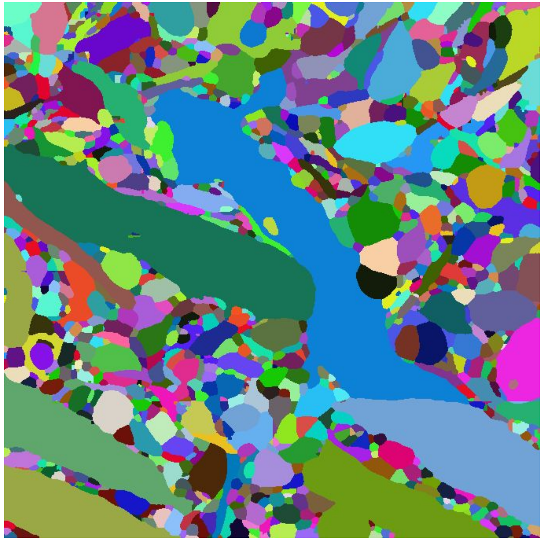

Neuroglancer

Specifically designed for label volumes

Exploits homogeneity by creating small blocks with *N* labels

- Reduces local entropy to $\log_2(N)$
- Lookup tables decode the values $[0, N)$ to the original 64-bit labels

Blocks are typically 8x8x8 voxels each

Google Open Source

https://opensource.google.com/projects/neuroglancer

Compression

Compresso Overview

Lossless compression

Compresso Overview

Lossless compression

Decouple per-segment shapes and per-pixel labels

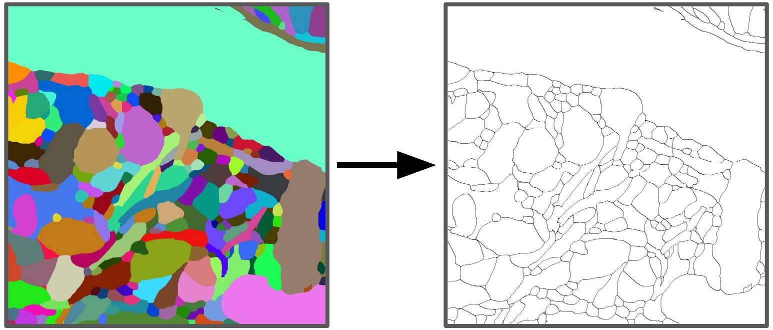

Compresso Overview

Lossless compression

Decouple per-segment shapes and per-pixel labels

Divide the volume into non-overlapping congruent 3D windows

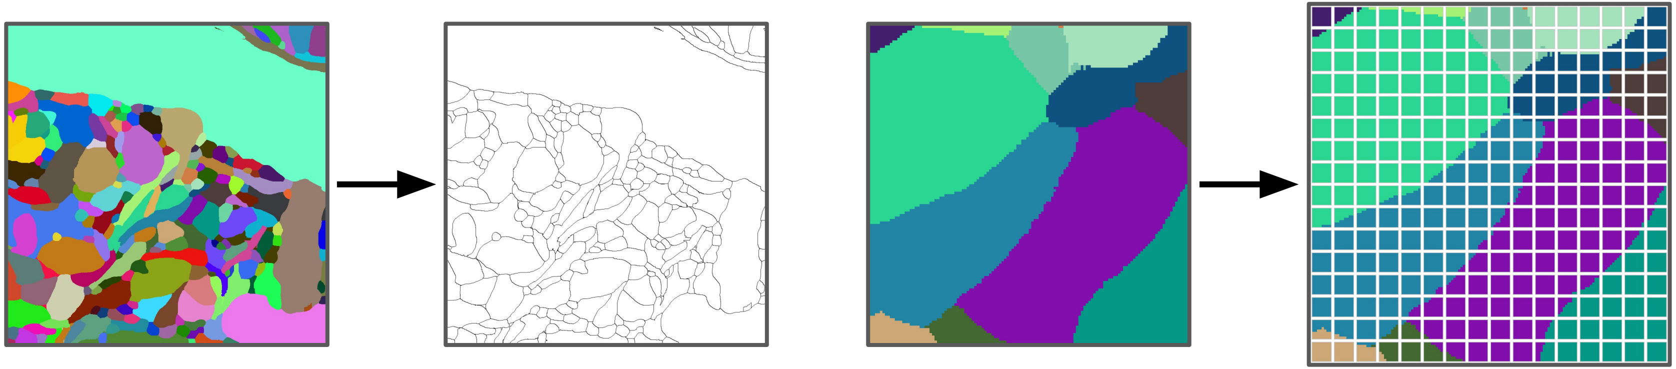

Boundary Map Generation

A pixel (x, y, z) is 1 if its neighbor $(x + 1, y, z)$ or $(x, y + 1, z)$ belongs to a different segment

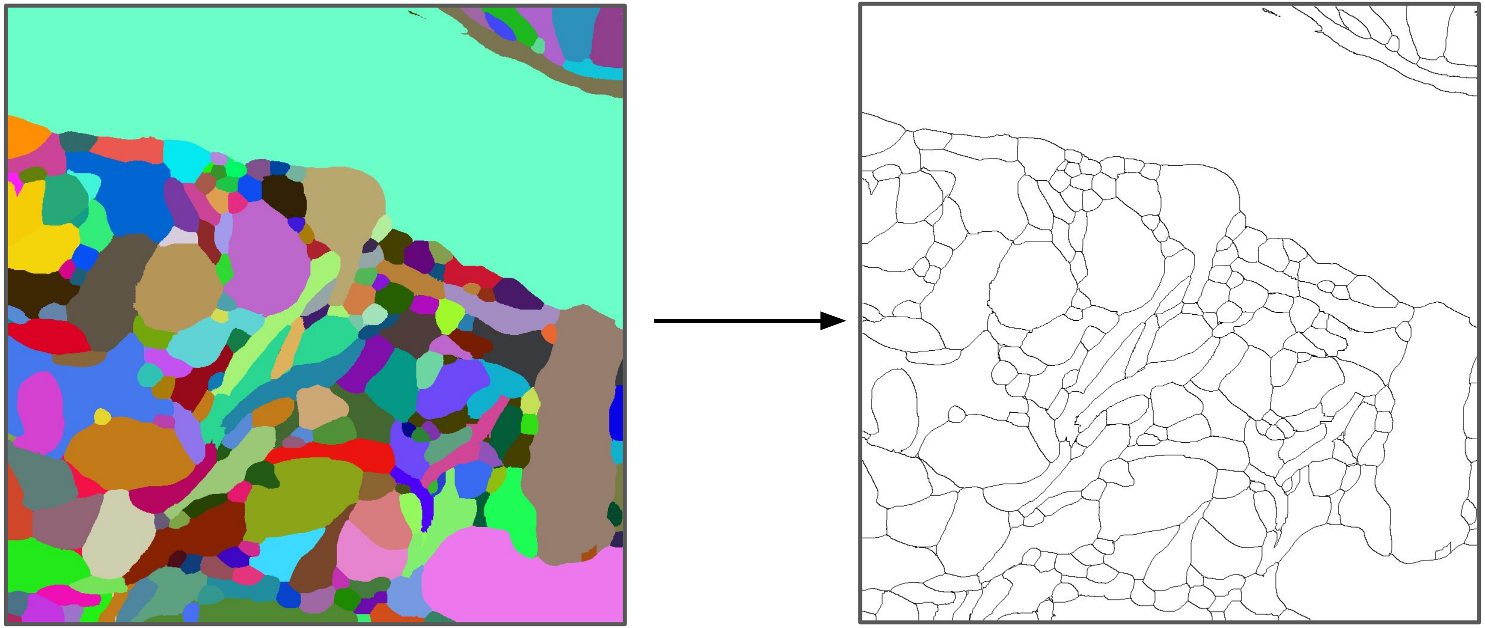

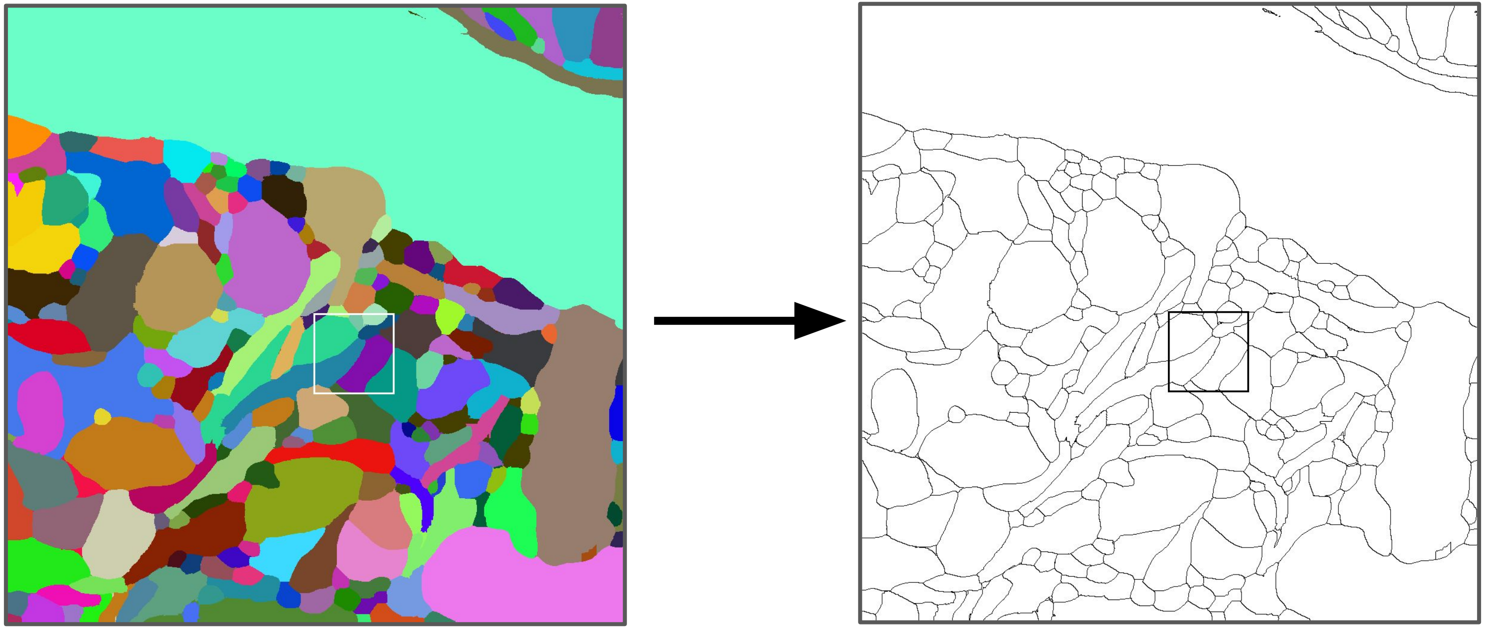

Boxed region divided into congruent windows

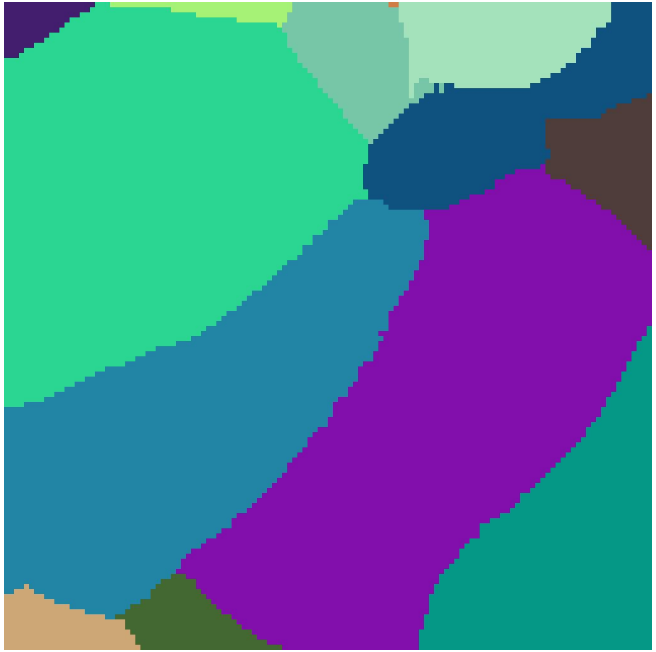

Boxed region divided into congruent windows

Each window is 8 x 8 x 1 voxels

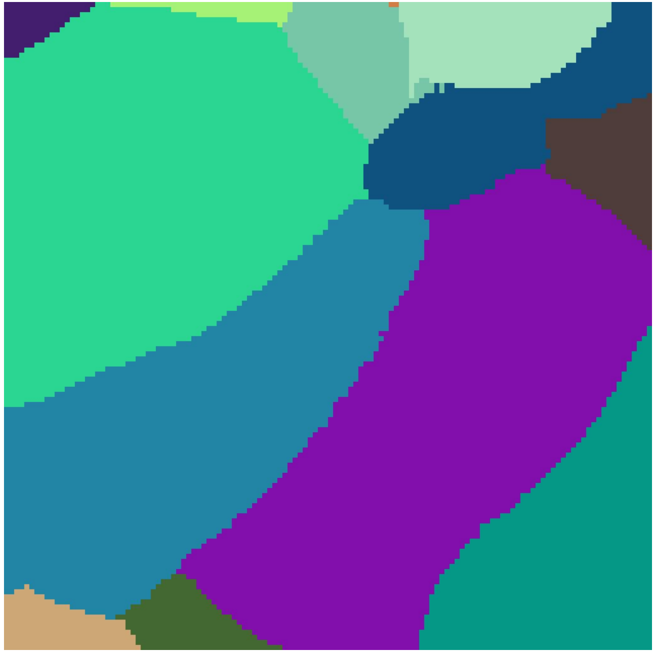

Accompanying boundary map

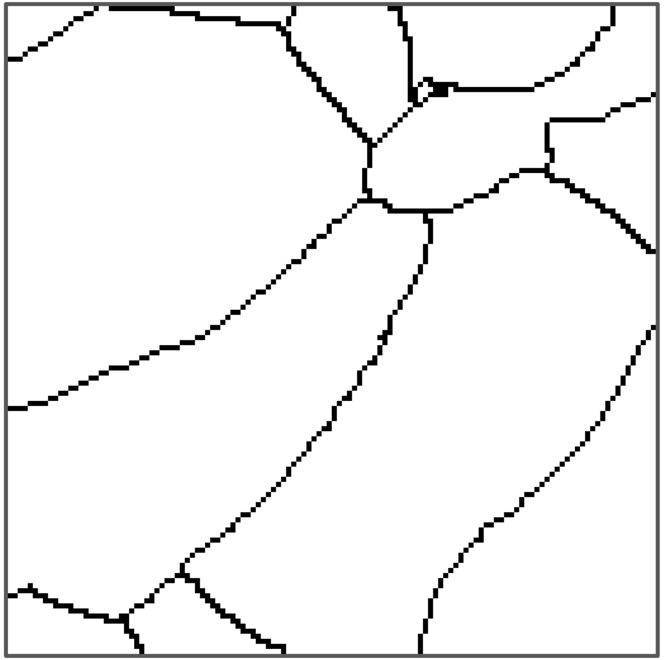

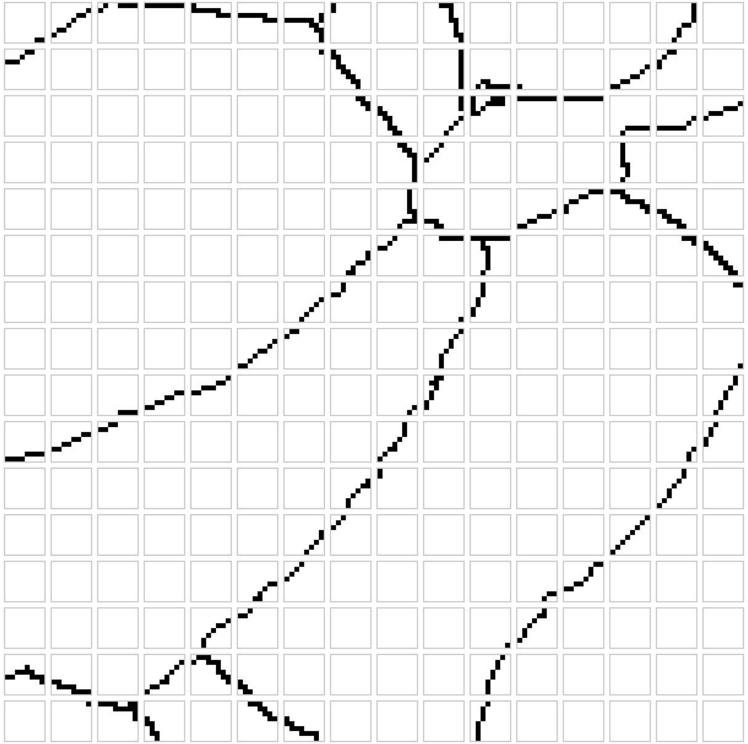

Accompanying boundary map

Goal: Store one 64-bit integer per window

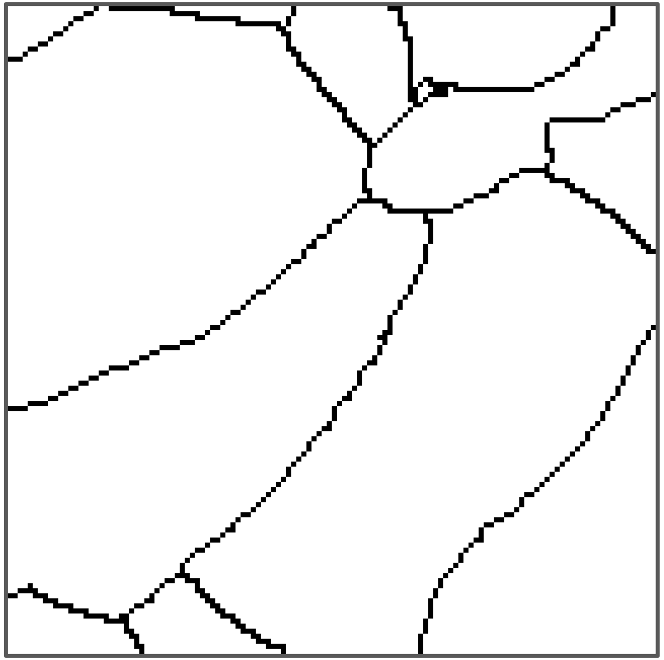

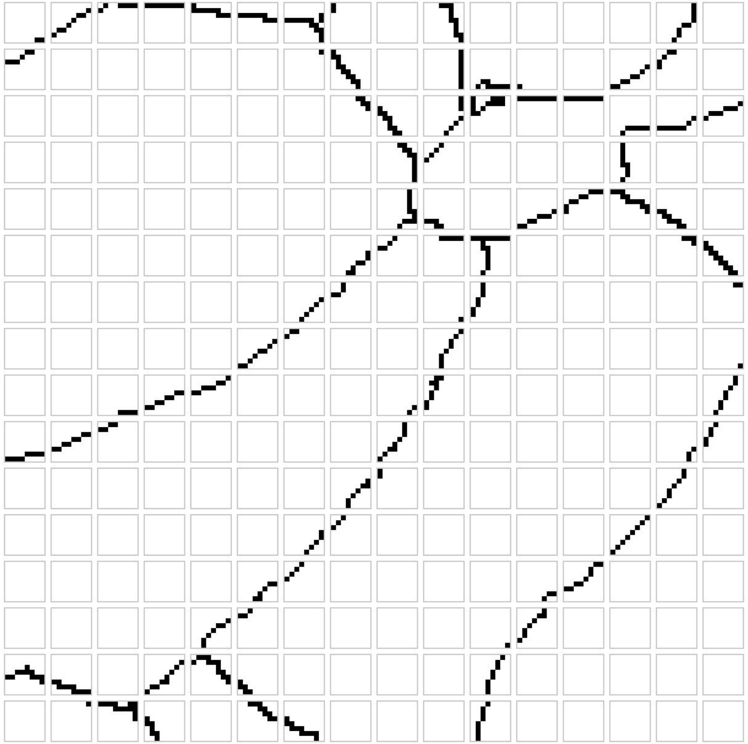

Assigning Values to Windows

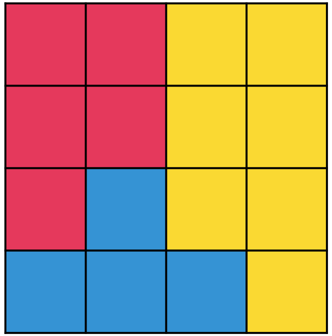

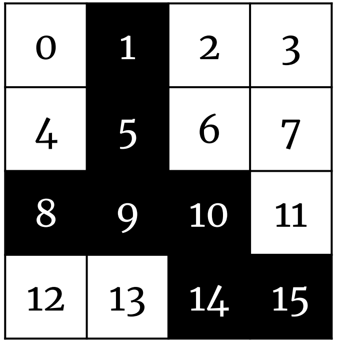

$$\sum_{i=0}^{n-1} \mathbb{I}(i) 2^i$$

$$2^1 + 2^5 + 2^8 + 2^9 + 2^{10} + 2^{14} + 2^{15} = 50978$$

Boundary Encoding

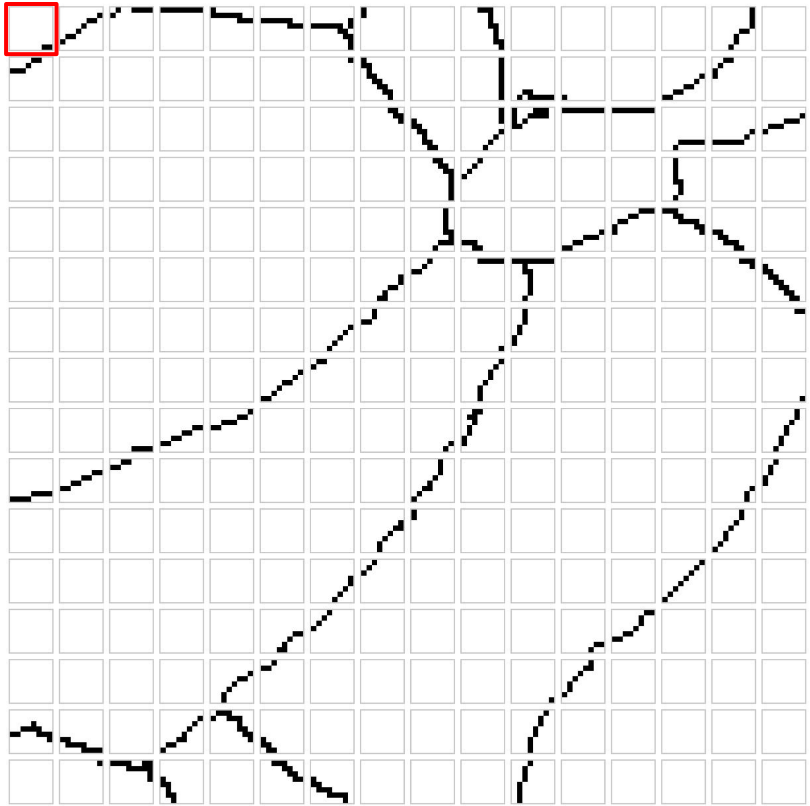

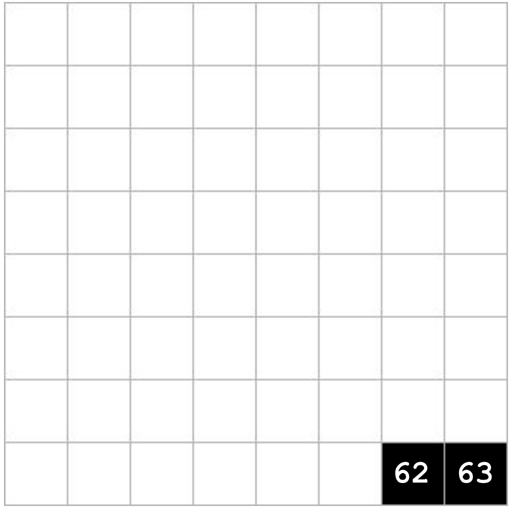

$$2^{62} + 2^{63} = 13835058055282163712$$

Boundary Encoding

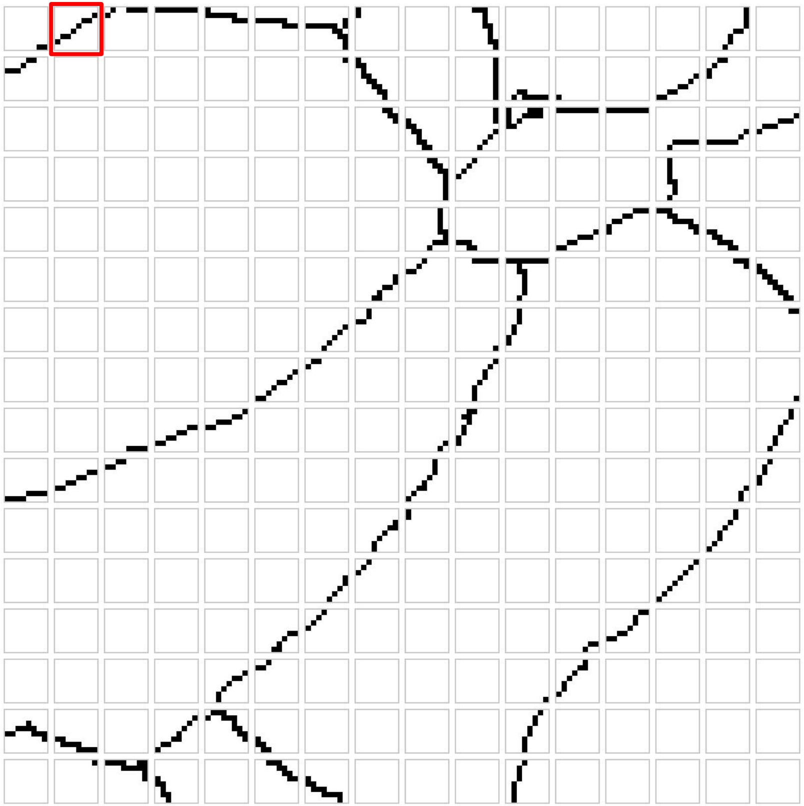

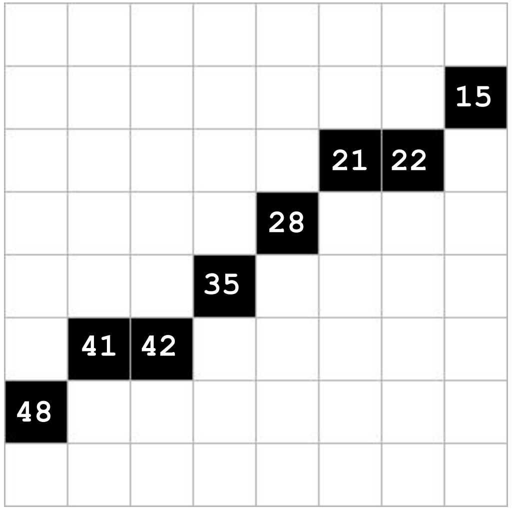

$$2^{15} + 2^{21} + 2^{22} + 2^{28} + 2^{35} + 2^{41} + 2^{42} + 2^{48} =$$
$$288106680975360$$

Boundary Encoding

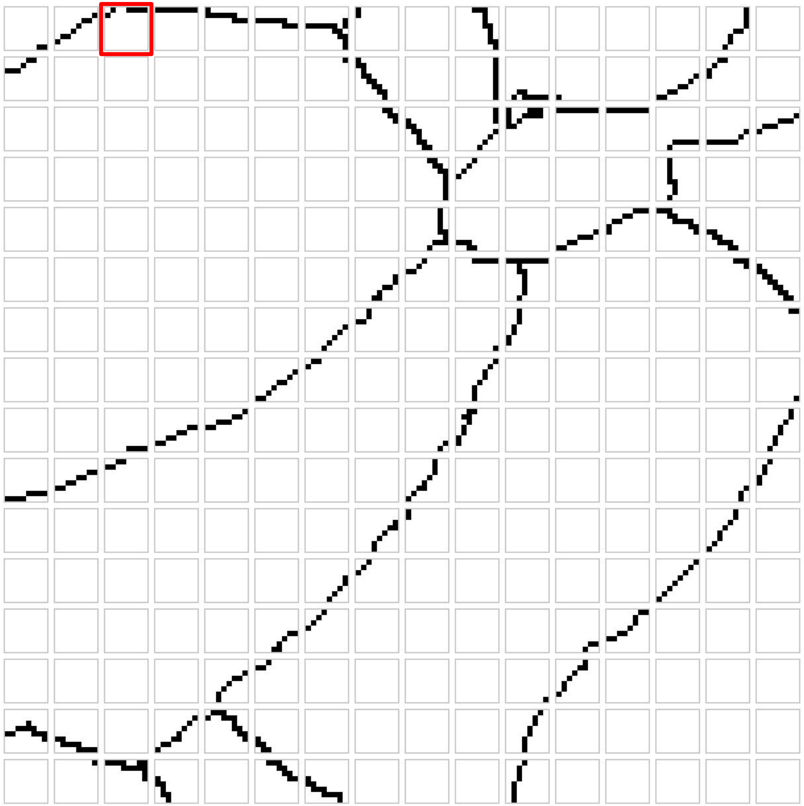

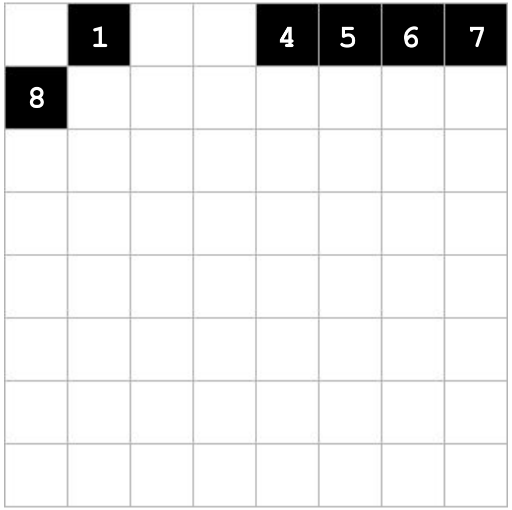

$$2^1 + 2^4 + 2^5 + 2^6 + 2^7 + 2^8 = 498$$

Boundary Encoding

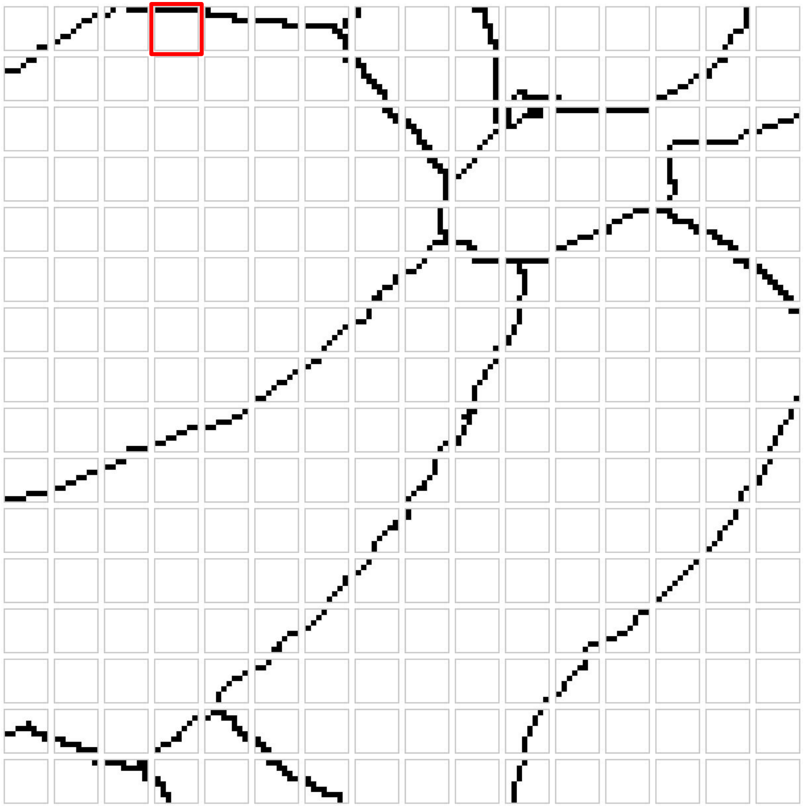

0	1	2	3	4	5	6	7

$$2^0 + 2^1 + 2^2 + 2^3 + 2^4 + 2^5 + 2^6 + 2^7 = 255$$

Boundary Encoding

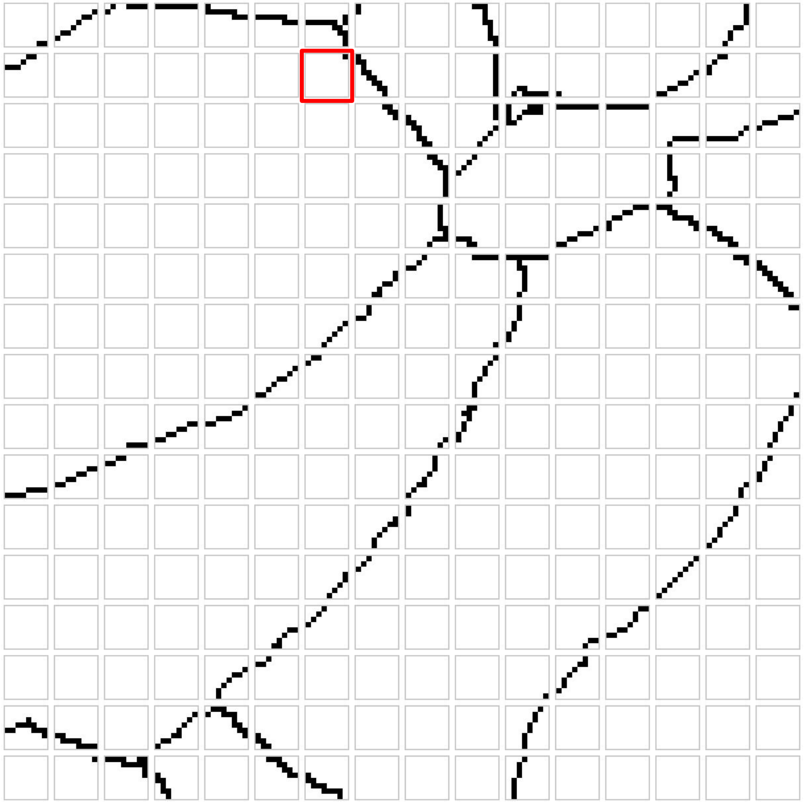

13835058055282163712
288106680975360
498
255
14696193
3762225152
9259612355635970048
257
0
9277485877618024504
0
0
0
0
580982358589603968
0
460848
0
0
0
0
0
128
.
.
.

Window Repetition

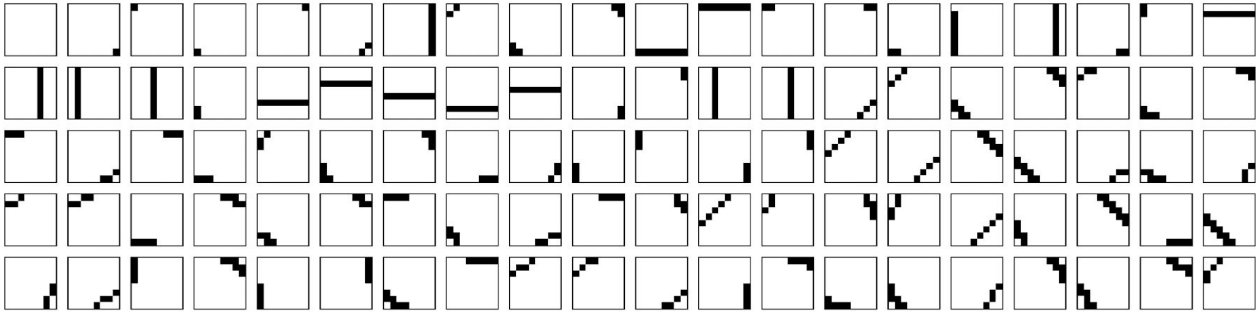

These 100 windows account for over 82% of all windows on a representative connectomics dataset

Window Repetition

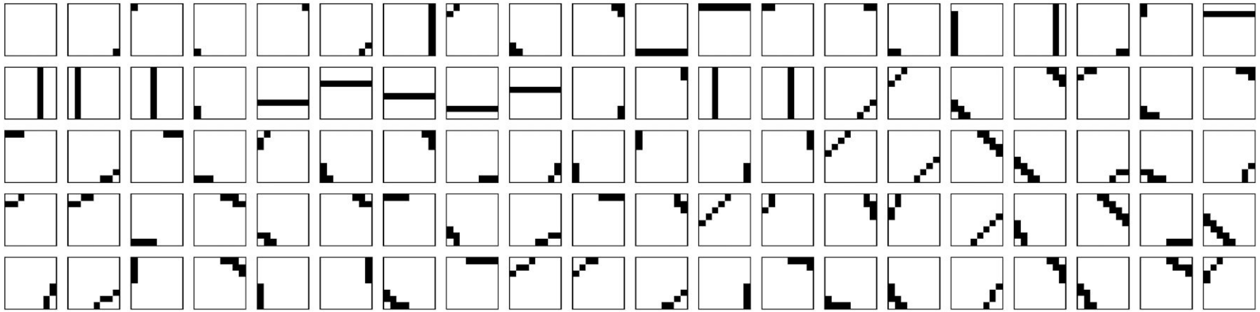

These 100 windows account for over 82% of all windows on a representative connectomics dataset

Typically there are only 100,000 unique windows in a given label volume

Window Repetition

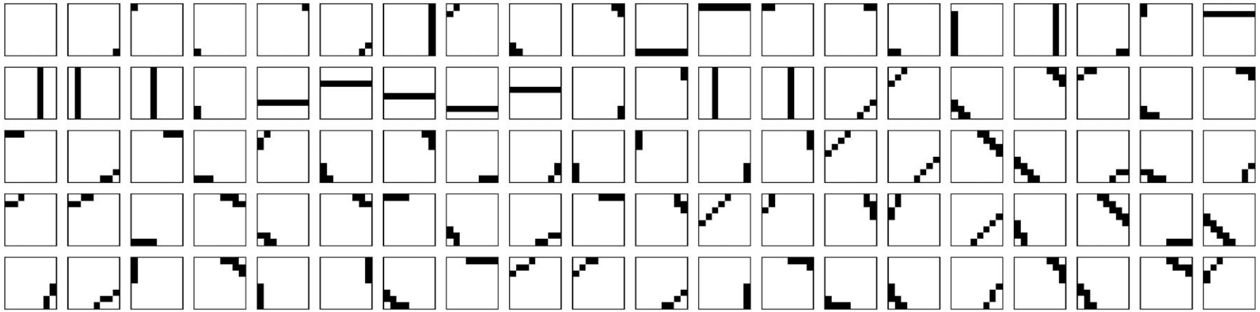

These 100 windows account for over 82% of all windows on a representative connectomics dataset

Typically there are only 100,000 unique windows in a given label volume

Map window values to this smaller subset to use 3 bytes per window

Compressing Per-Pixel Labels

Goal: Store one 64-bit label per component

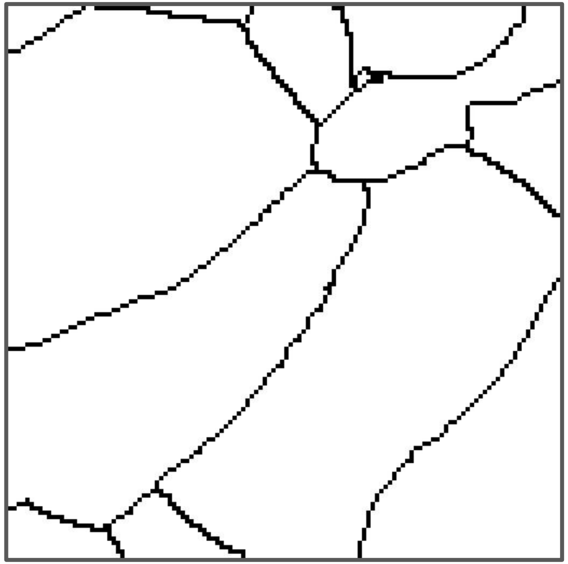

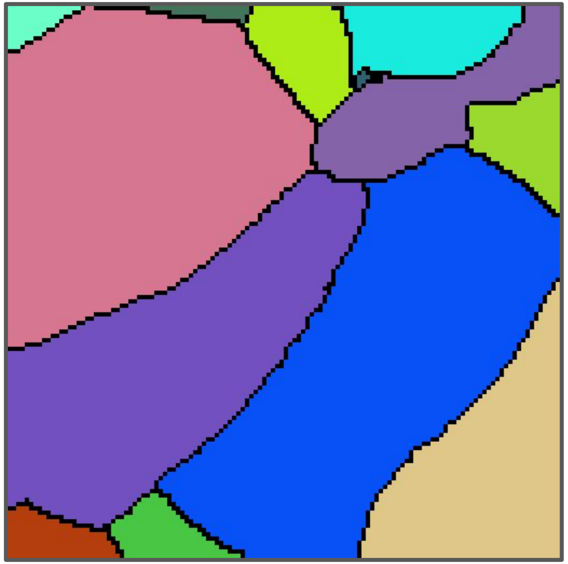

Compressing Per-Pixel Labels

Goal: Store one 64-bit label per component

Solution: Identify continuous regions using a connected components algorithm

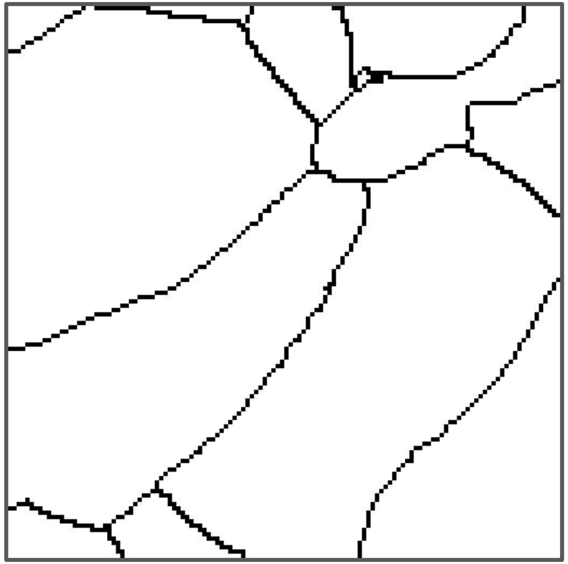

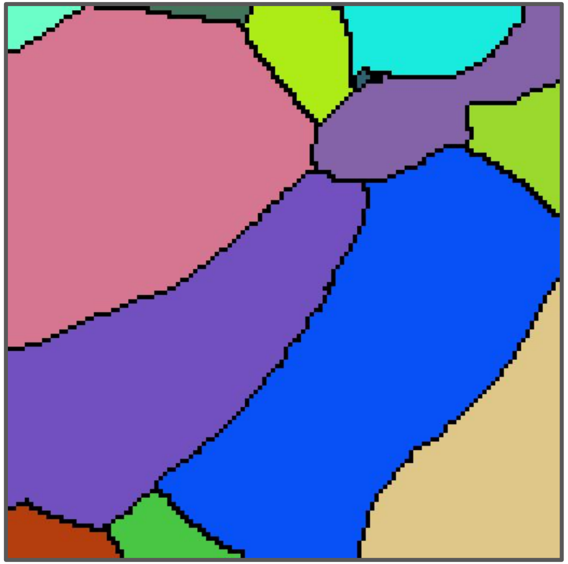

Per-Pixel Label Encoding

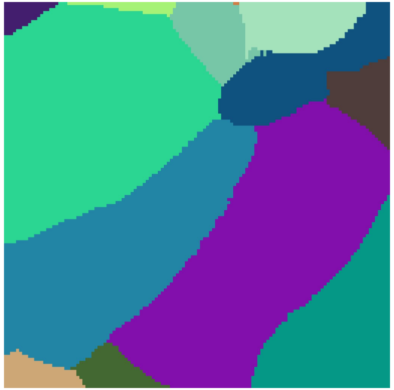

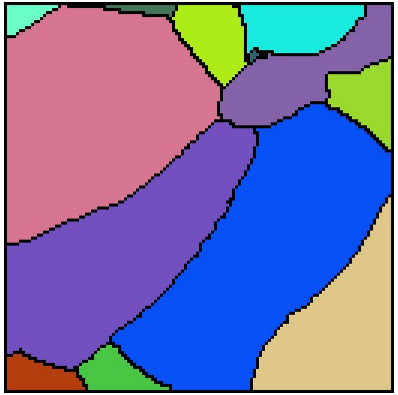

Per-Pixel Label Encoding

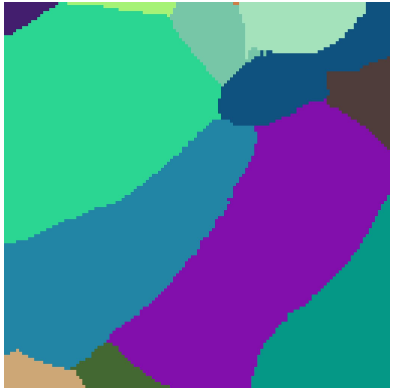

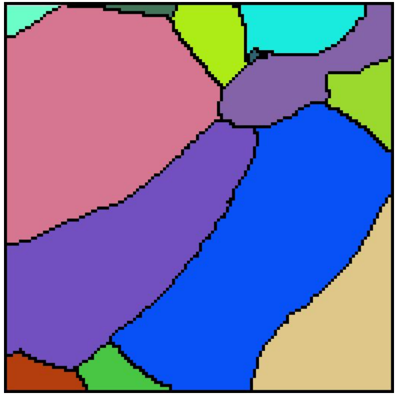

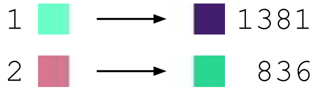

Per-Pixel Label Encoding

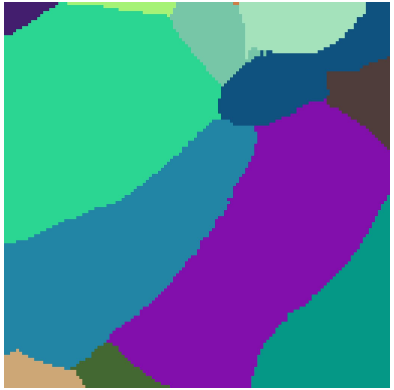

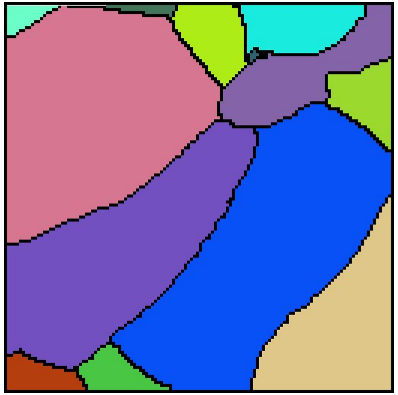

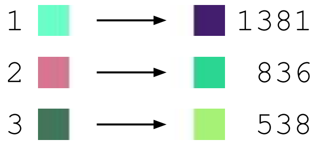

Per-Pixel Label Encoding

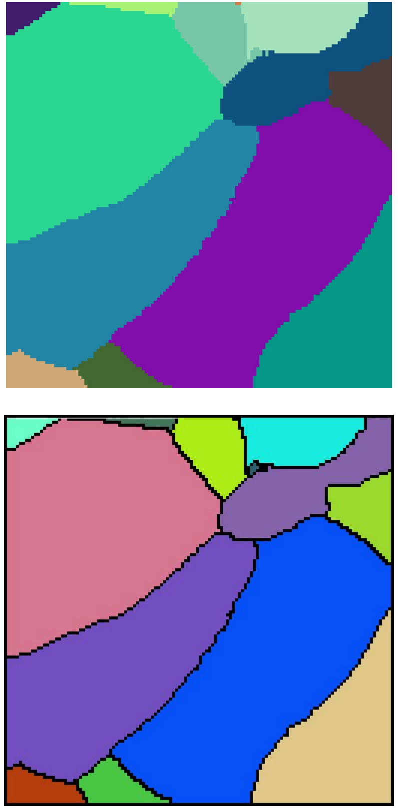

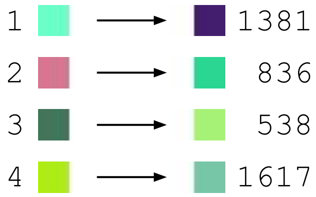

Per-Pixel Label Encoding

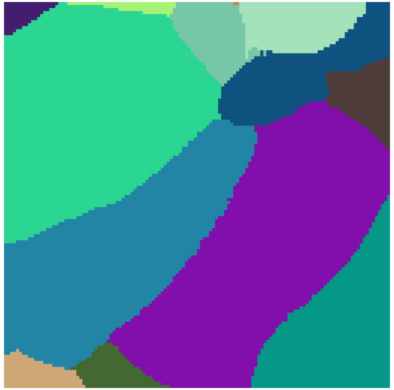

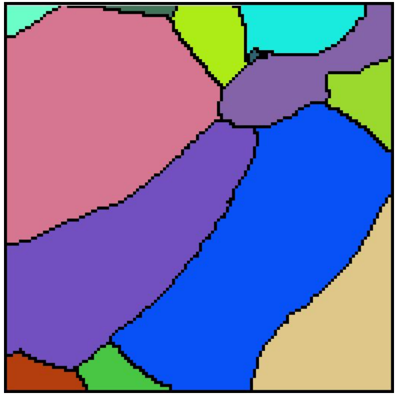

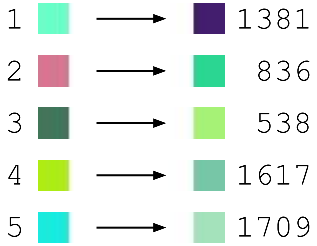

Per-Pixel Label Encoding

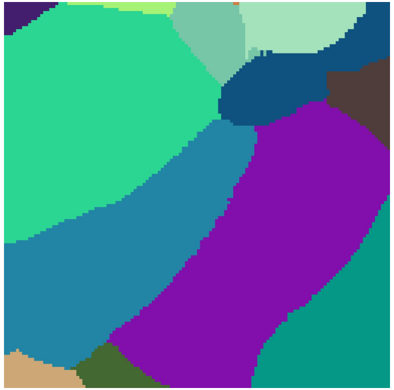

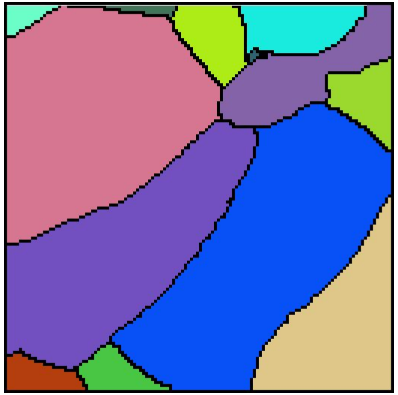

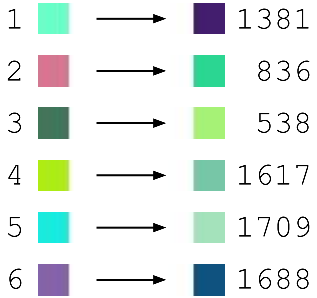

Per-Pixel Label Encoding

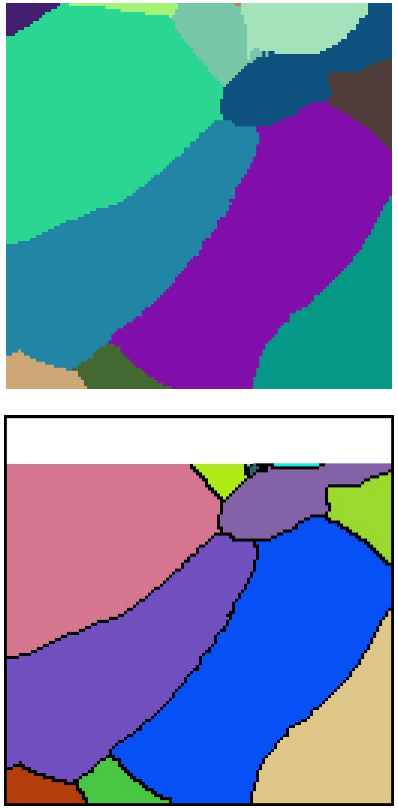

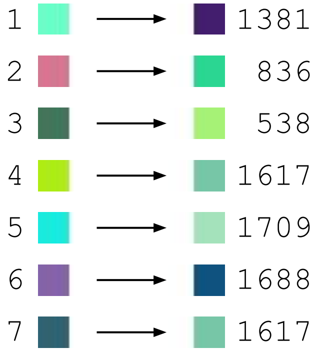

Per-Pixel Label Encoding

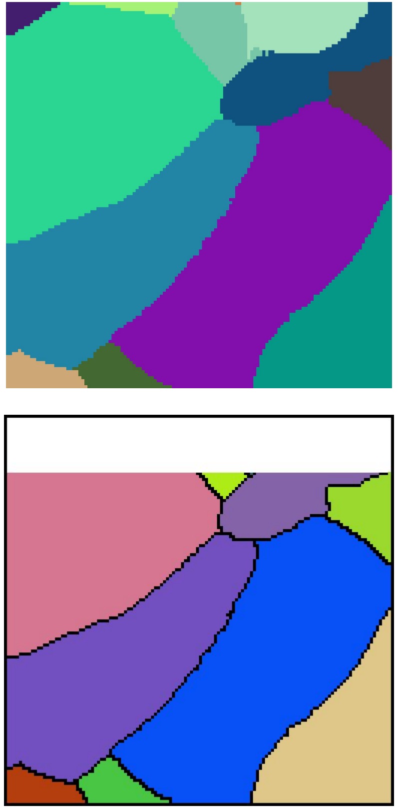

Per-Pixel Label Encoding

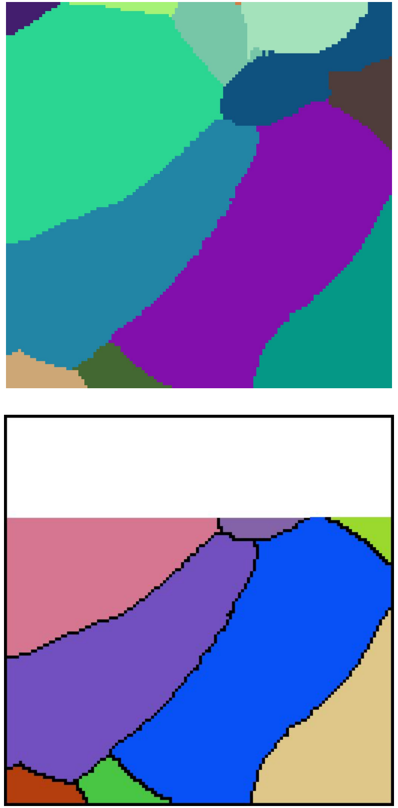

Per-Pixel Label Encoding

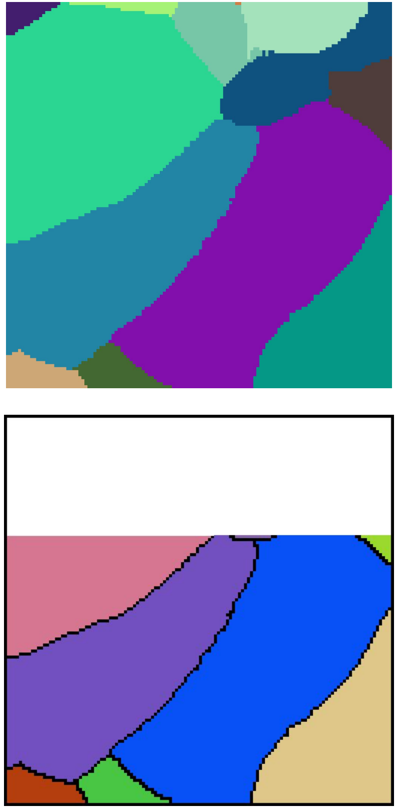

1	1381
2	836
3	538
4	1617
5	1709
6	1688
7	1617
8	1619
9	1020
10	827

Per-Pixel Label Encoding

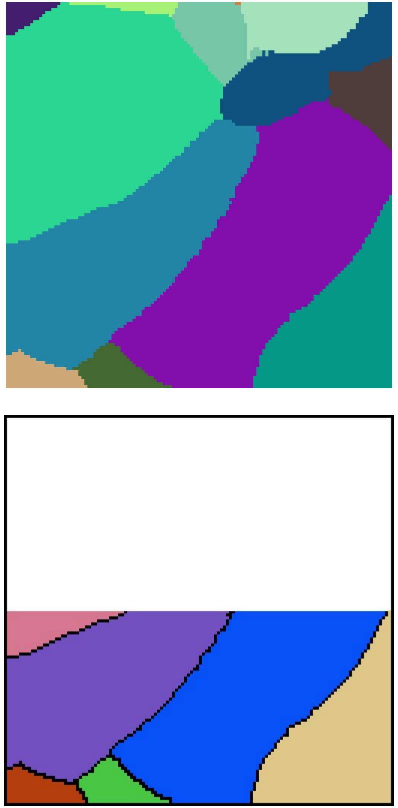

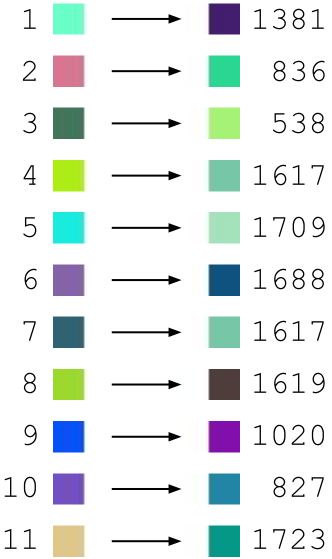

Per-Pixel Label Encoding

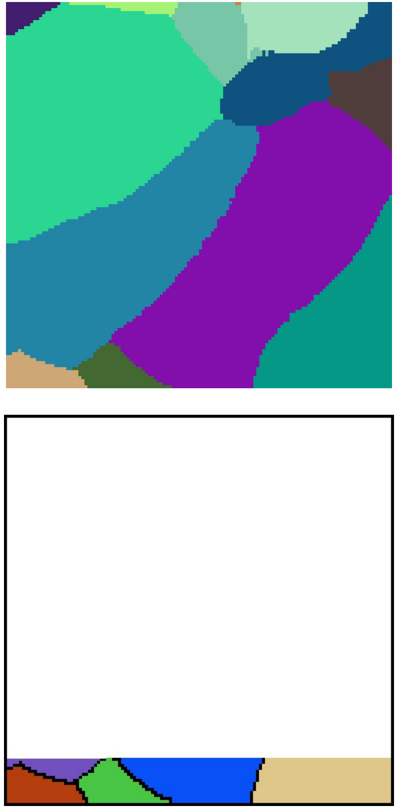

1	1381
2	836
3	538
4	1617
5	1709
6	1688
7	1617
8	1619
9	1020
10	827
11	1723
12	1246

Per-Pixel Label Encoding

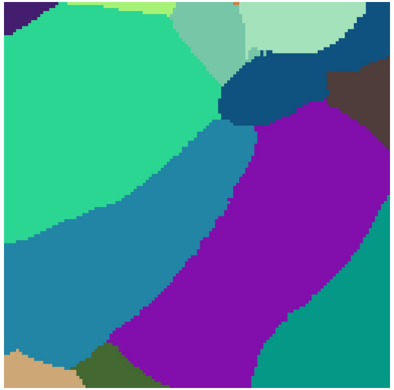

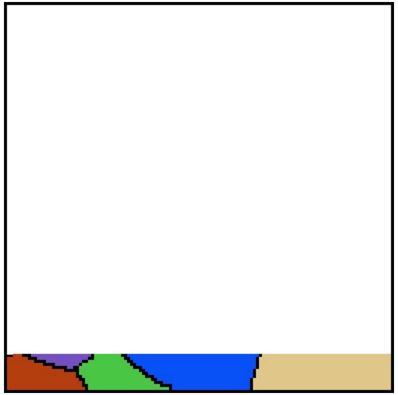

1	1381
2	836
3	538
4	1617
5	1709
6	1688
7	1617
8	1619
9	1020
10	827
11	1723
12	1246
13	1258

Per-Pixel Label Encoding

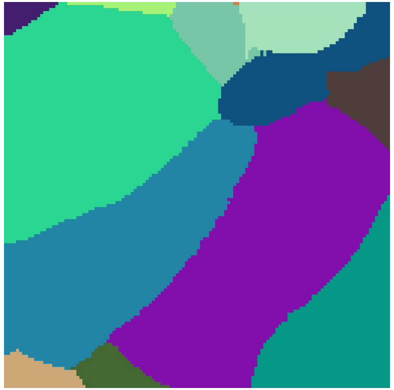

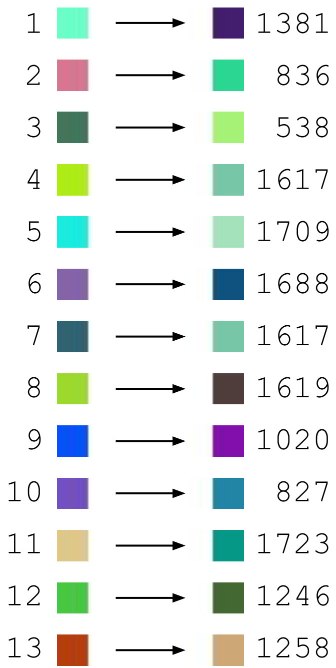

Indeterminate Locations

So far, we assumed the boundary map and connected component mapping is enough for decompression

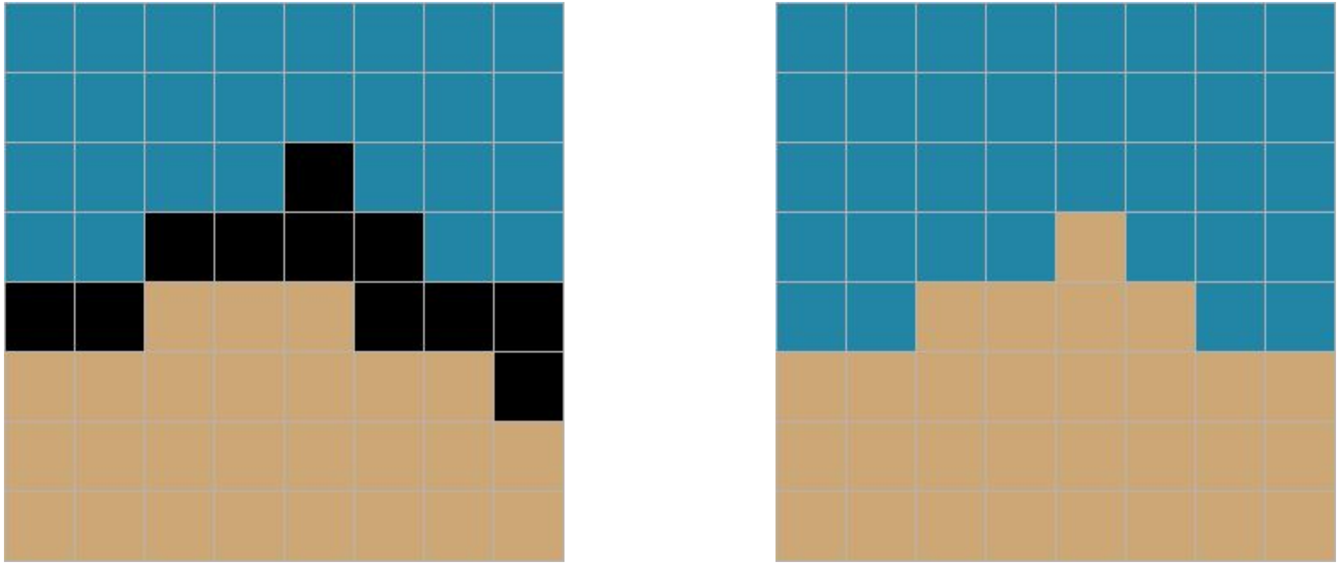

Indeterminate Locations

So far, we assumed the boundary map and connected component mapping is enough for decompression

One additional corner case to consider:

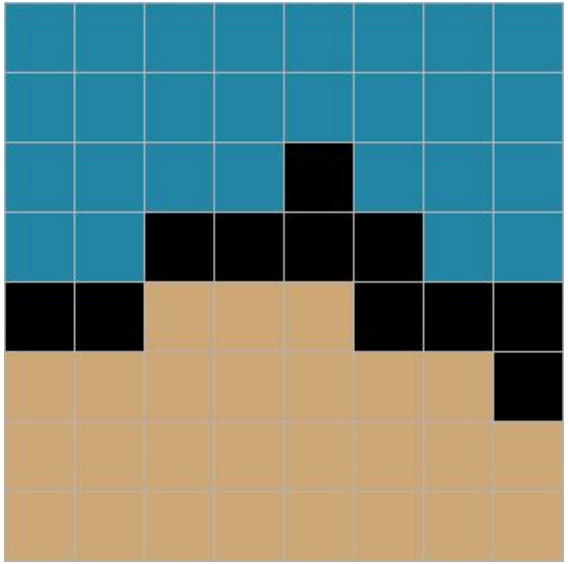

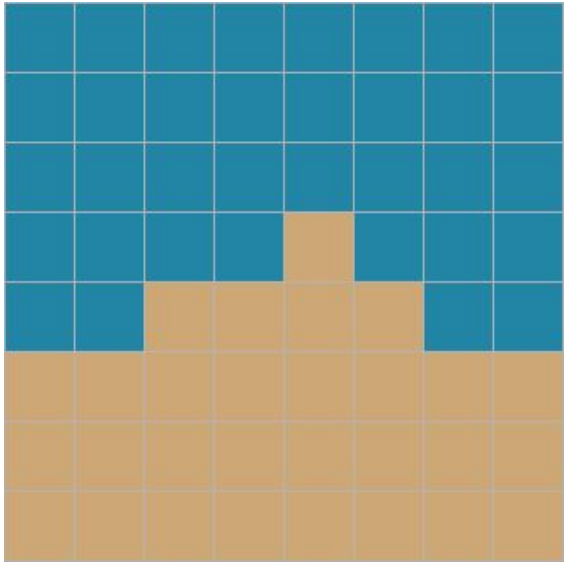

Indeterminate Locations

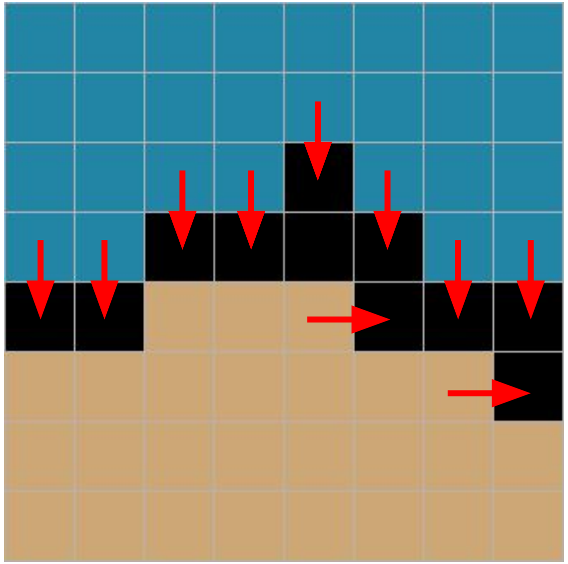

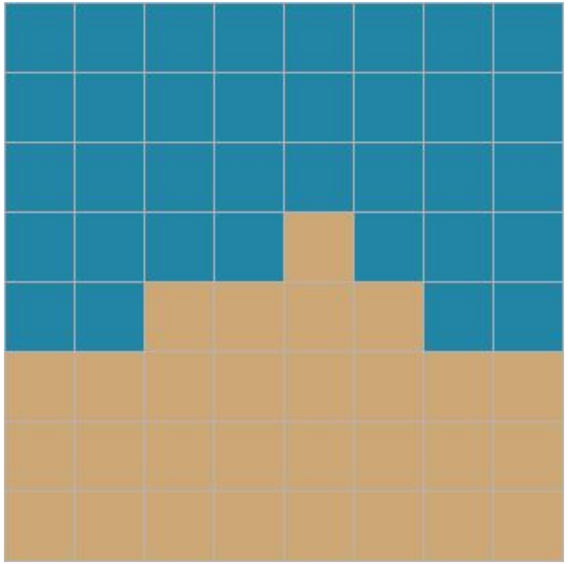

Indeterminate Locations

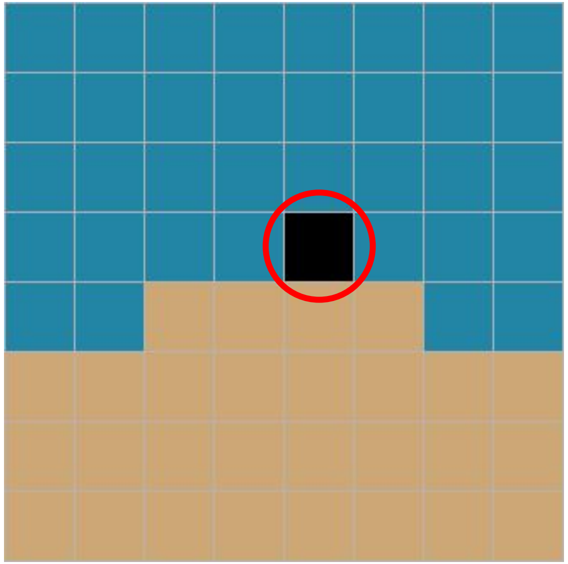

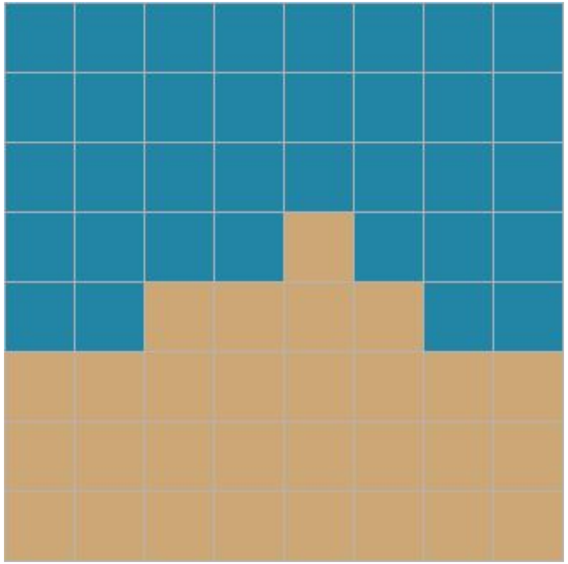

Indeterminate Locations

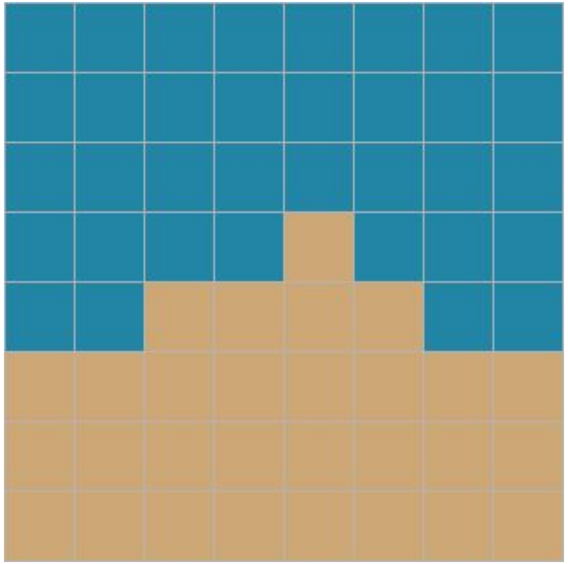

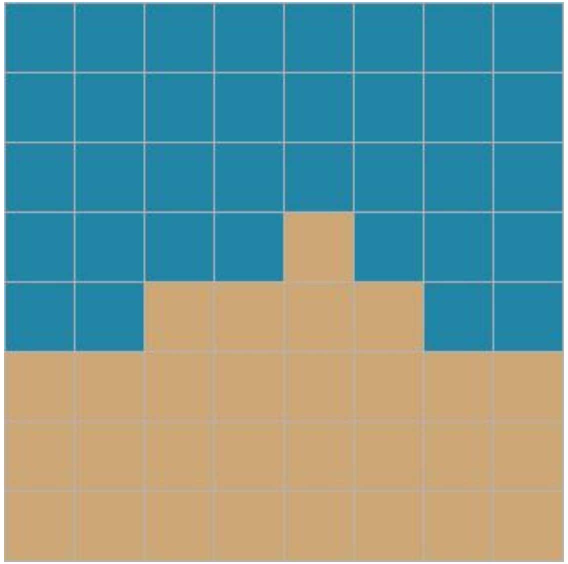

Decompression

Decompressing Boundary Map

```
13835058055282163712
     288106680975360
                 498
                 255
            14696193
          3762225152
 9259612355635970048
                 257
                   0
 9277485877618024504
                   0
                   0
                   0
                   0
  580982358589603968
                   0
              460848
                   0
                   0
                   0
                   0
                   0
                 128
                   .
                   .
                   .
```

Decompressing Boundary Map

13835058055282163712	1100000000000000000000000000000000000000000000000000000000000000
288106680975360	0000000000000001000001100000100000010000011000001000000000000000
498	0000000000000000000000000000000000000000000000000000000111110010
255	0000000000000000000000000000000000000000000000000000000011111111
14696193	0000000000000000000000000000000000000000111000000011111100000001
3762225152	0000000000000000000000000000000011100000001111110000000000000000
9259612355635970048	1000000010000000110000000110000010111111100000000000000000000000
257	0000000000000000000000000000000000000000000000000000000100000001
0	0000000000000000000000000000000000000000000000000000000000000000
9277485877618024504	1000000011000000010000000100000001100000001000000010000000111000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
580982358589603968	0000100000010000000100000110000010000000100000001000000010000000
0	0000000000000000000000000000000000000000000000000000000000000000
460848	0000000000000000000000000000000000000000000001110000100000110000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
128	0000000000000000000000000000000000000000000000000000000010000000

.
.
.

Decompressing Boundary Map

13835058055282163712	1100000000000000000000000000000000000000000000000000000000000000
288106680975360	0000000000000001000001100000100000010000011000001000000000000000
498	**0000000000000000000000000000000000000000000000000000000111110010**
255	0000000000000000000000000000000000000000000000000000000011111111
14696193	0000000000000000000000000000000000000000111000000011111100000001
3762225152	0000000000000000000000000000000011100000001111110000000000000000
9259612355635970048	1000000010000000110000000110000010111111100000000000000000000000
257	0000000000000000000000000000000000000000000000000000000100000001
0	0000000000000000000000000000000000000000000000000000000000000000
9277485877618024504	1000000011000000010000000100000001100000001000000010000000111000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
580982358589603968	0000100000010000000100000110000010000000100000001000000010000000
0	0000000000000000000000000000000000000000000000000000000000000000
460848	0000000000000000000000000000000000000000000001110000100000110000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
0	0000000000000000000000000000000000000000000000000000000000000000
128	0000000000000000000000000000000000000000000000000000000010000000

.
.
.

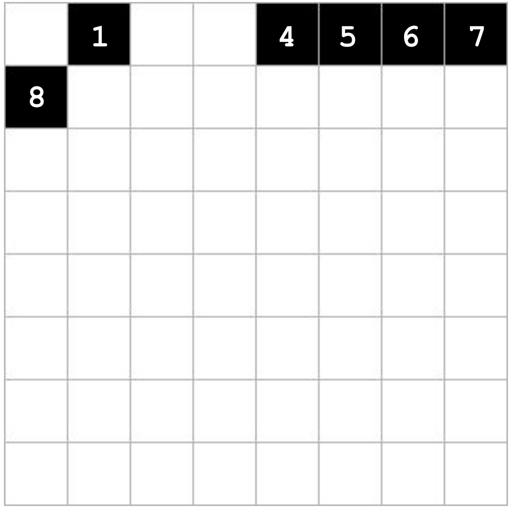

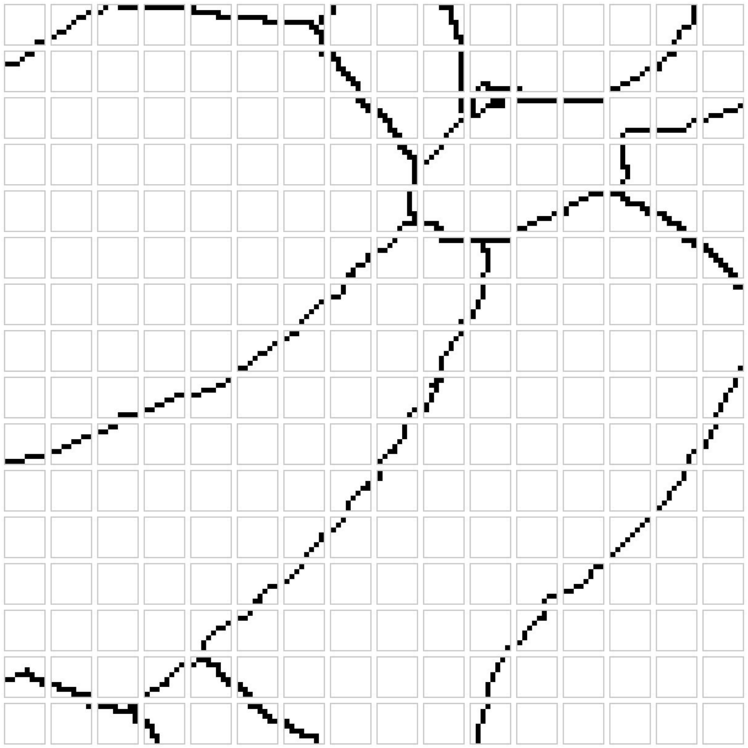

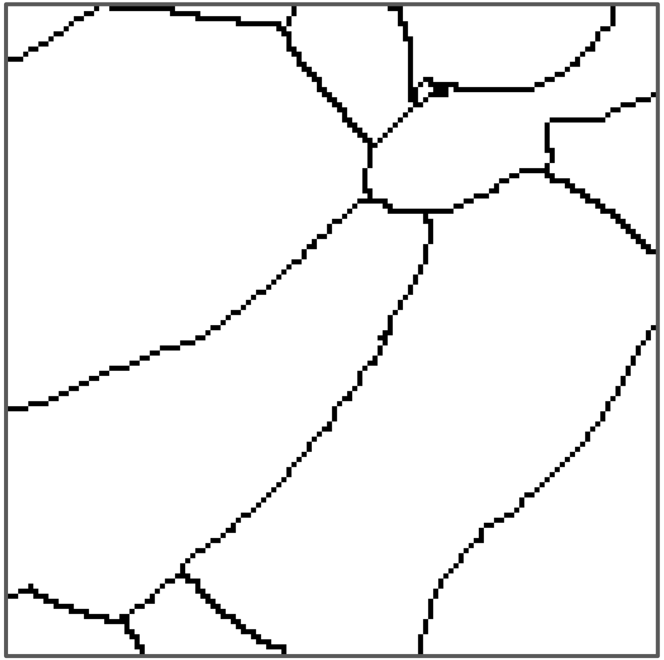

Decompressing Per-Pixel Labels

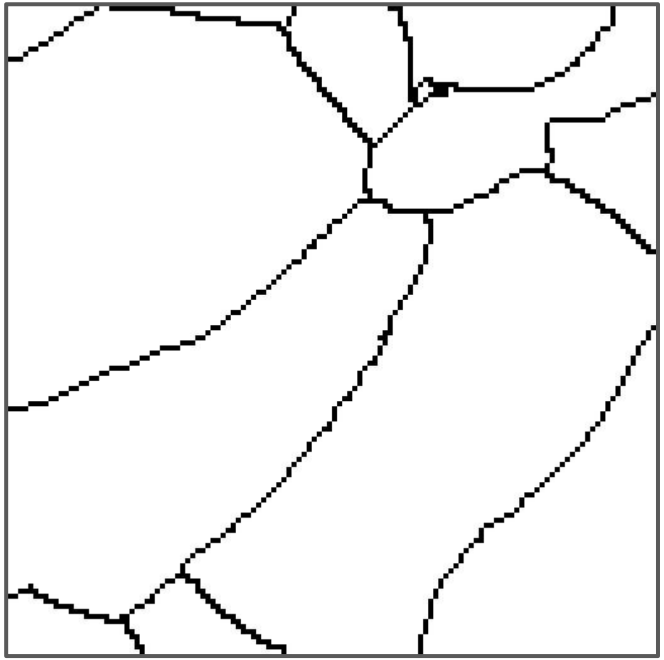

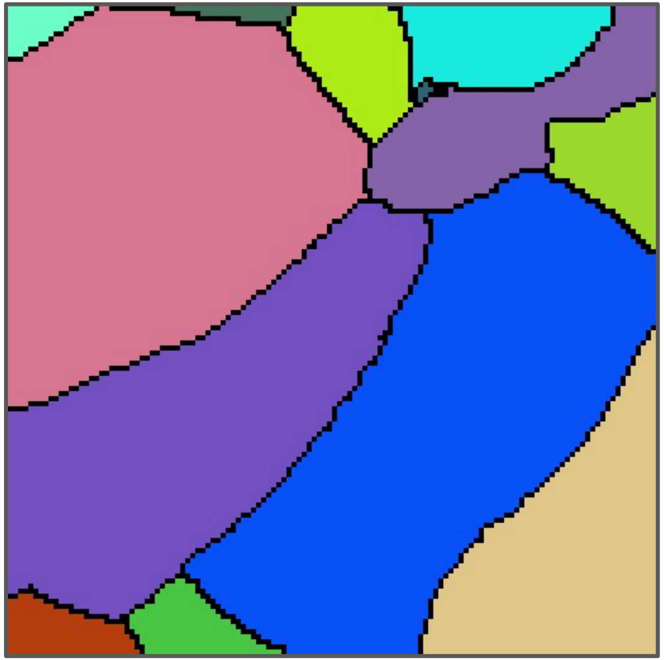

Decompressing Per-Pixel Labels

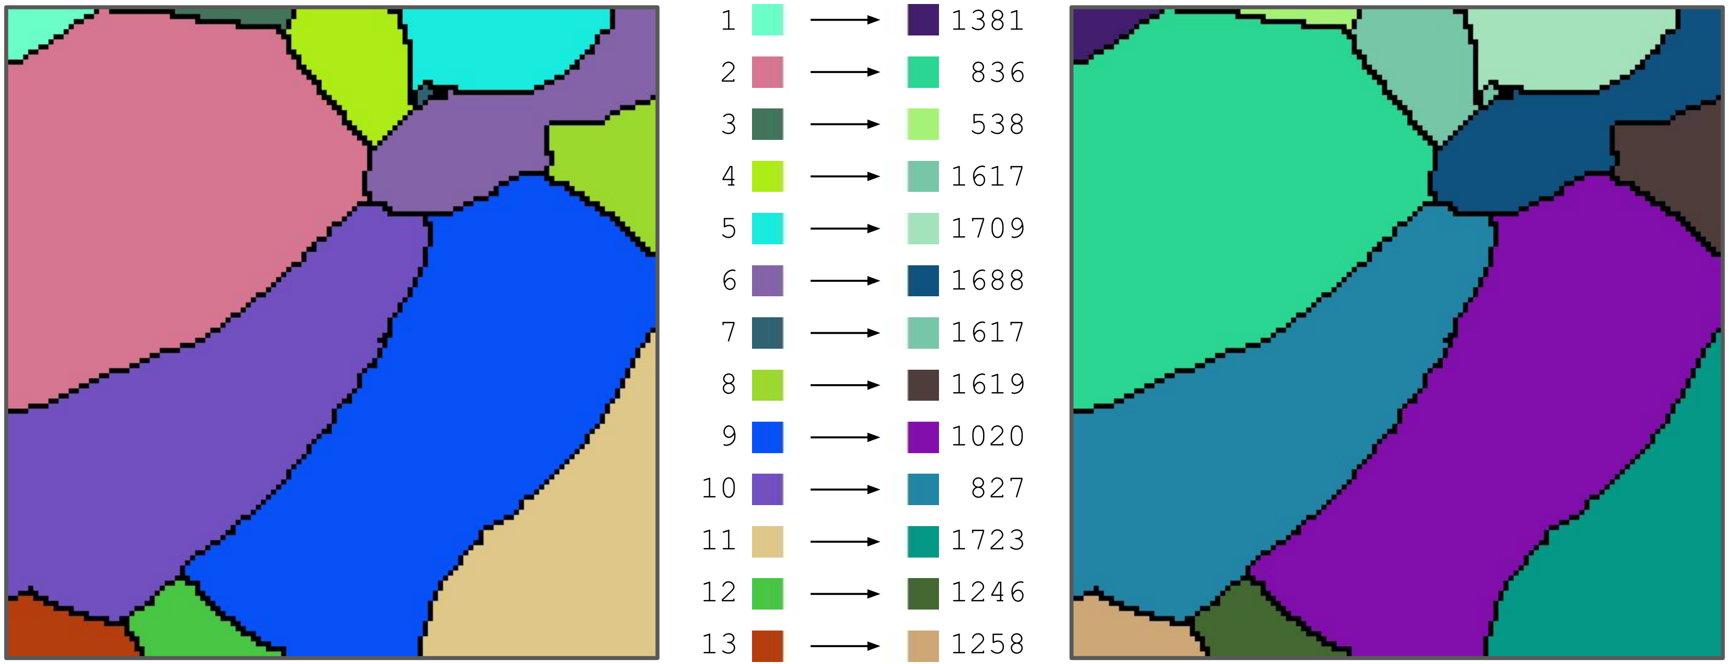

Decompressing Per-Pixel Labels

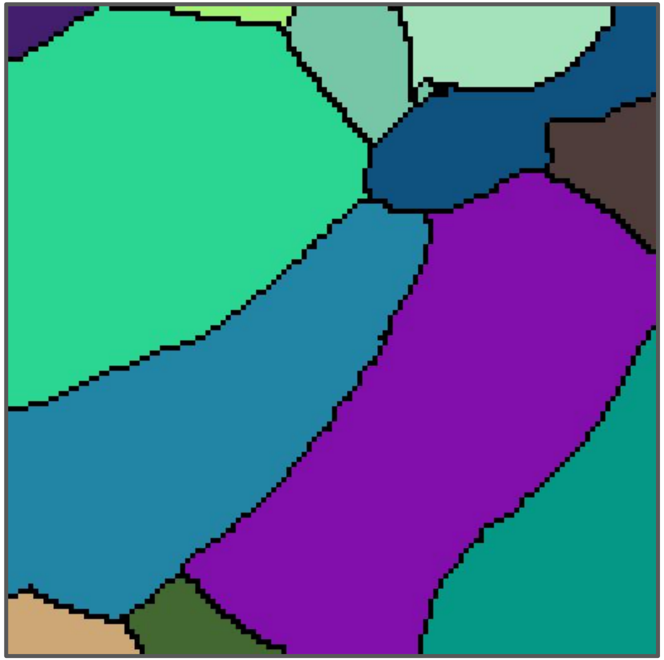

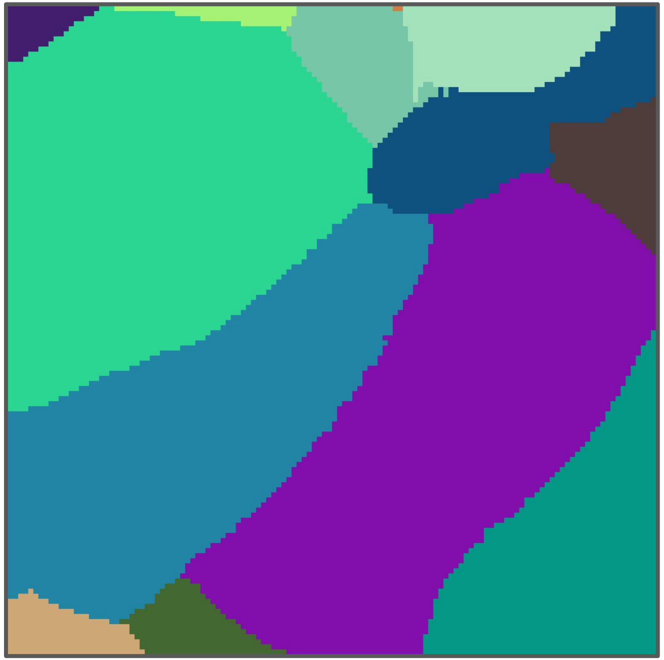

Variable 3D Window Sizes

Compresso can use different sized windows depending on the input data

Variable 3D Window Sizes

Compresso can use different sized windows depending on the input data

4x4x4 windows outperform 8x8x1 windows by 12.5% on an isotropic dataset

Isotropic Data

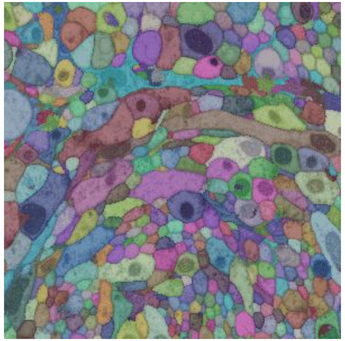

XY-plane

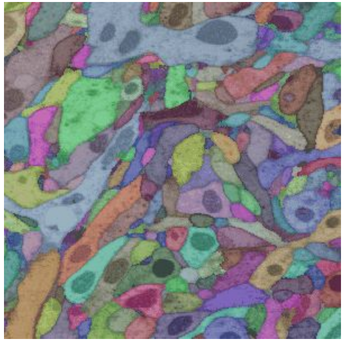

YZ-plane

Extends to Other Segmentation Datasets

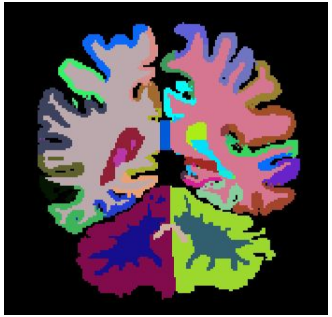

SPL Brain Atlas

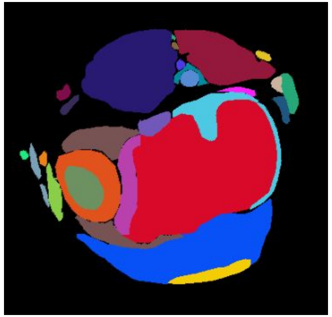

SPL Knee Atlas

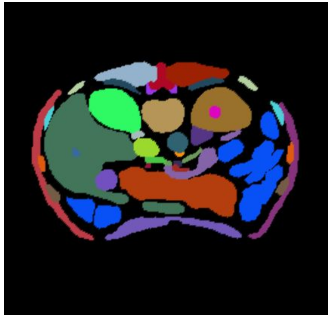

SPL Abdominal Atlas

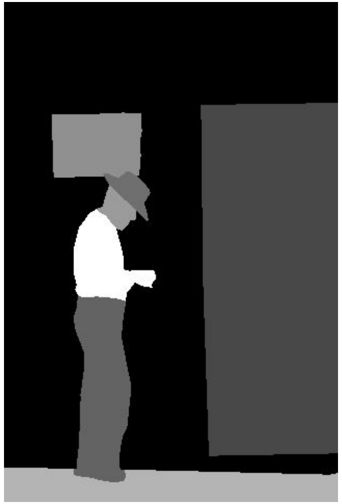

Berkeley Segmentation Dataset

PASCAL Visual Object Classes Dataset

Results

Two Stage Compression with LZMA

Follow Compresso with a general-purpose compression scheme such as LZMA

Two Stage Compression with LZMA

Follow Compresso with a general-purpose compression scheme such as LZMA

LZMA uses complex models for probability predictions of bits

Dataset	Uncompressed	Neuroglancer + LZMA	Compresso + LZMA
AC3 Mouse cortex, EM	1.26 GB	550x	**771x**
AC4 Mouse cortex, EM	838.86 MB	479x	**660x**
CREMI A, B, C Drosophila brain, EM	1.56 GB	465x, 629x, 496x	**804x, 1158x, 899x**
L. Cylinder Mouse cortex, EM	10.07 GB	425x	**889x**
SPL Brain Atlas T1/T2-weighted MRI	135.27 MB	**764x**	645x
SPL Knee Atlas MRI	249.56 MB	1172x	**1562x**
SPL Abdominal Atlas CT	59.24 MB	417x	**482x**

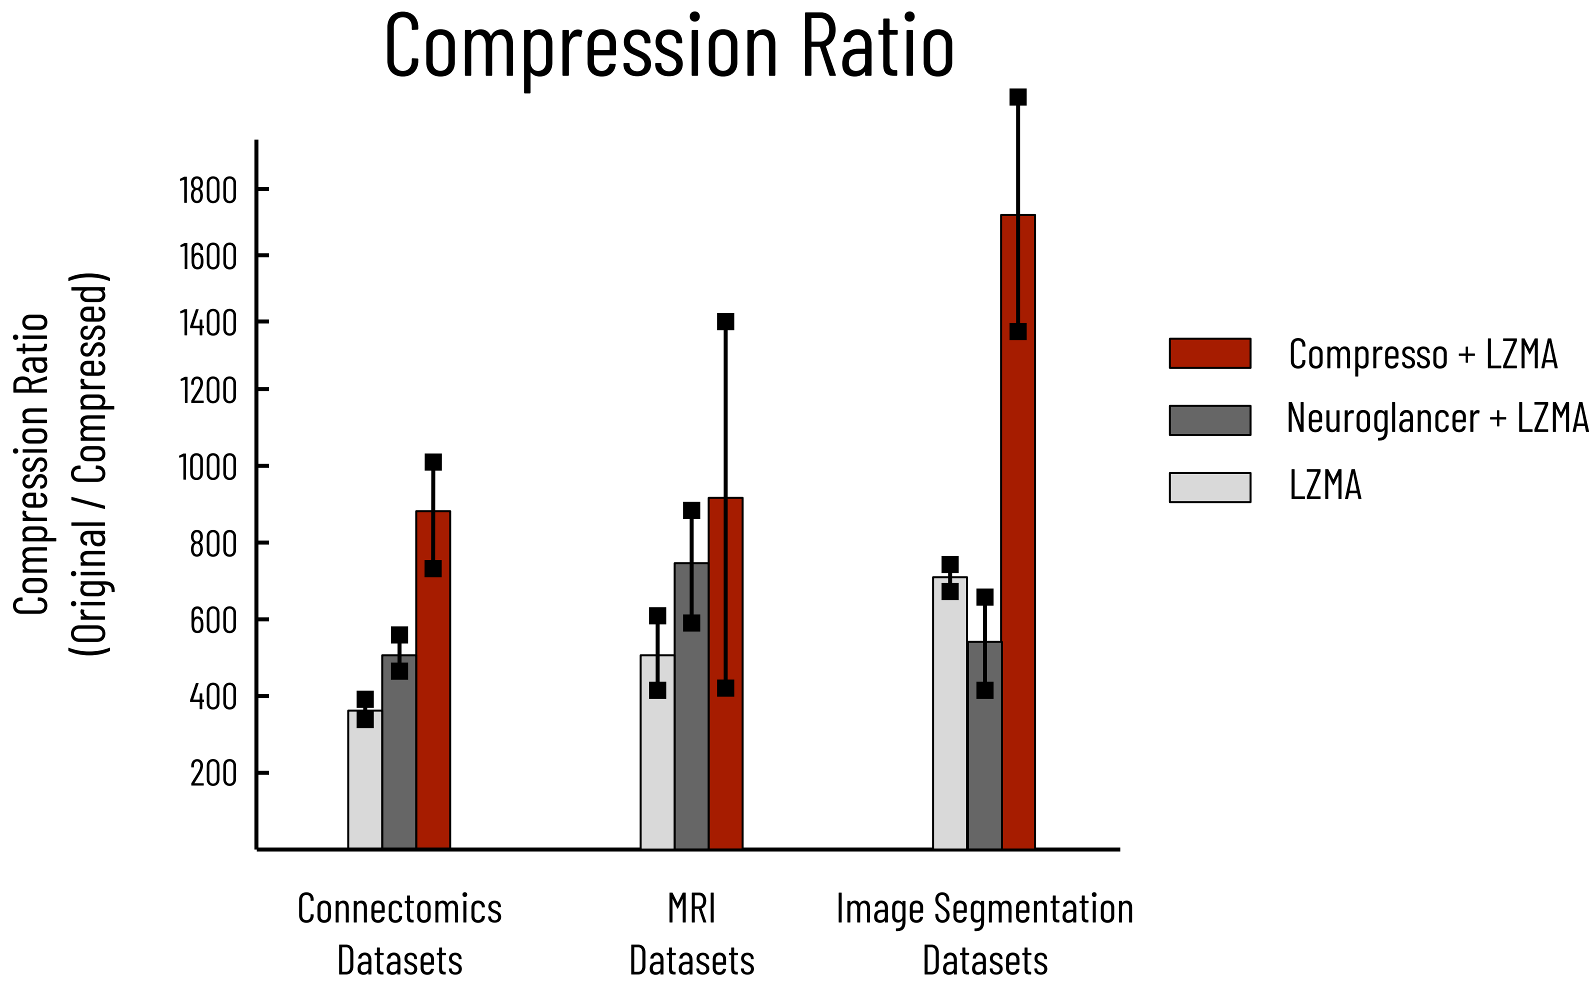
Compression Ratio
Compression Ratio
(Original / Compressed)
1800
1600
1400
1200
1000
800
600
400
200
Connectomics
Datasets
MRI
Datasets
Image Segmentation
Datasets
Compresso + LZMA
Neuroglancer + LZMA
LZMA

Method	Compression Speed	Decompression Speed
LZMA	9.89 MB / s	366.13 MB / s
Neuroglancer + LZMA	43.80 MB / s	164.32 MB / s
Compresso + LZMA	131.15 MB / s	206.60 MB / s

Complexity

P is the number of pixels; N is the number of distinct window values; X and Y are the size of the x and y dimensions of the input data; and α is the inverse Ackermann function.

Compression:

$$O(P(1 + \alpha(XY)) + N \log N)$$

Decompression:

$$O(P(1 + \alpha(XY)))$$

Compression of 100 microns cubed

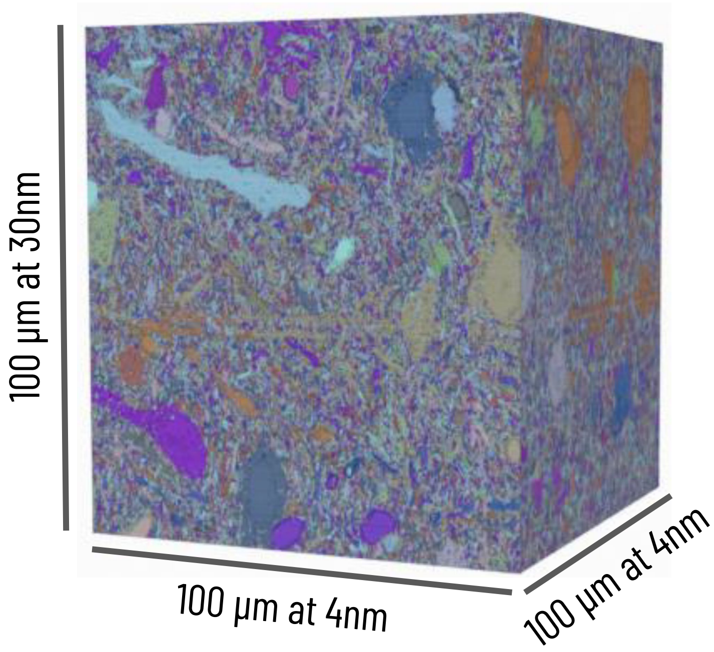

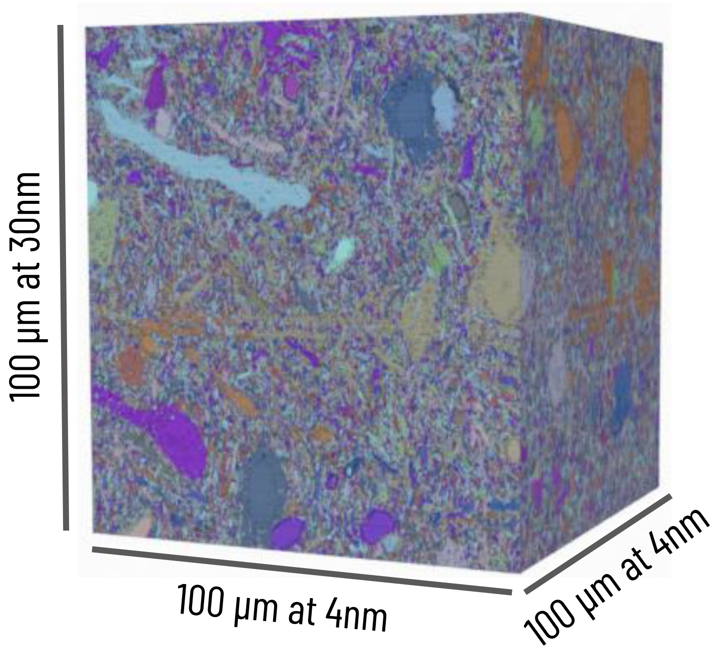

Uncompressed: 19.25 terabytes

Compression of 100 microns cubed

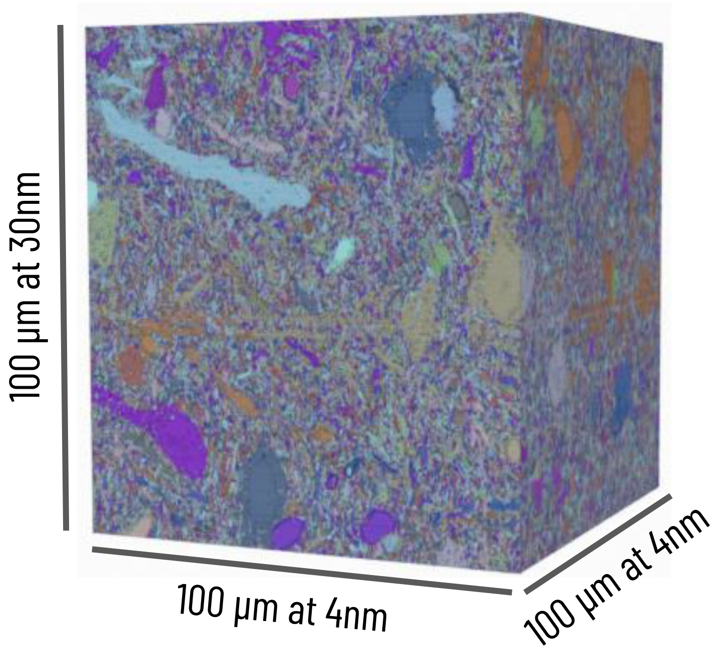

Uncompressed:	19.25 terabytes
With Compresso + LZMA:	25.94 gigabytes
Ratio:	742x

Compression of 100 microns cubed

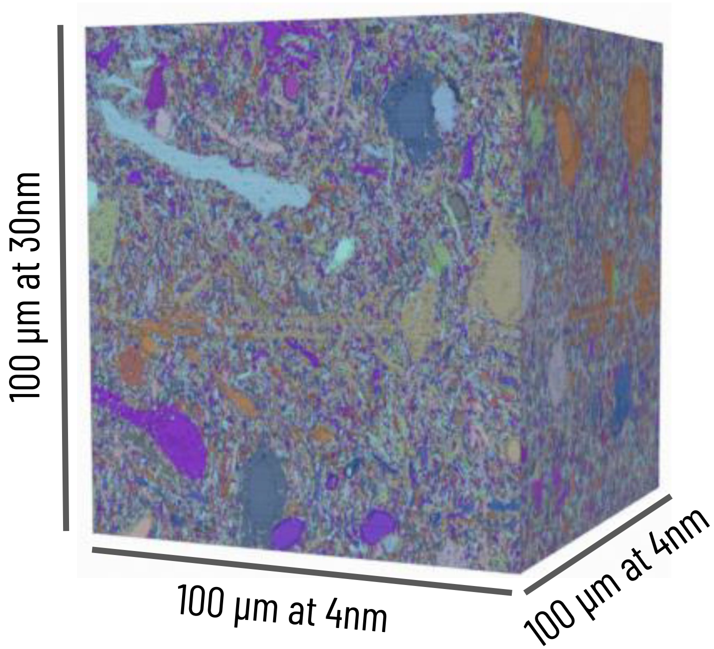

Uncompressed:	19.25 terabytes
With Compresso + LZMA:	25.94 gigabytes
Ratio:	742x

AWS Storage Costs (S3 Standard Storage):

Uncompressed:	\$442.75 / month
Compressed:	\$0.60 / month

Connectomics Pipeline

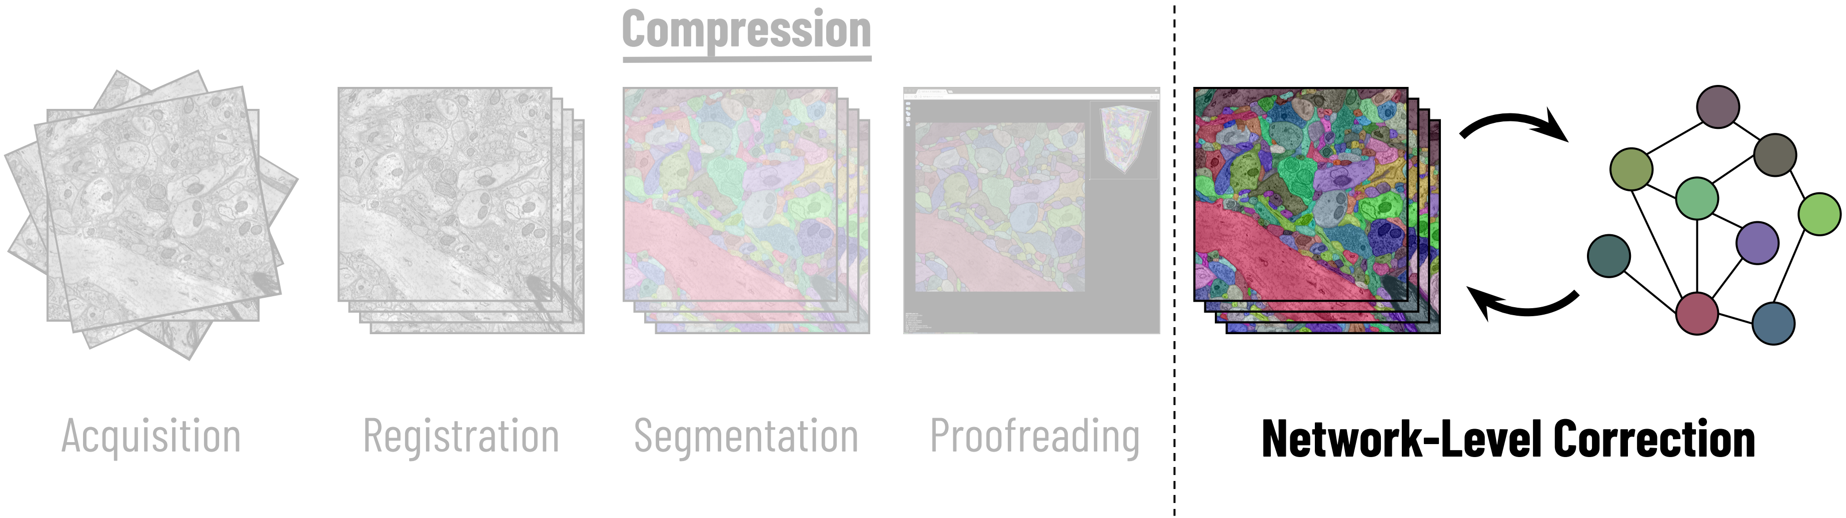

Biologically-Constrained Region Merging for Connectome Reconstruction

Brian Matejek, Daniel Haehn, Donglai Wei, Toufiq Parag, Hanspeter Pfister

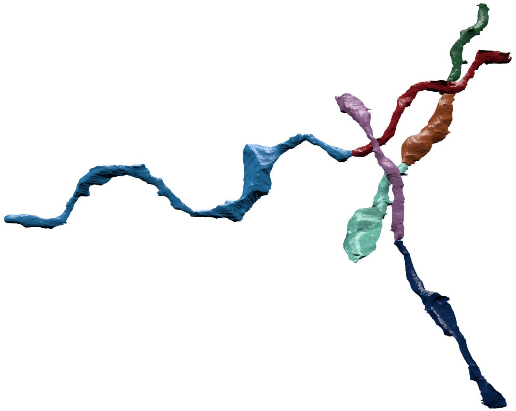

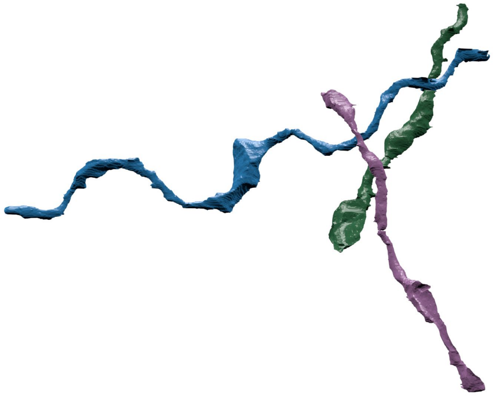

Under Review, ECCV 2018

U-Net

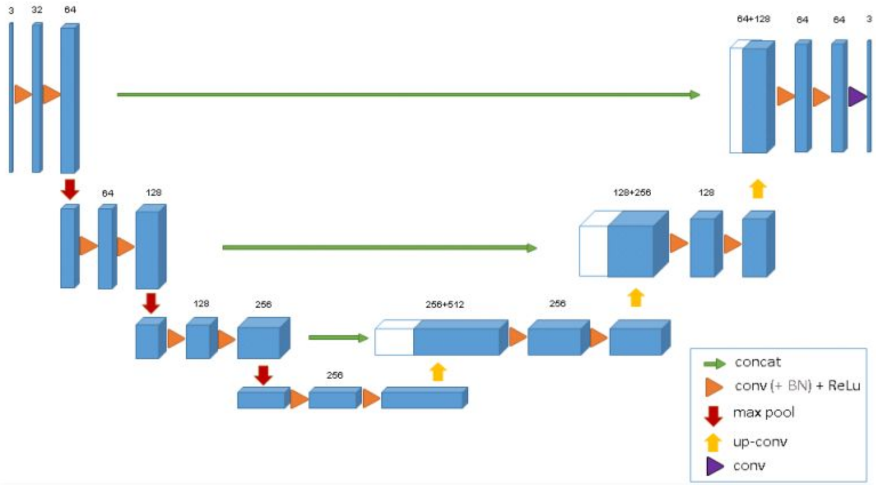

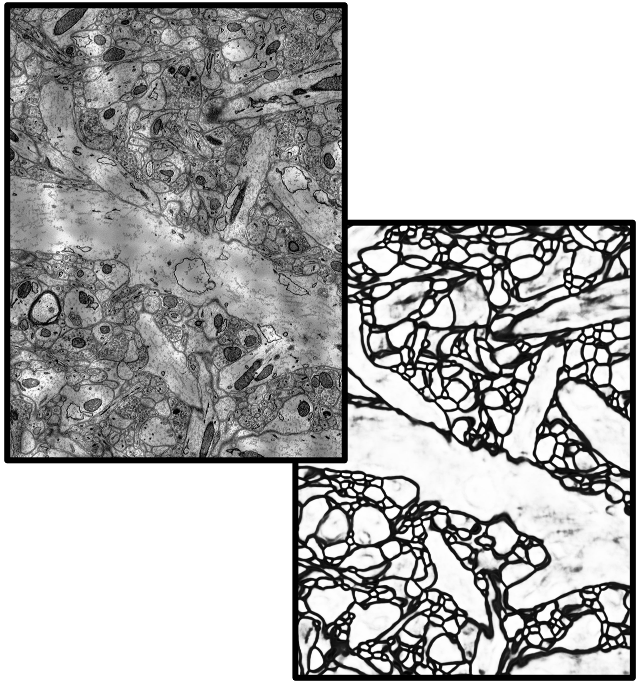

Ronneberger et al., U-Net: Convolutional Networks for Biomedical Image Segmentation, MICCAI 2015
Cicek et al., 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation, MICCAI 2016

3D Watershed on Affinities

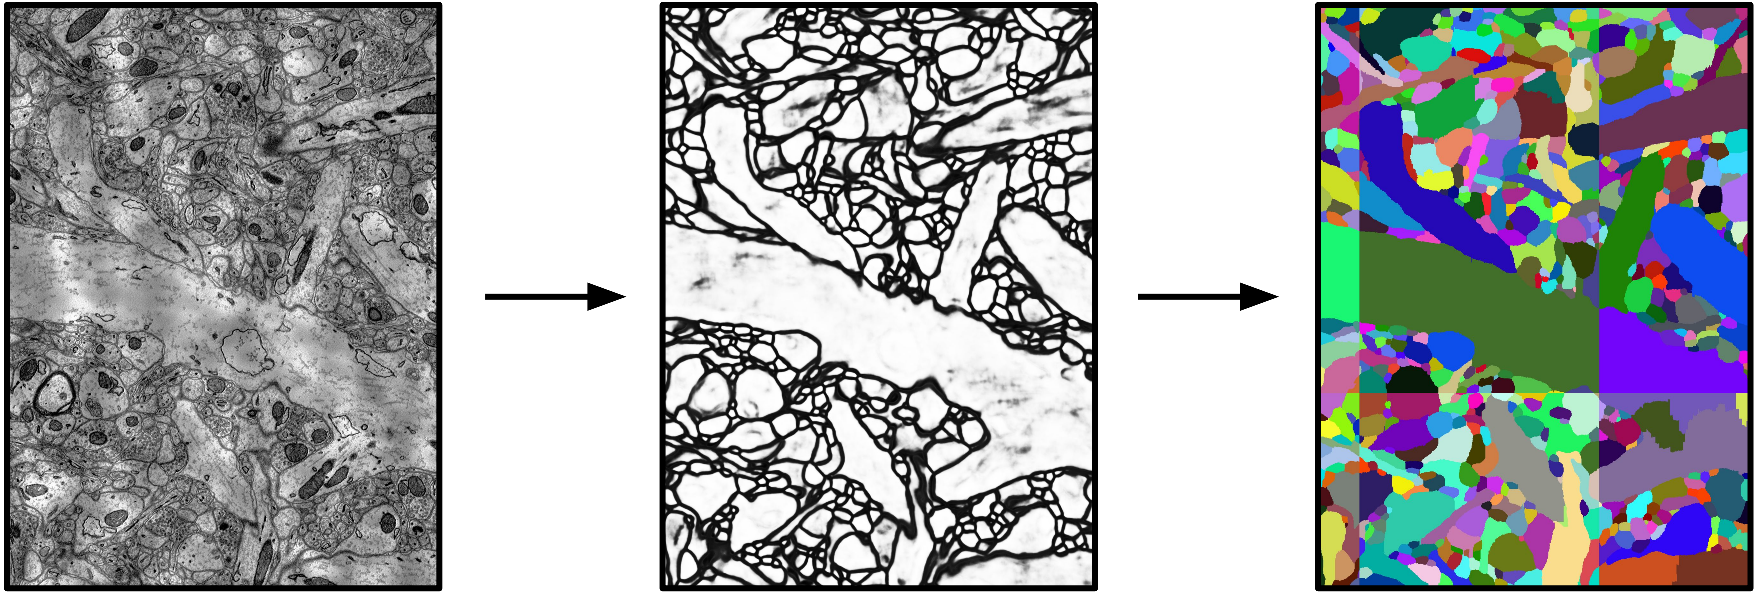

Zlateski et al., Image Segmentation by Size-Dependent Single Linkage Clustering of a Watershed Basin Graph, 2015
Funke et al., A Deep Structured Learning Approach Towards Automating Connectome Reconstruction from 3D Electron Micrographs, 2017
Zeng et al., DeepEM3D: Approaching Human-Level Performance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

Context-Aware Delayed Agglomeration

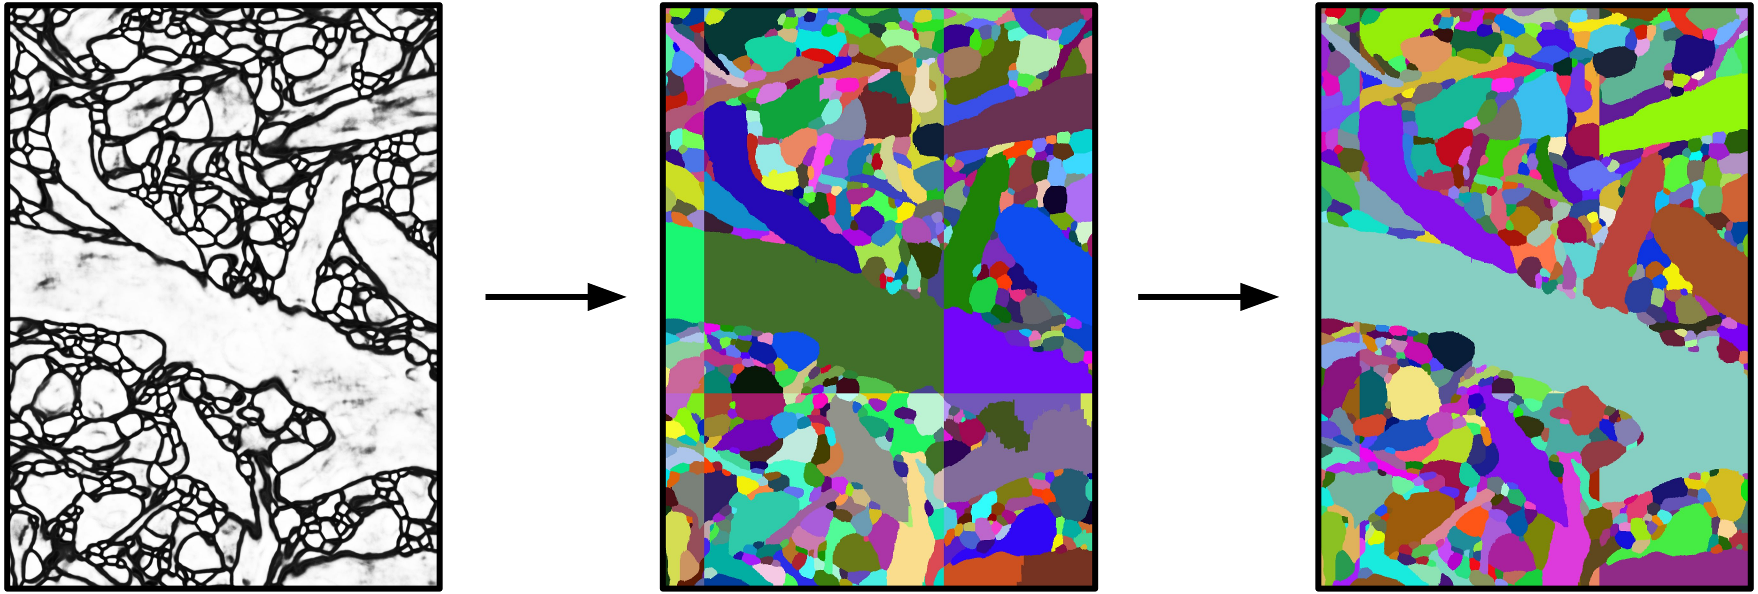

Nunez-Iglesias et al., Machine Learning of Hierarchical Clustering to Segment 2D and 3D Images, PLoS ONE, 2013
Parag et al., A Context-Aware Delayed Agglomeration Framework for Electron Microscopy Segmentation, PLoS ONE 2015

Errors

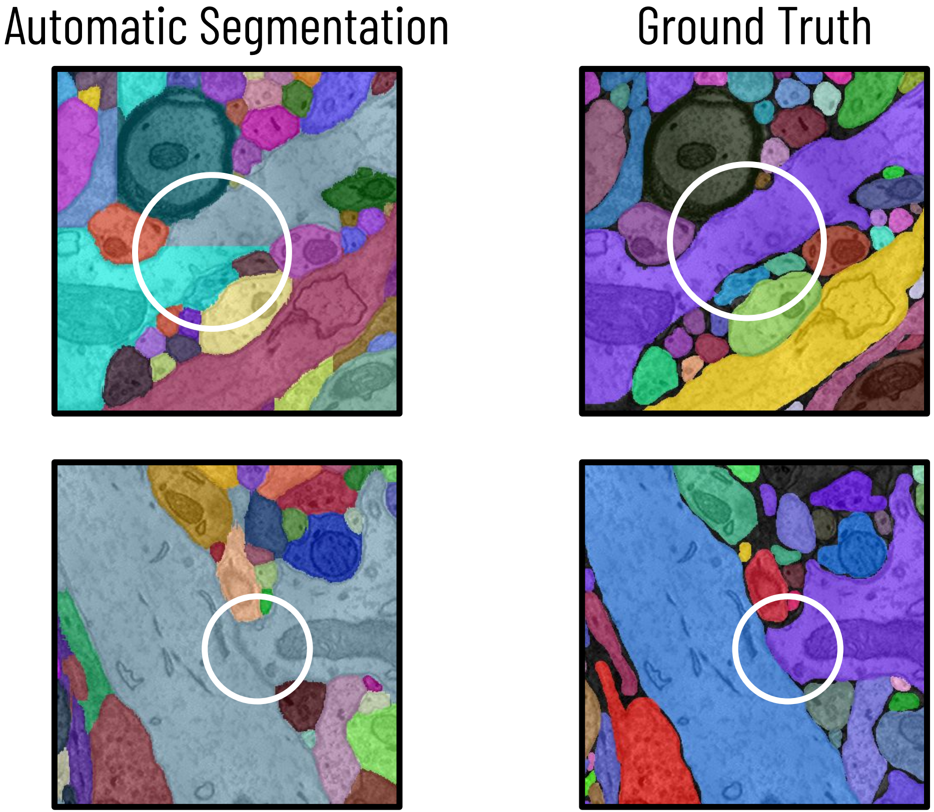

Proofreading and Error Correction

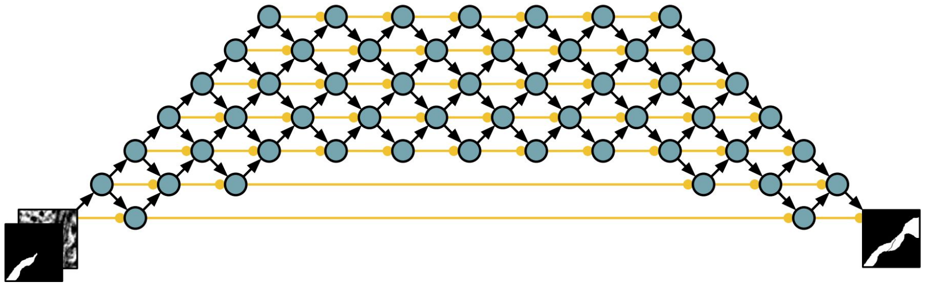

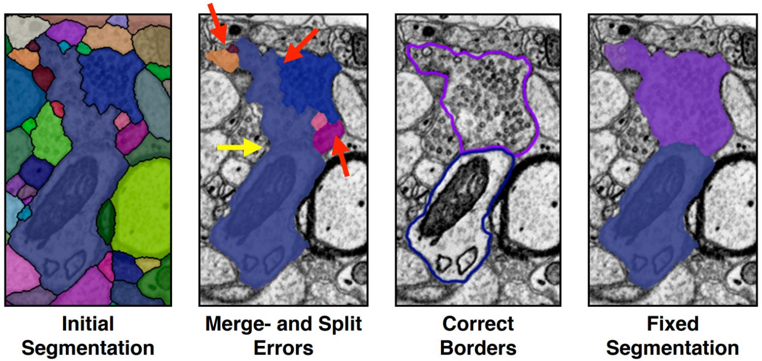

Zung et al., An Error Detection and Correction Framework for Connectomics, NIPS 2017
Haehn et al., Guided Proofreading of Automatic Segmentations for Connectomics, CVPR 2018

Need for Global Context

Stained Images

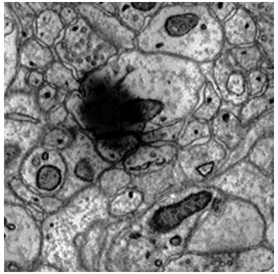

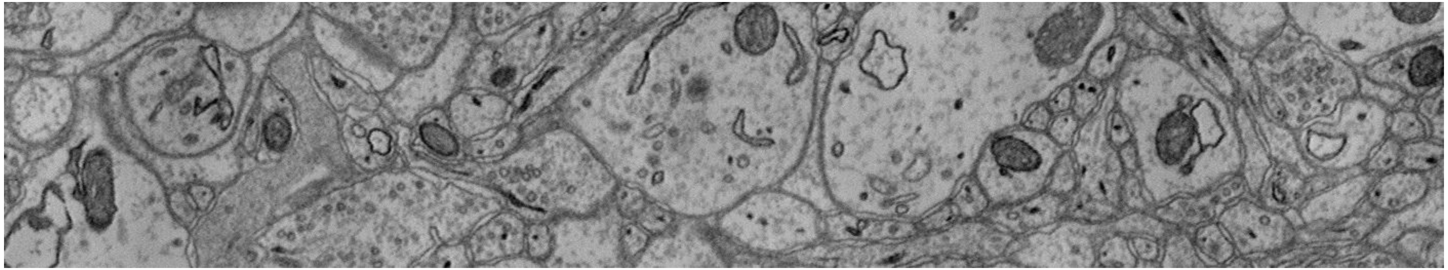

Missing Sections

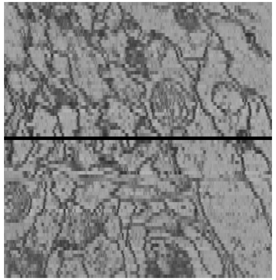

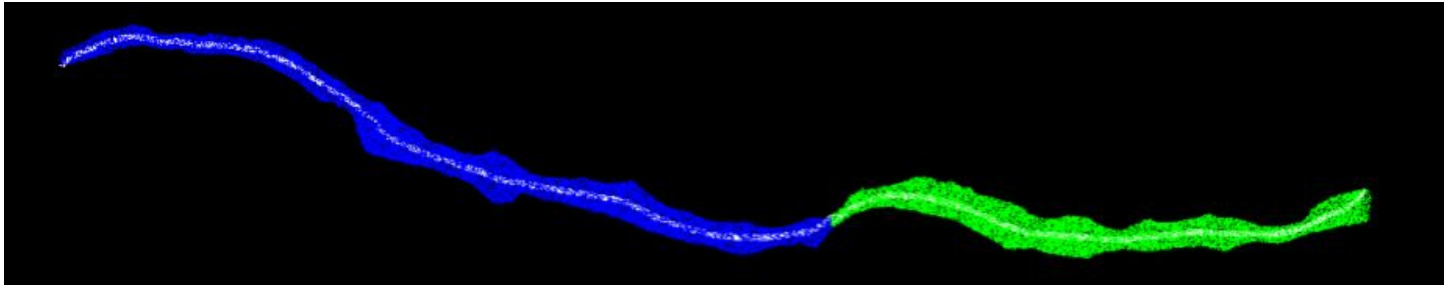

Variable Image Data

xy-slice

yz-slice

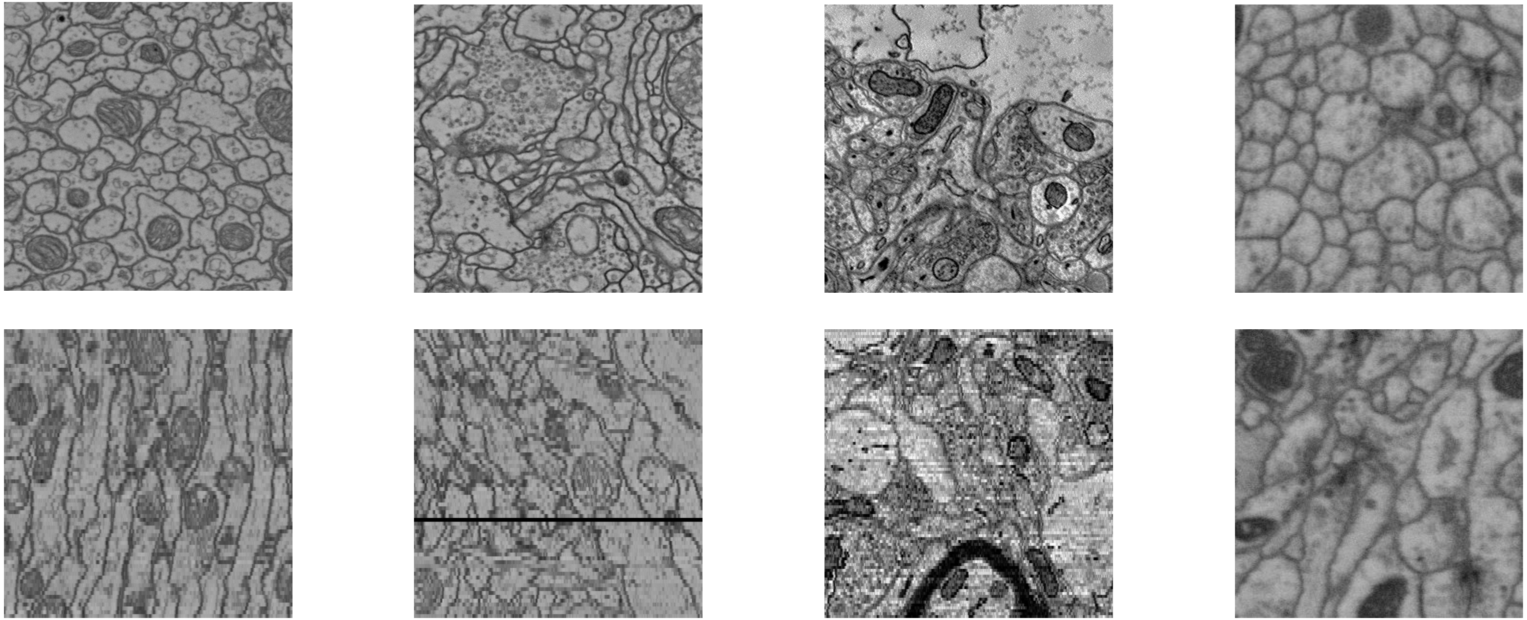

Cremi Vol. A	Cremi Vol. C	Kasthuri	FlyEM
Drosophila melanogaster	*Drosophila melanogaster*	*Mouse*	*Drosophila melanogaster*
4 x 4 x 40 nm^3 / vx	4 x 4 x 40 nm^3 / vx	6 x 6 x 30 nm^3 / vx	10 x 10 x 10 nm^3 / vx

Proposed Region Merging

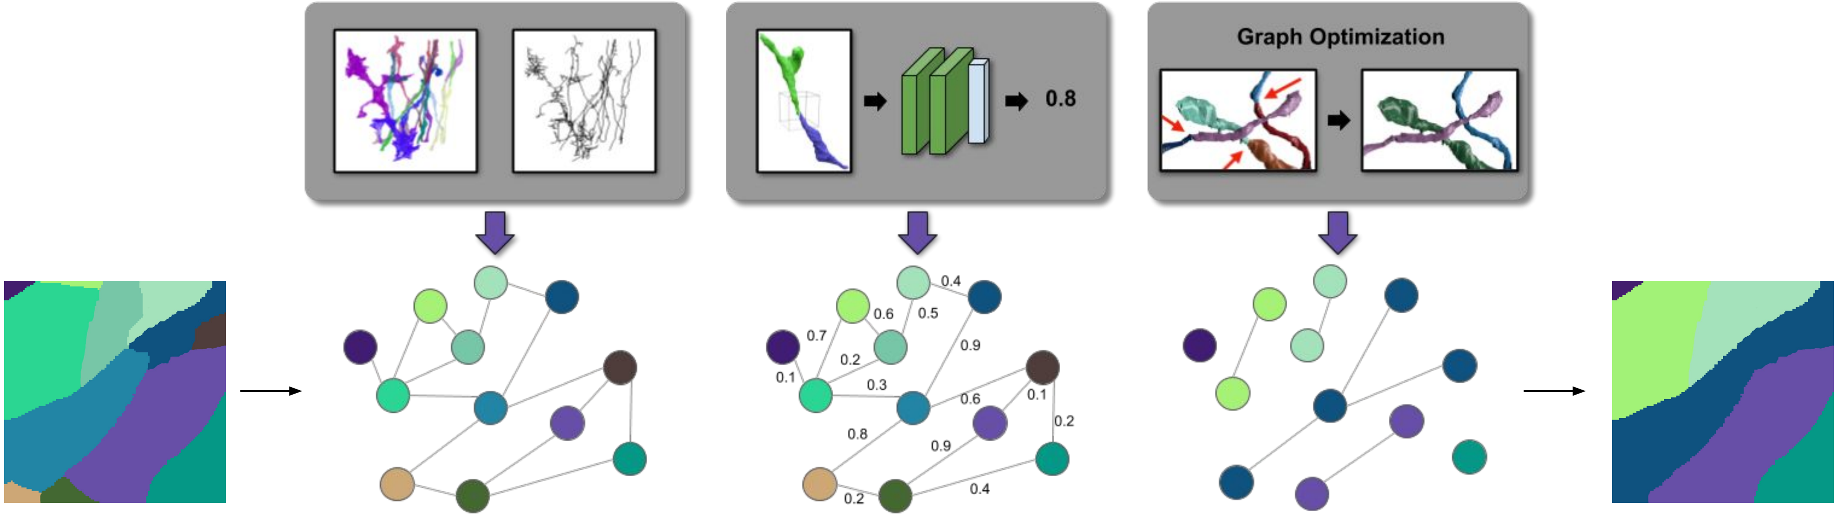

Proposed Region Merging with Biological Constraints

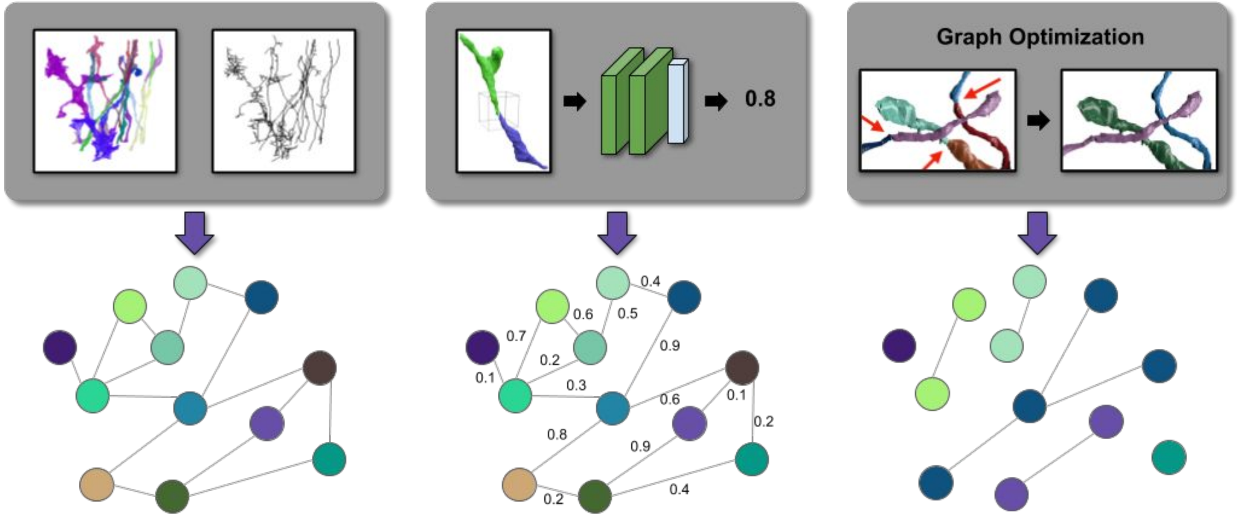

Geometric Priors | Learned Constraints | Topological Restrictions

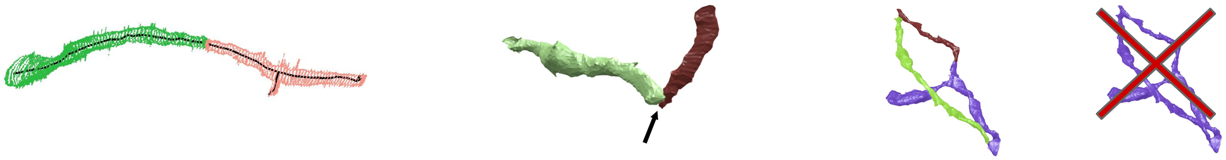

Goal: Construct a graph with as few extra edges as possible

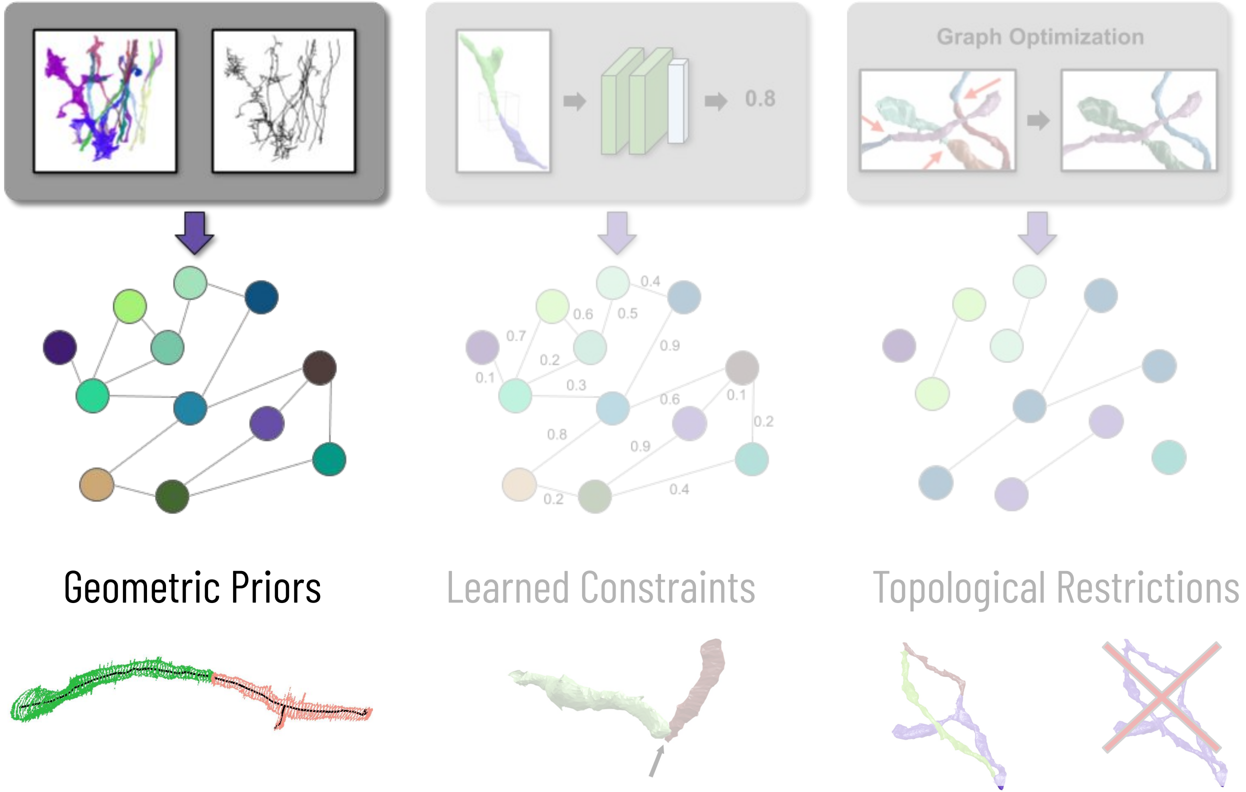

Adjacency Graphs

Every segment in the label volume receives a node

Segments with a pair of neighboring voxels receive an edge between the corresponding nodes

Adjacency Graphs

Every segment in the label volume receives a node

Segments with a pair of neighboring voxels receive an edge between the corresponding nodes

Typical Segment

103 Adjacent Neighbors

Non-Adjacent Split Errors

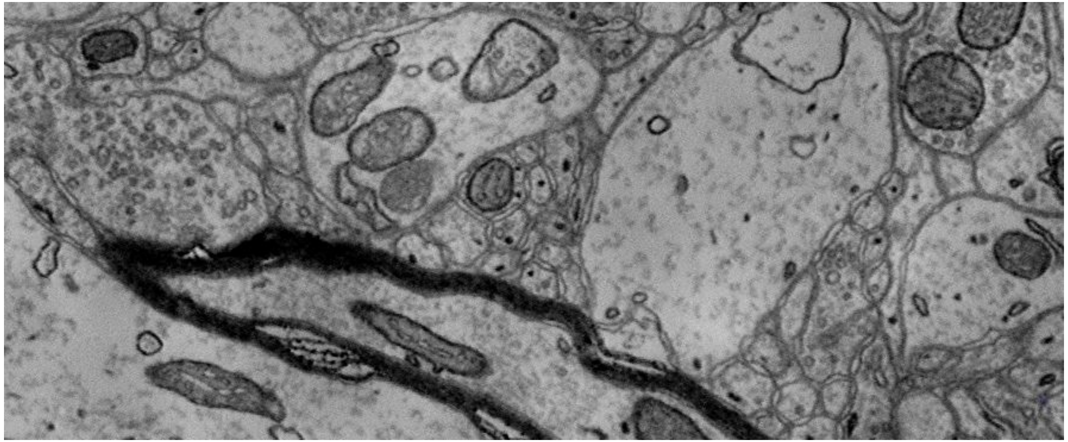

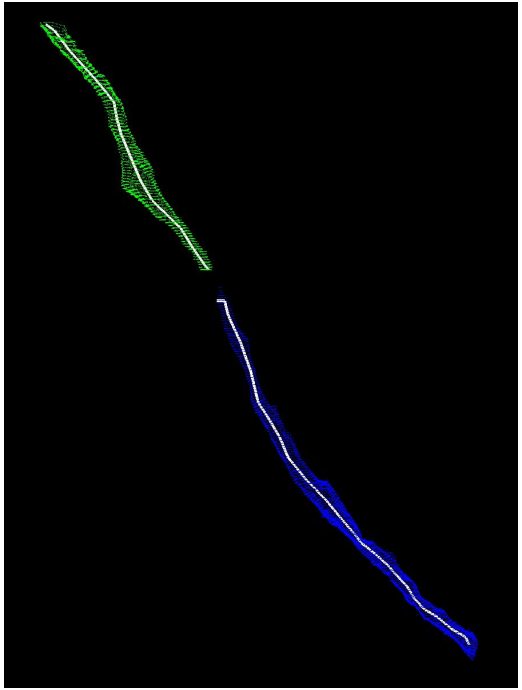

Correctly Segmented

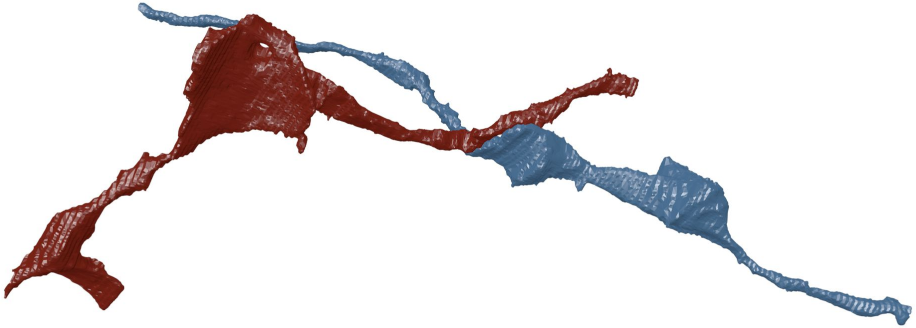

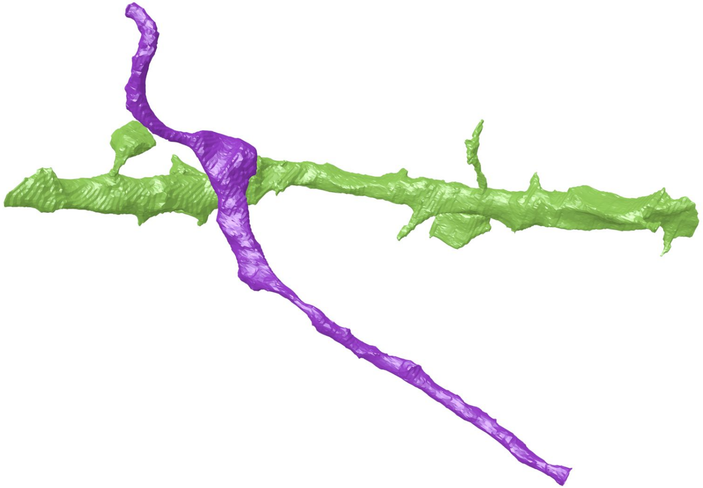

Incorrectly Split

Incorrectly Split

Skeletonization

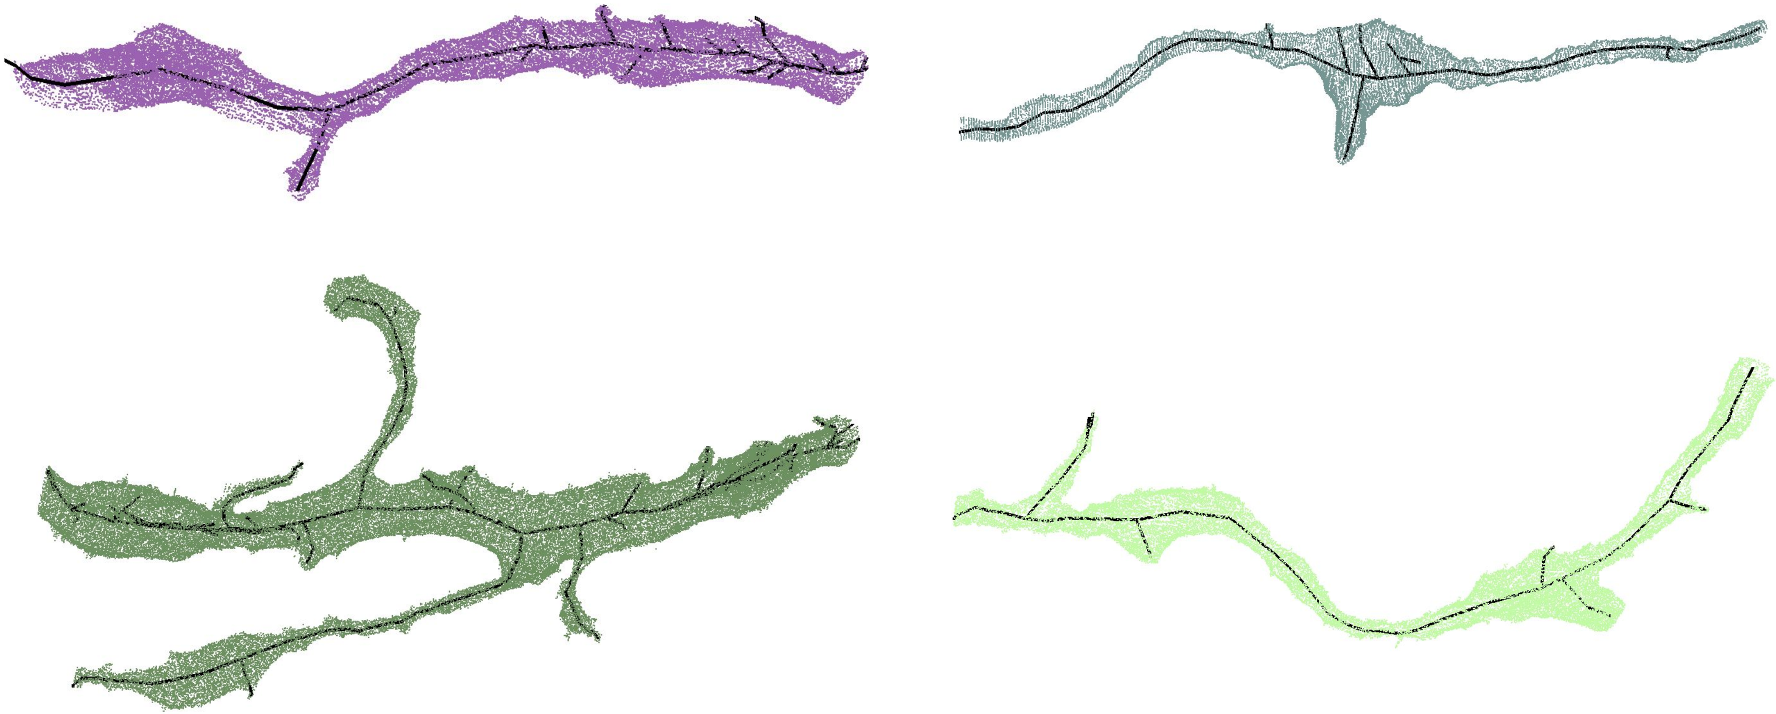

Merge Candidate

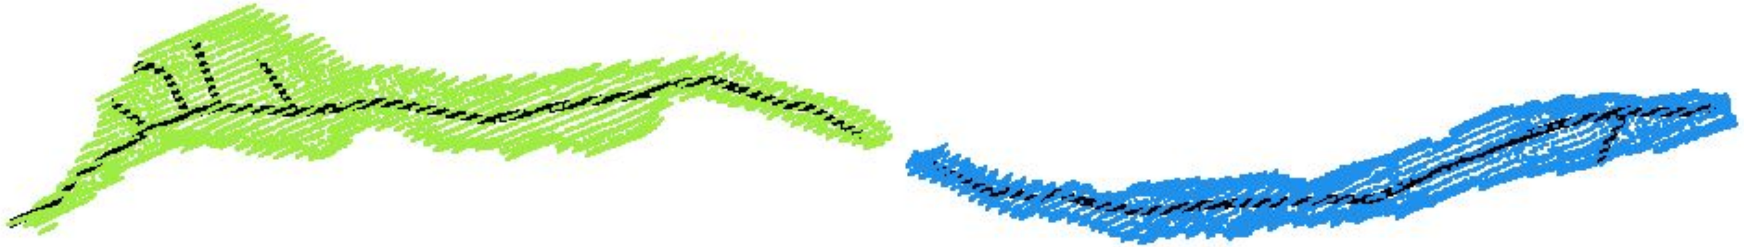

Merge Candidate

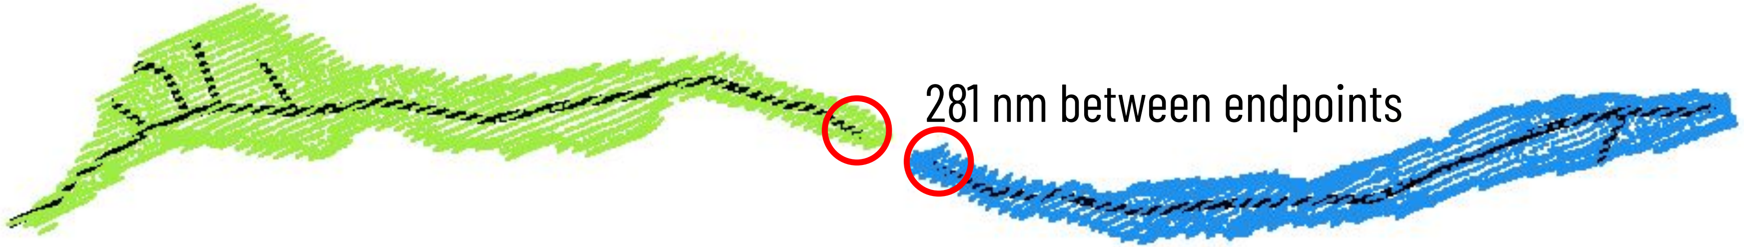

Pruned Candidate

Pruned Candidate

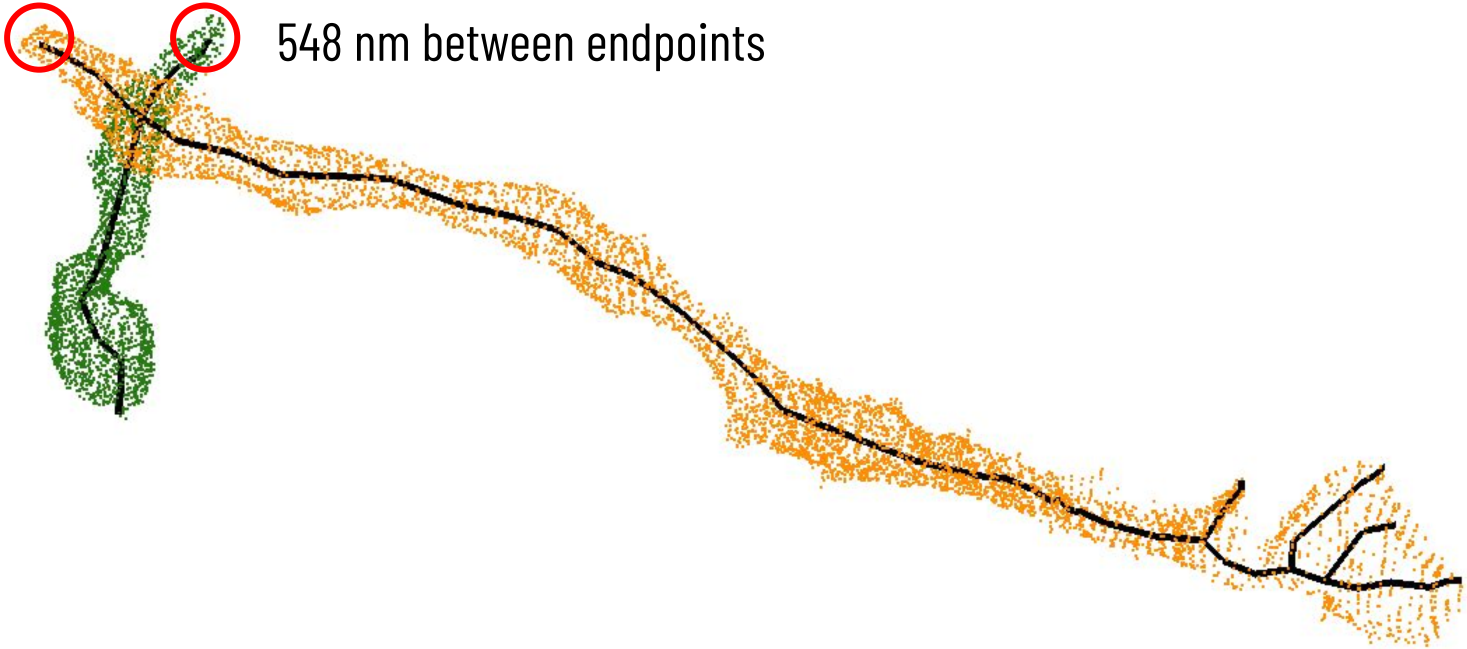

548 nm between endpoints

Number of Edges

Typical Segment

103

Before

Number of Edges

Typical Segment

103

Before

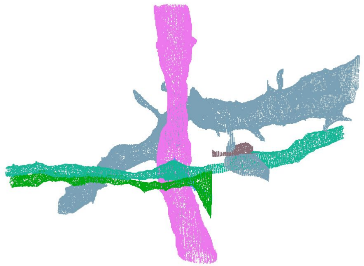

5

After

Goal: Populate edge weights for the graph

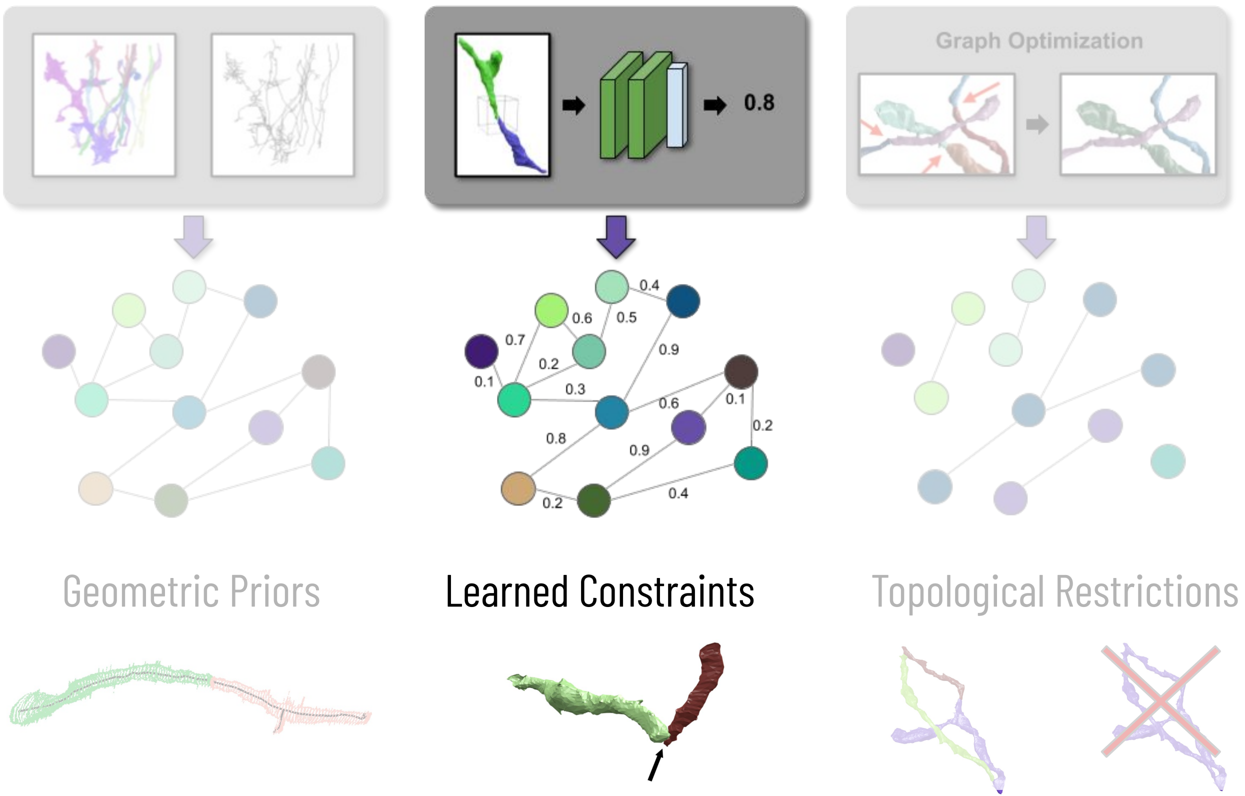

Learned Constraints

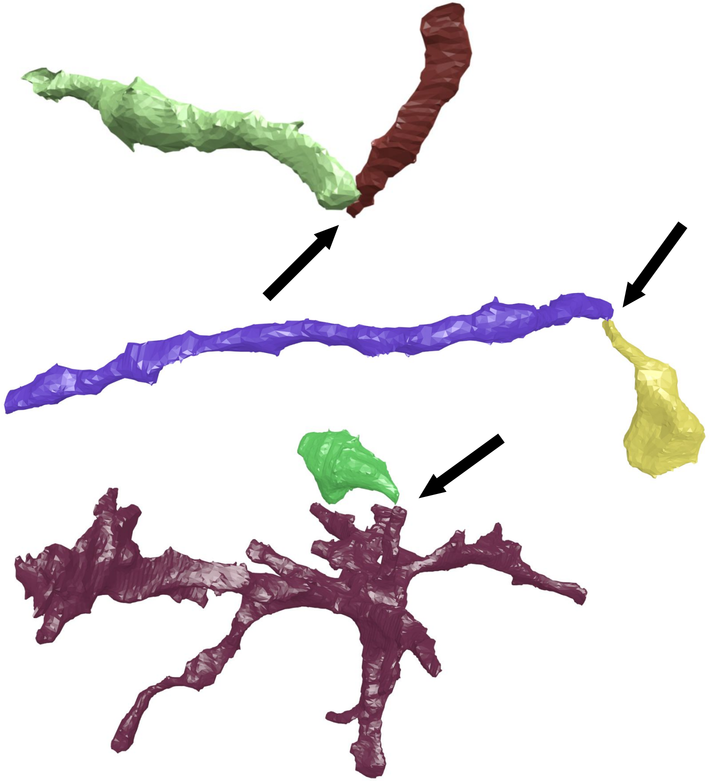

Learned Constraints

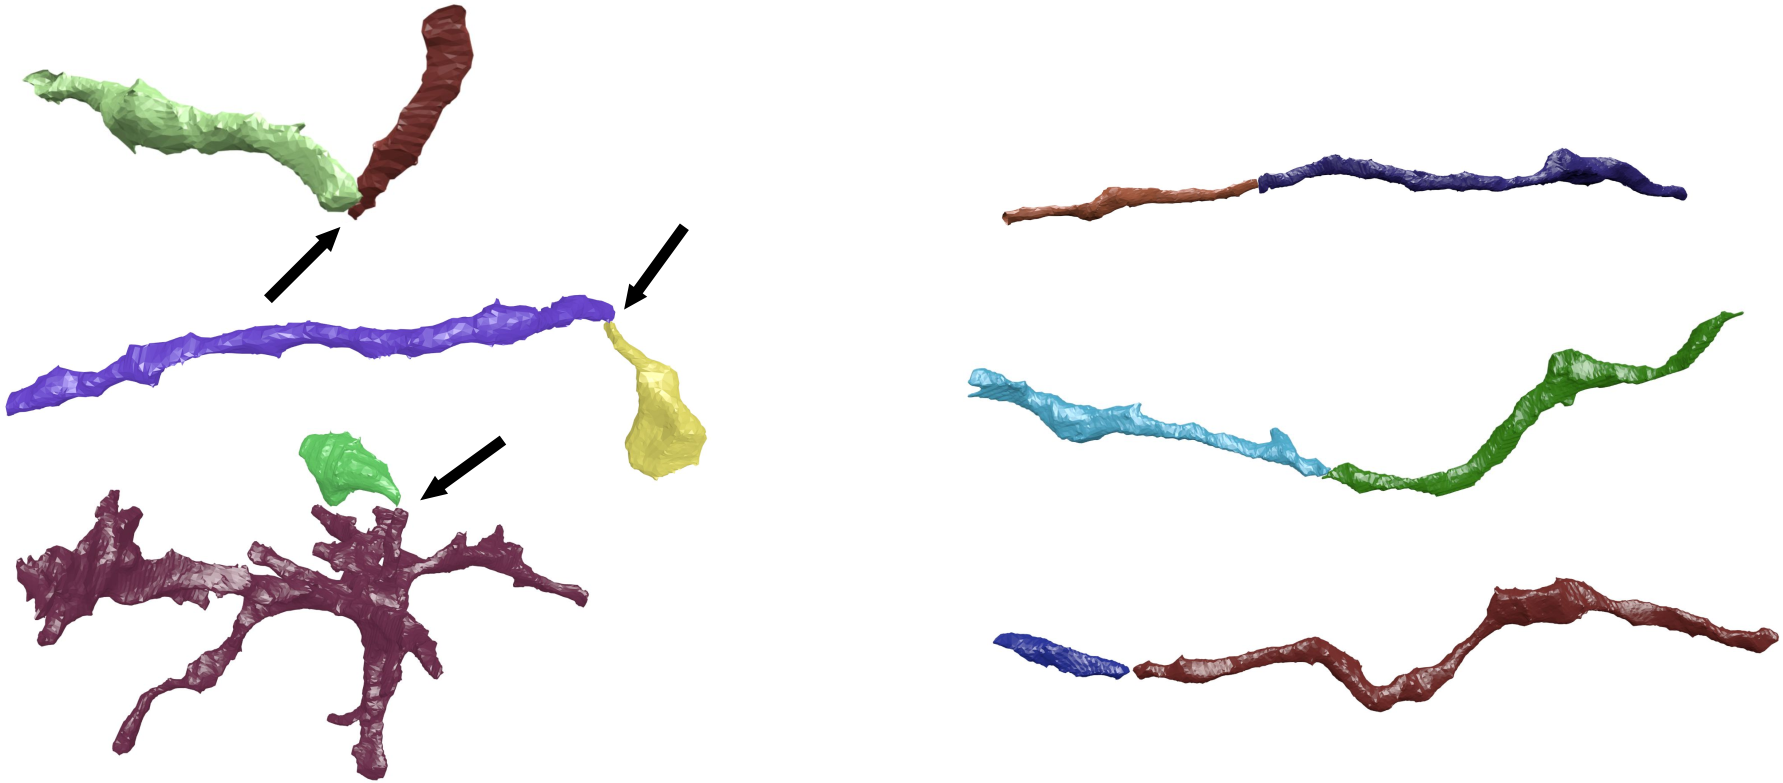

Architecture and Training Parameters

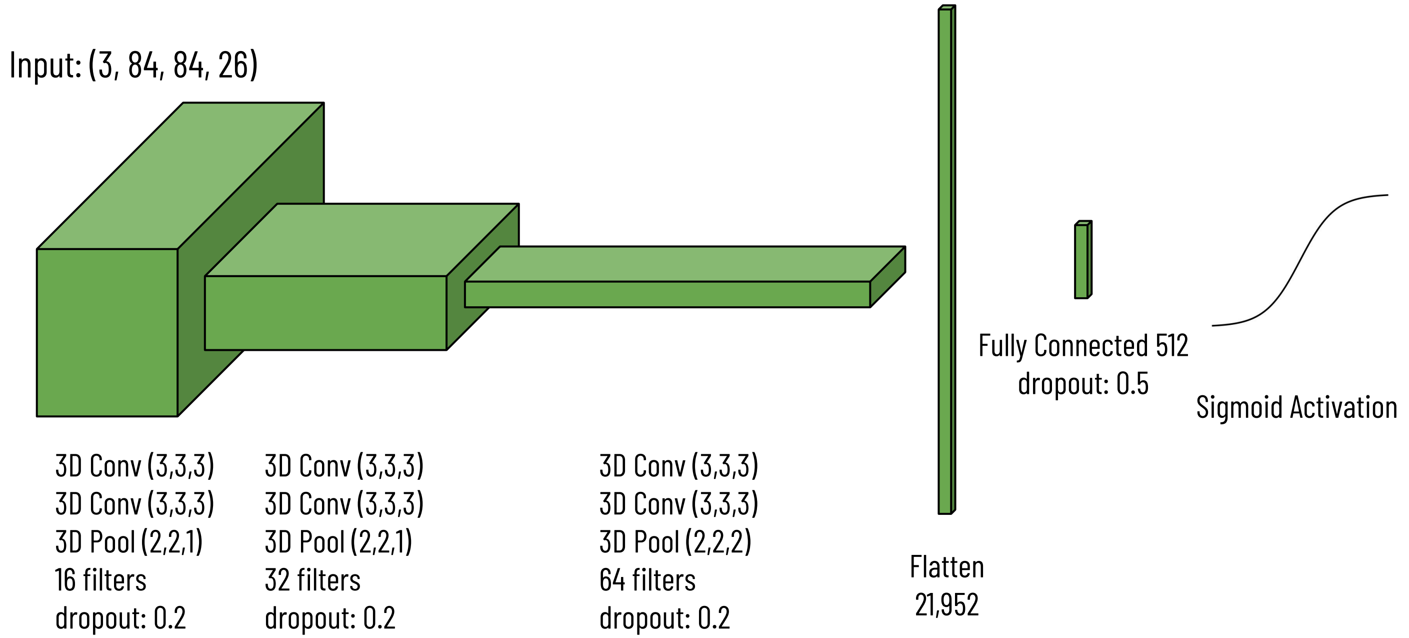

Architecture and Training Parameters

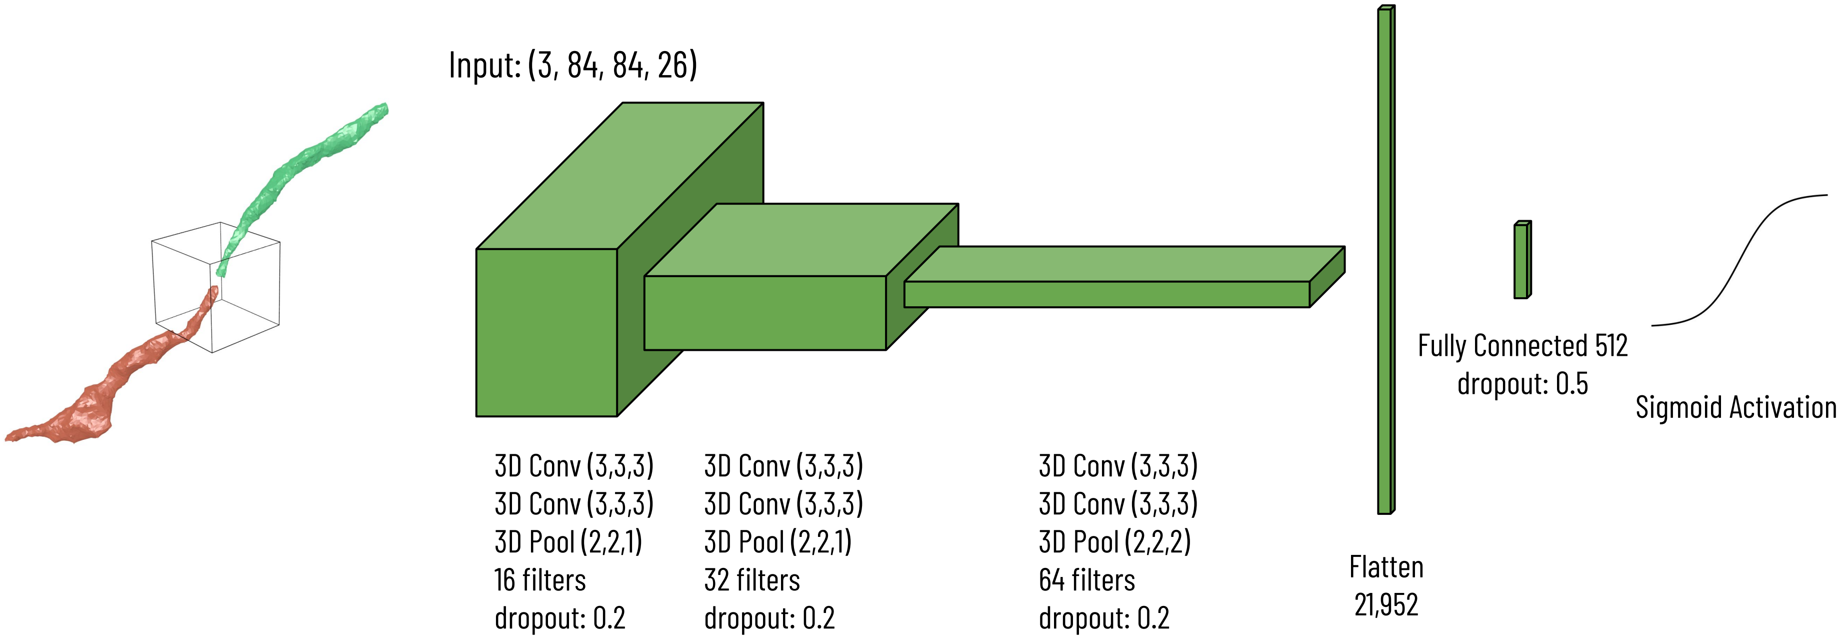

Architecture and Training Parameters

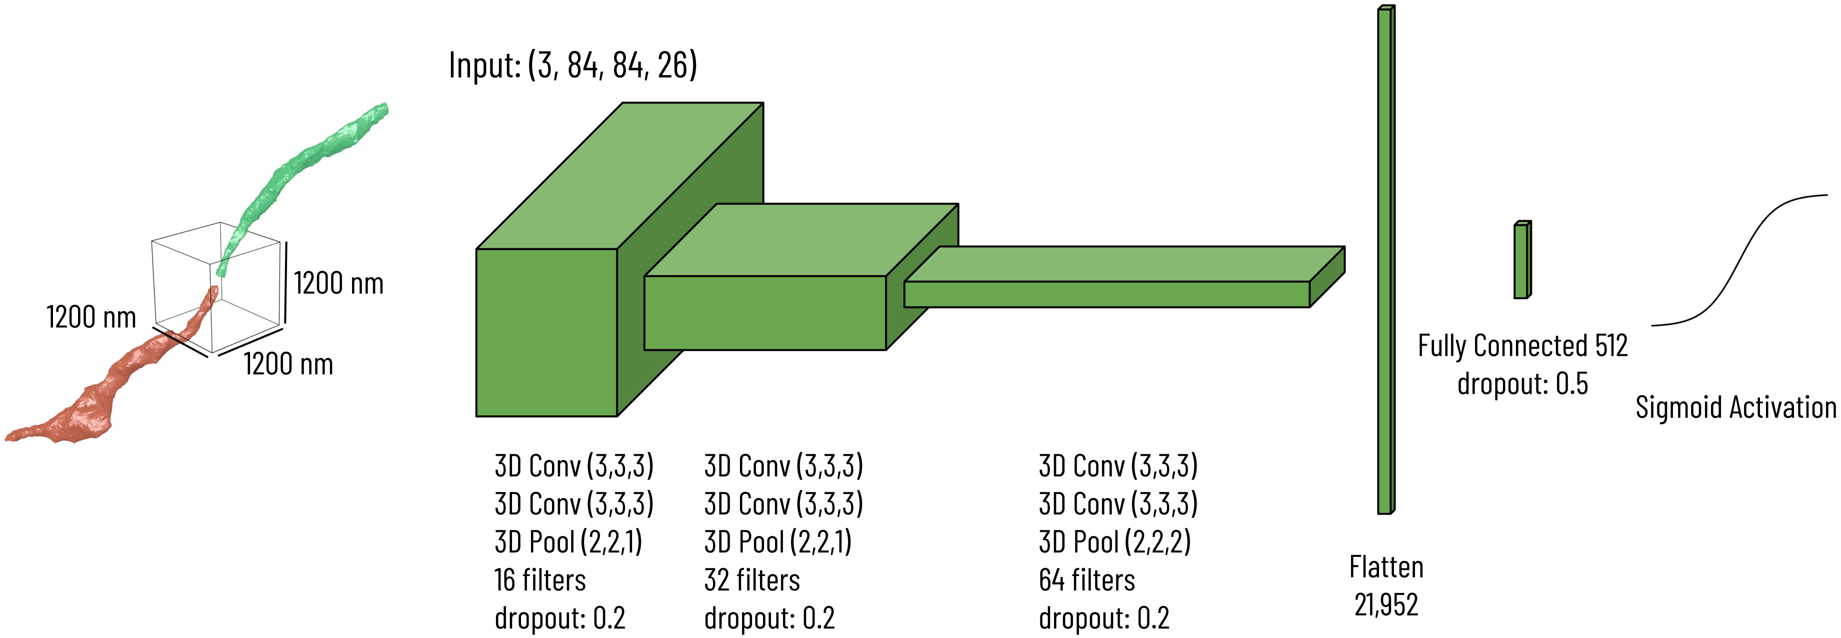

Architecture and Training Parameters

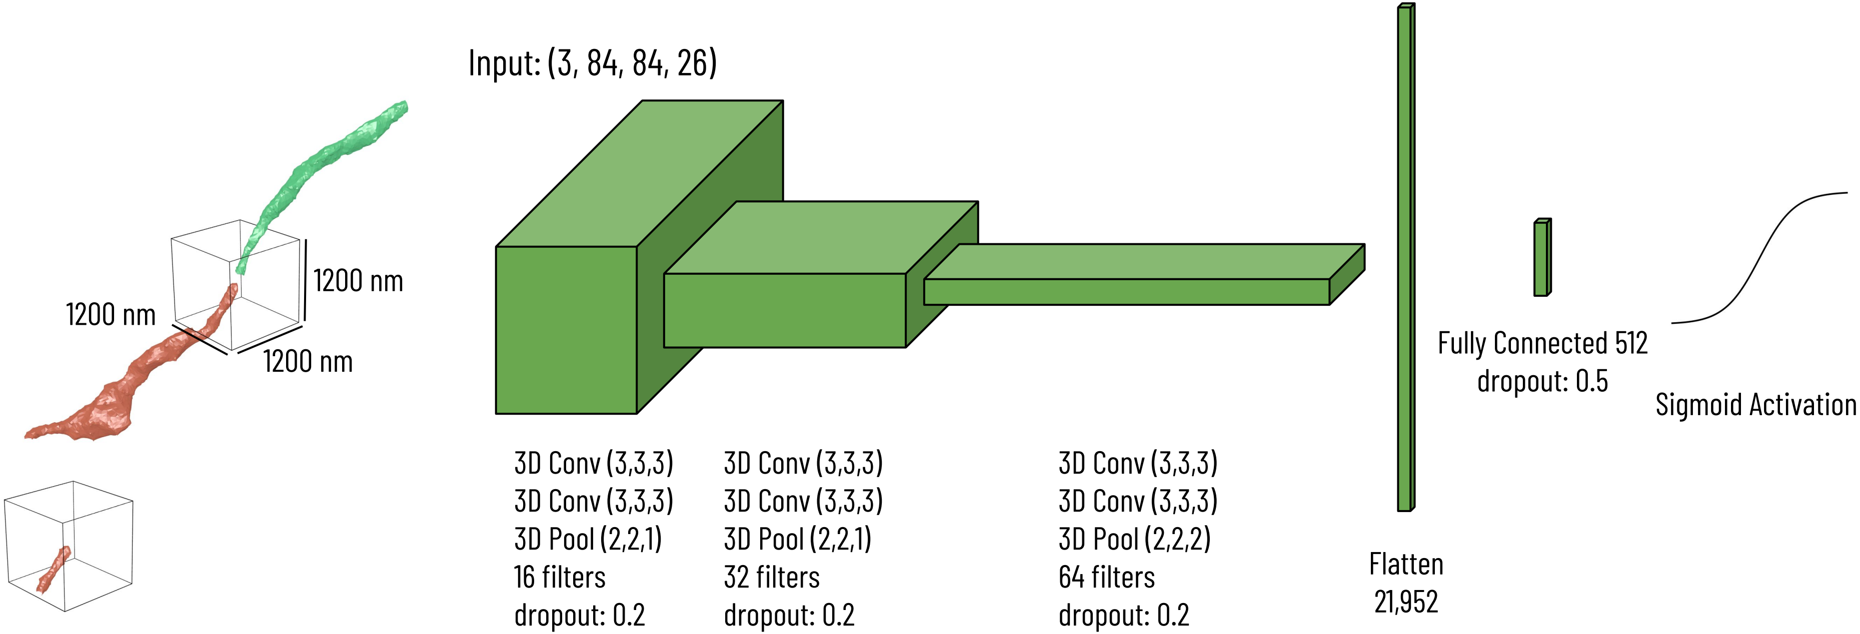

Architecture and Training Parameters

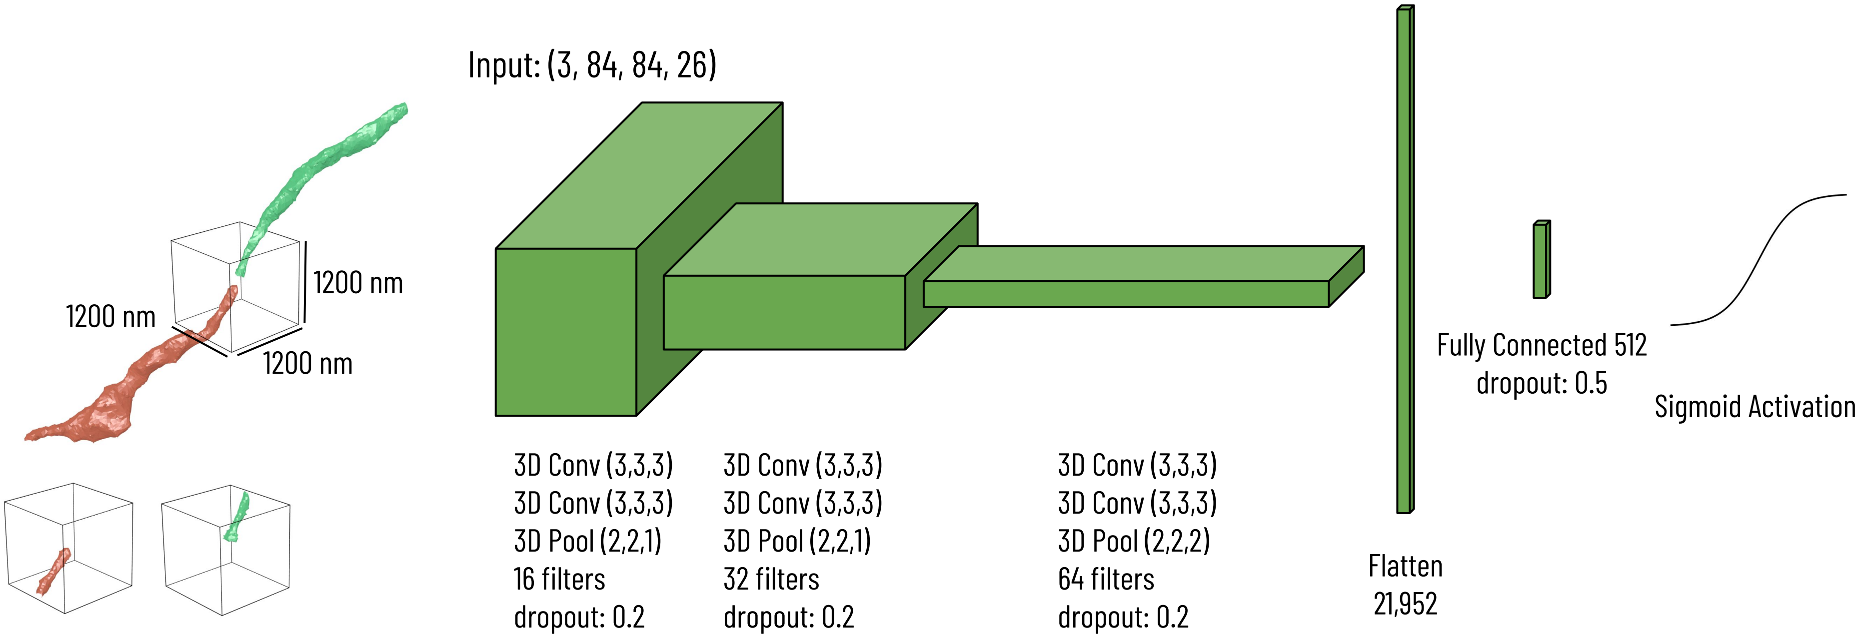

Architecture and Training Parameters

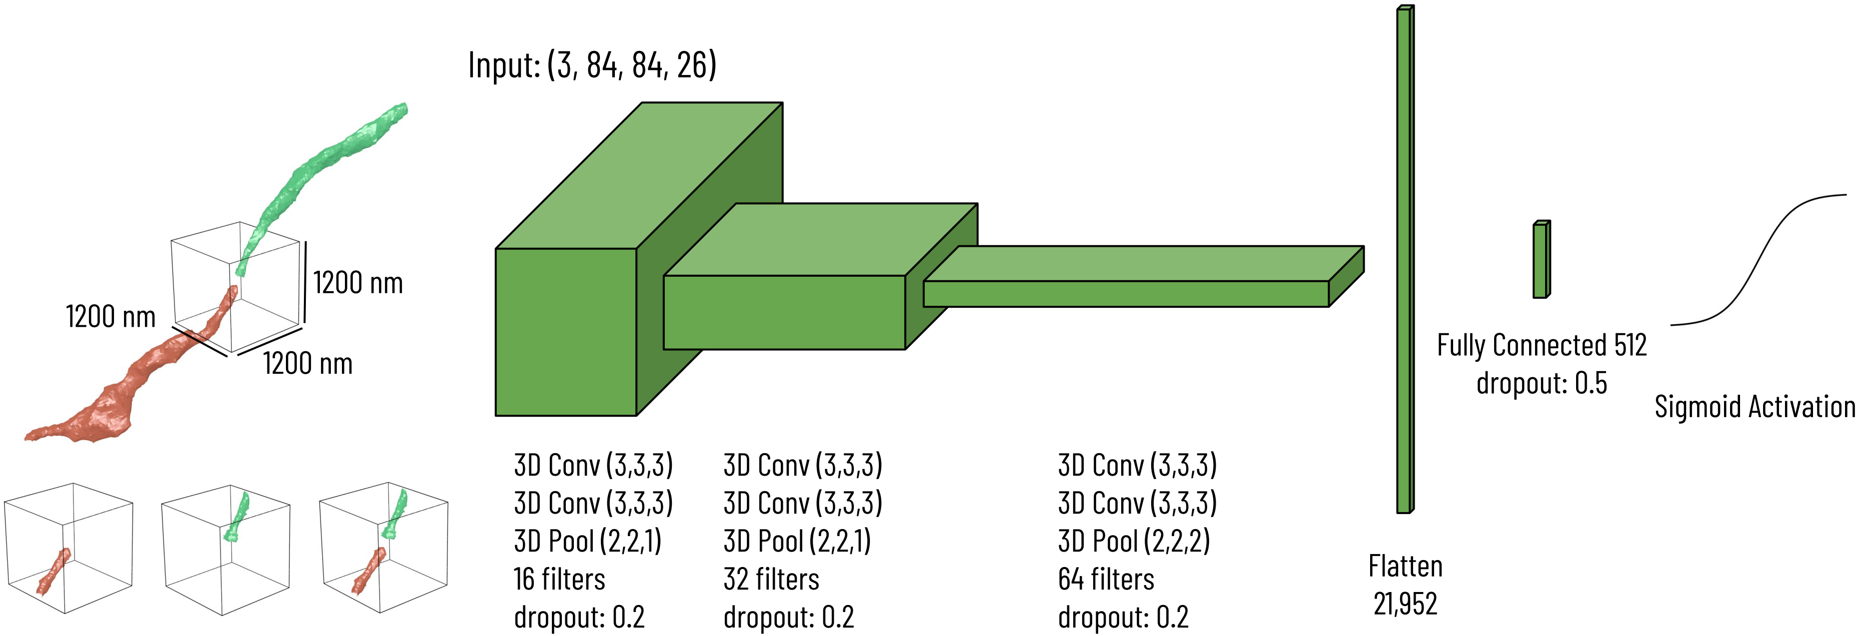

Parameters

Parameters	**Values**
Loss Function	Mean Squared Error
Optimizer	SGD with Nesterov Momentum
Momentum	0.9
Initial Learning Rate	0.01
Decay Rate	$5 * 10^{-8}$
Activation	LeakyReLU ($\alpha = 0.001$)
Kernel Sizes	$3 \times 3 \times 3$
Filter Sizes	$16 \to 32 \to 64$

Table 1: Training parameters.

Architectures

Depth	Input Size	No. Parameters	Output Size	Accuracy	Precision	Recall
3	(3, 18, 52, 52)	1,101,553	(64, 3, 3, 3)	91.30	58.06	92.81
3	(3, 20, 60, 60)	2,313,969	(64, 4, 4, 4)	92.41	61.70	92.41
3	(3, 22, 68, 68)	4,312,817	(64, 5, 5, 5)	92.33	61.49	92.34
3	(3, 24, 76, 76)	7,294,705	(64, 6, 6, 6)	93.51	65.78	93.13
3	**(3, 26, 84, 84)**	**11,456,241**	**(64, 7, 7, 8)**	**95.38**	**74.43**	**92.34**
3	(3, 28, 92, 92)	16,994,033	(64, 8, 7, 8)	91.87	59.70	94.22
3	(3, 30, 100, 100)	24,104,689	(64, 9, 9, 9)	92.01	60.24	93.75
4	(3, 28, 92, 92)	1,404,913	(128, 2, 2, 2)	91.70	60.24	85.94
4	(3, 32, 108, 108)	2,650,097	(128, 3, 3, 3)	92.80	64.28	86.88

Table 2. The results of various network architectures trained on the Kasthuri data.

Independent of Image Data

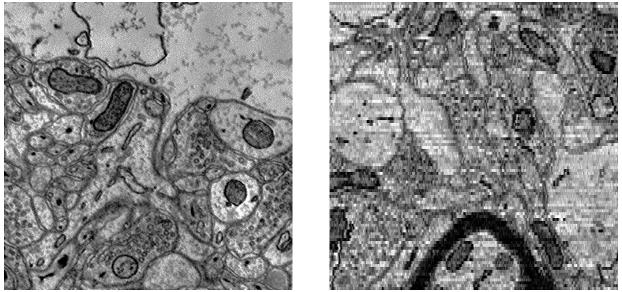

Kasthuri

Mouse

$6 \times 6 \times 30\ nm^3$ / vx

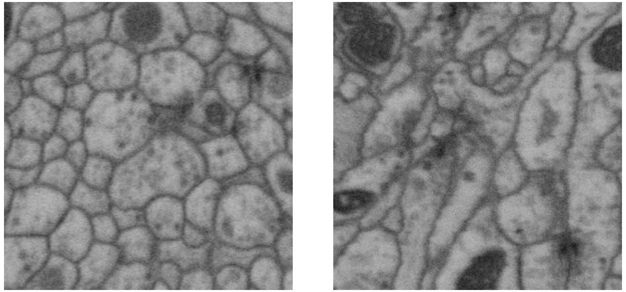

FlyEM

Drosophila melanogaster

$10 \times 10 \times 10\ nm^3$ / vx

Regions of Interest

Too small and there is not enough local context

Regions of Interest

Too large and unnecessary detail inhibits learning

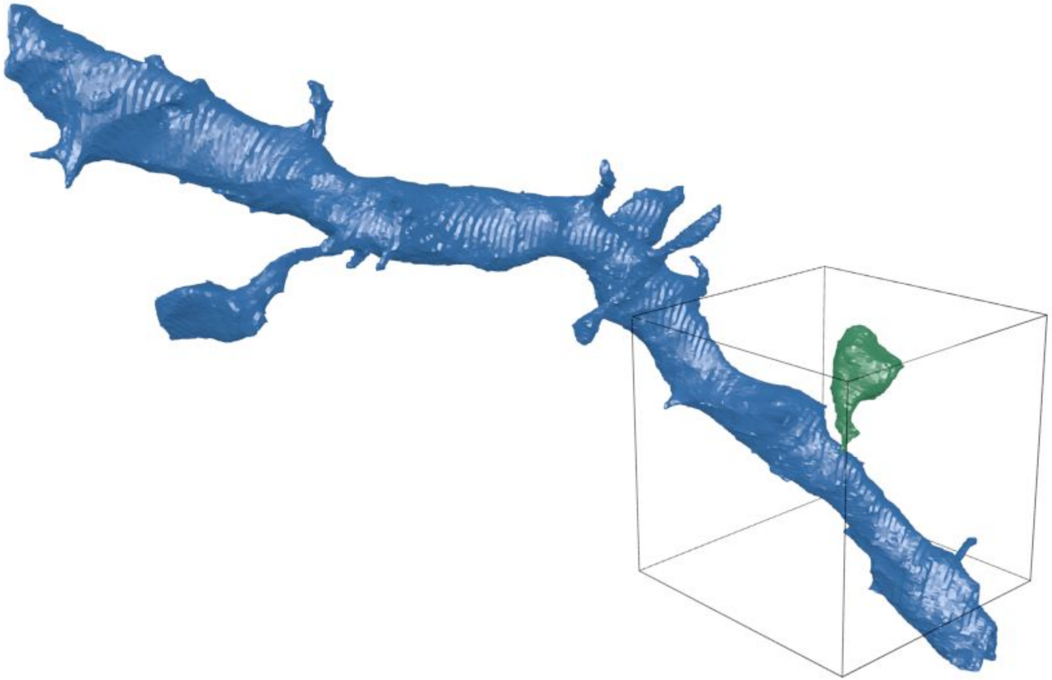

Regions of Interest

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

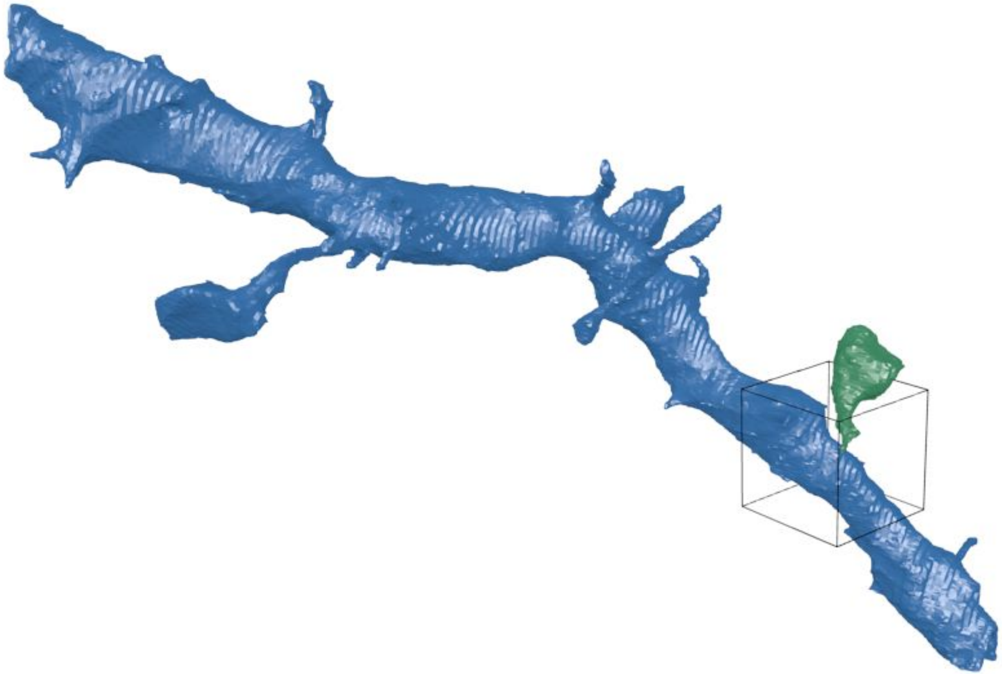

Input Examples

Should Merge

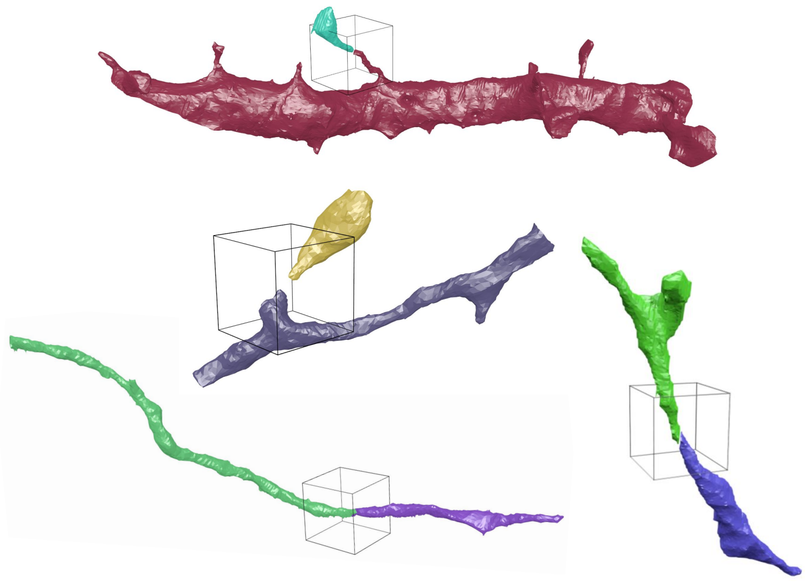

Input Examples

Should Merge

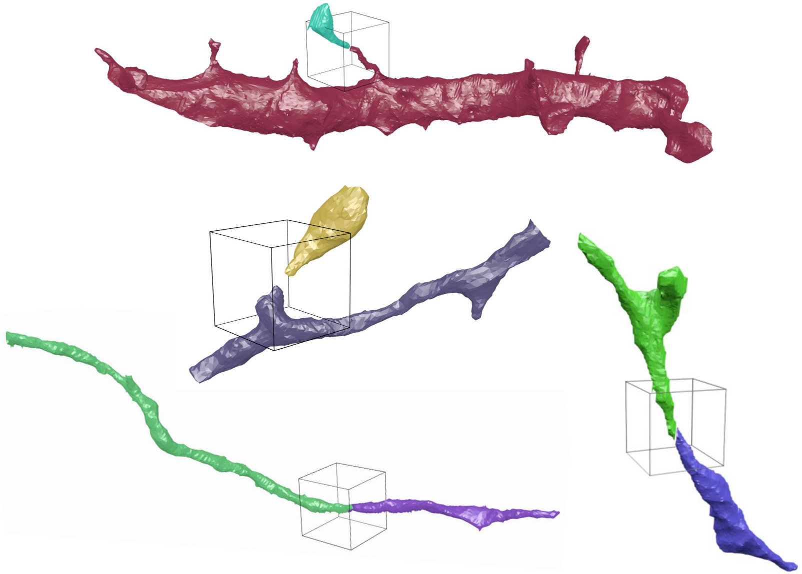

Should Not Merge

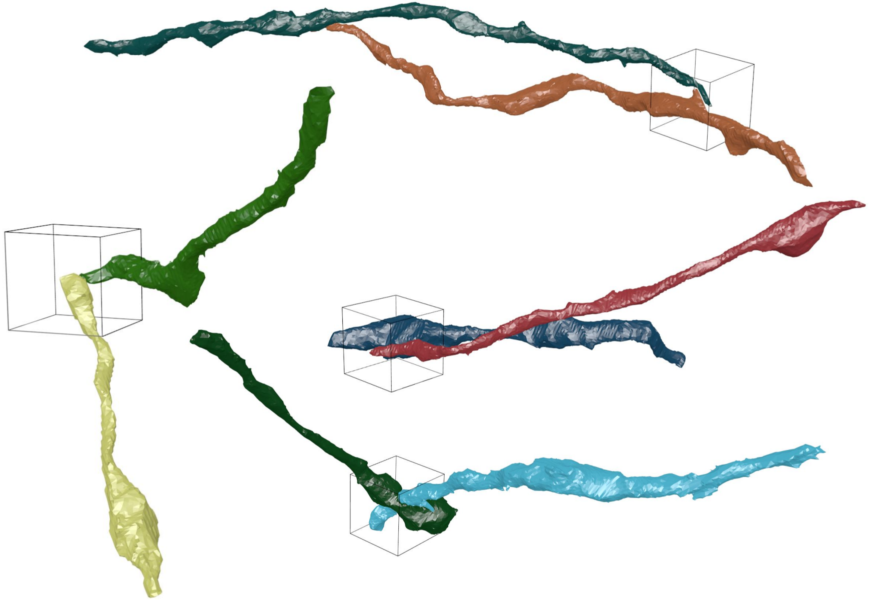

Goal: Partition the graph into individual neurons

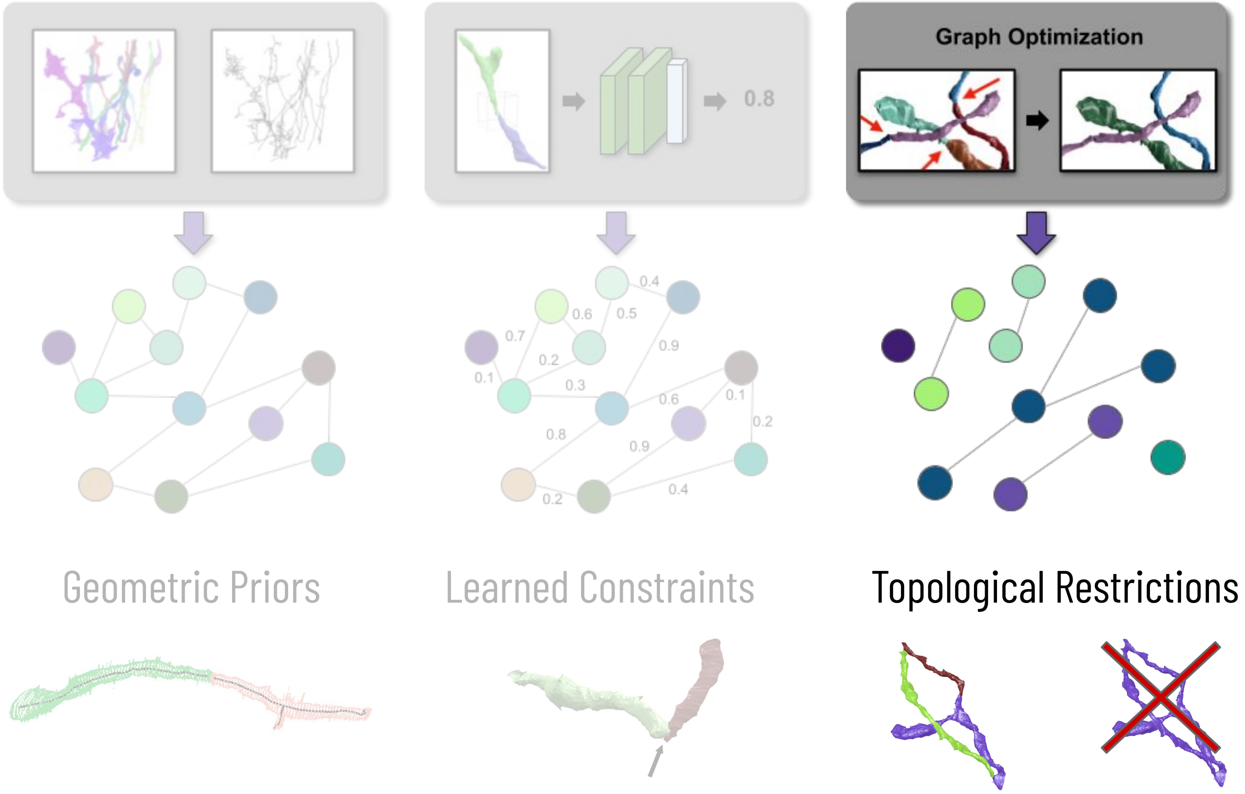

Multicut

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

Enforce topological constraints on neurons

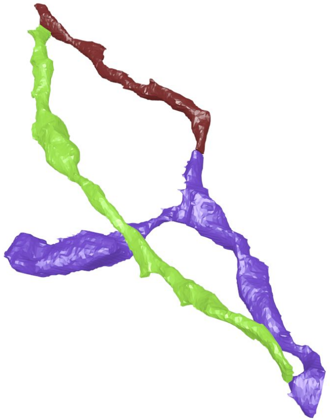

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

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

Enforce topological constraints on neurons

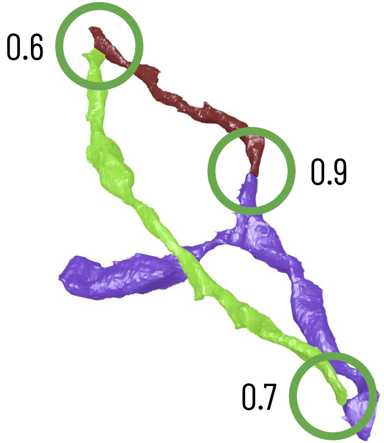

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

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

Enforce topological constraints on neurons

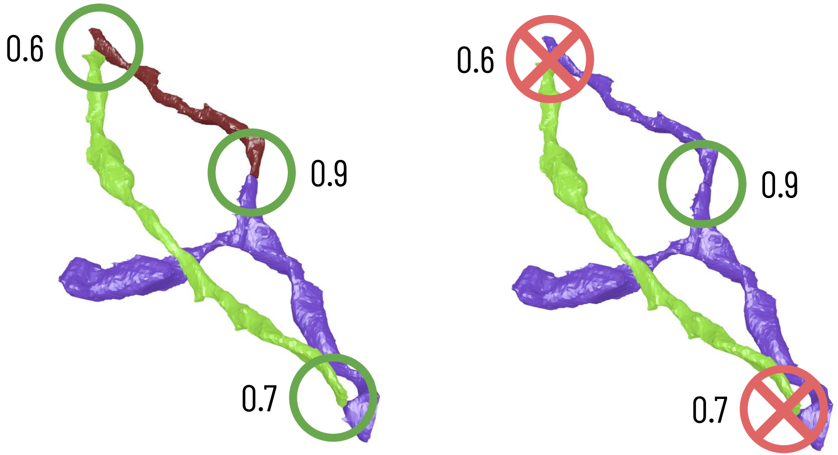

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

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

Enforce topological constraints on neurons

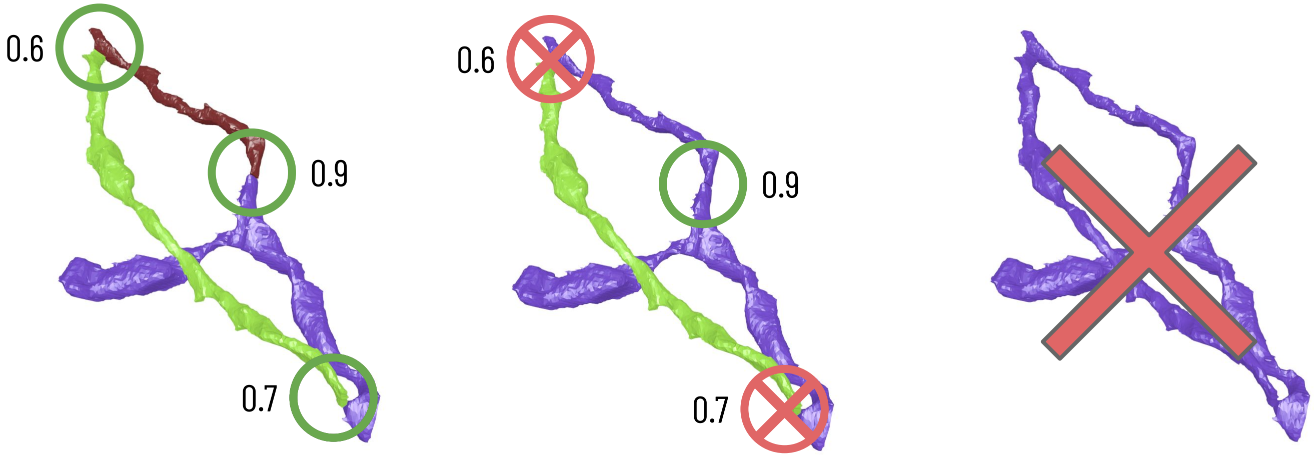

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

Guarantees a globally consistent solution

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

Guarantees a globally consistent solution

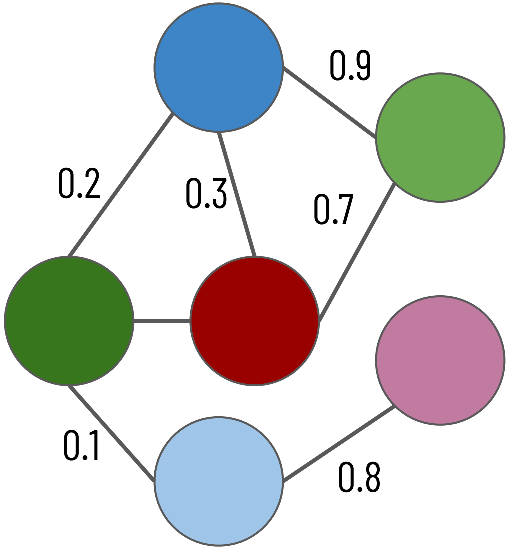

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

Guarantees a globally consistent solution

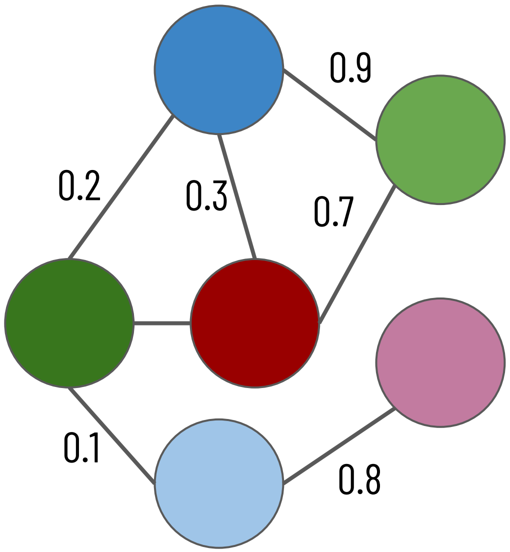

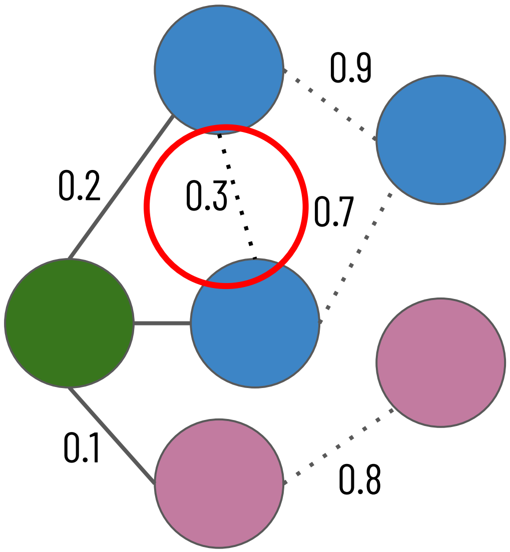

Multicut

Reformulate the segmentation problem as a multicut graph partitioning one

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

Lifted Multicut Extension

Costs associated with segmenting any pair of nodes

Lifted Multicut Extension

Costs associated with segmenting any pair of nodes

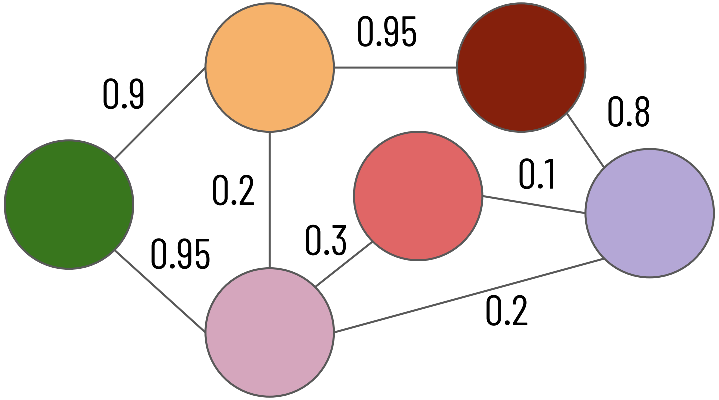

Lifted Multicut Extension

Costs associated with segmenting any pair of nodes

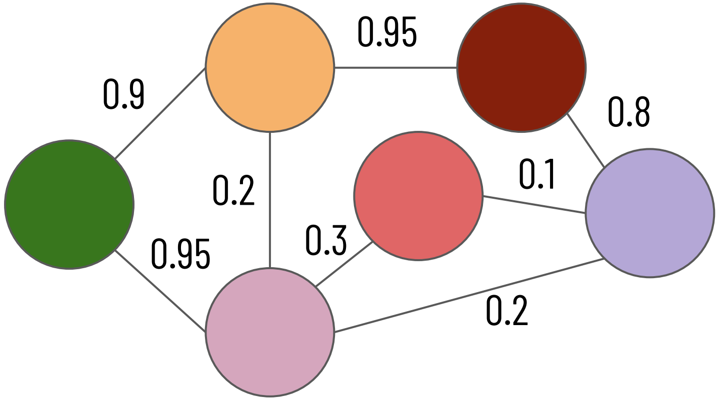

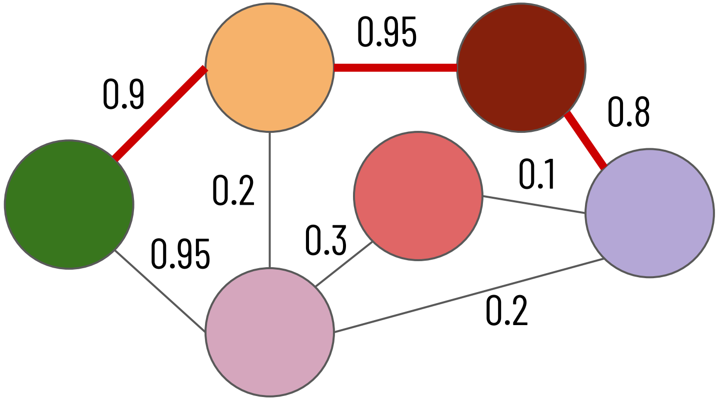

Lifted Multicut Extension

Costs associated with segmenting any pair of nodes

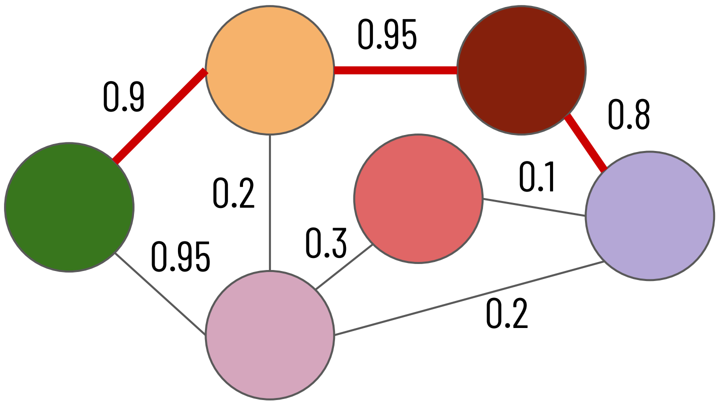

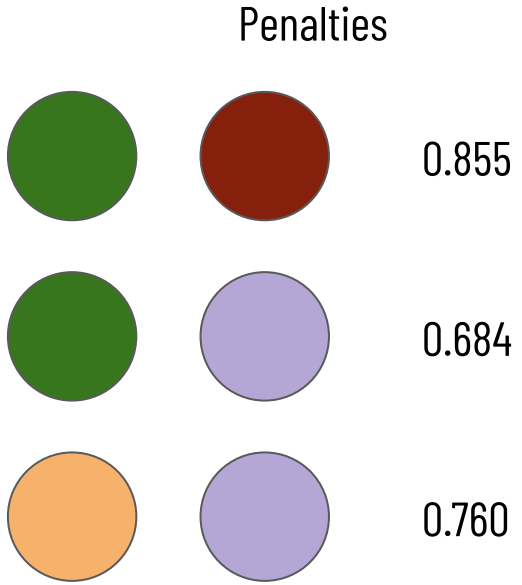

Results

Datasets

Training Data

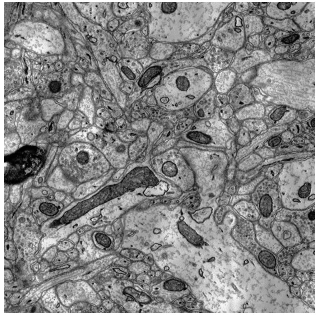

Kasthuri Vol. 1

Mouse

$6 \times 6 \times 30\ nm^3$ / vx

1335 x 1809 x 338 vx

$8.01 \times 10.85 \times 10.14\ \mu m^3$

Testing Data

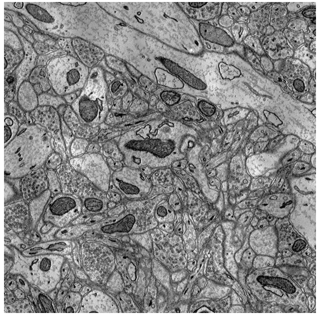

Kasthuri Vol. 2

Mouse

$6 \times 6 \times 30\ nm^3$ / vx

1336 x 1809 x 338 vx

$8.02 \times 10.85 \times 10.14\ \mu m^3$

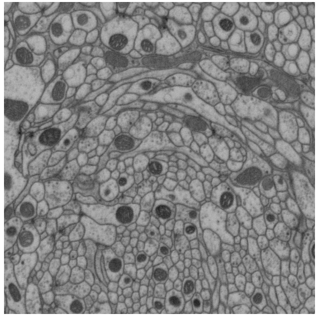

FlyEM Vol. 1

Drosophila melanogaster

$10 \times 10 \times 10\ nm^3$ / vx

999 x 999 x 998 vx

$9.99 \times 9.99 \times 9.98\ \mu m^3$

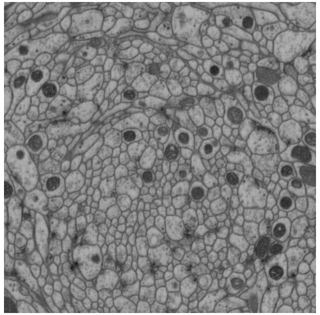

FlyEM Vol. 2

Drosophila melanogaster

$10 \times 10 \times 10\ nm^3$ / vx

999 x 999 x 999 vx

$9.99 \times 9.99 \times 9.99\ \mu m^3$

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

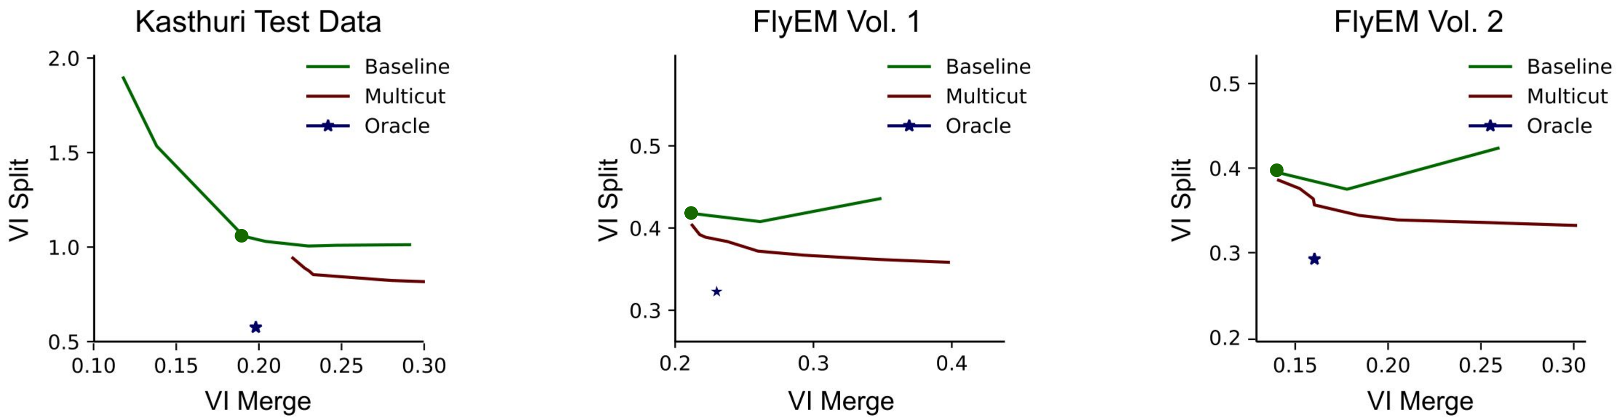

Baseline curve is generated by varying an agglomeration parameter in the Neuroproof algorithm

Green dot represents the input segmentation

Oracle correctly partitions the graph that we extract from the input segmentation

Variation of Information

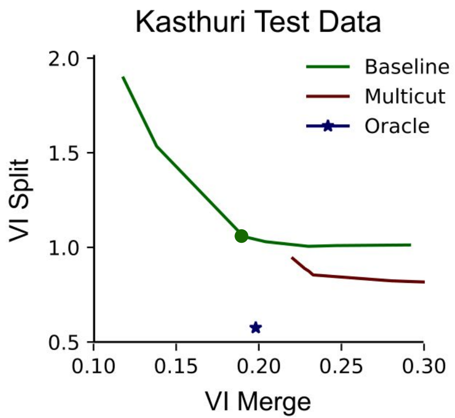

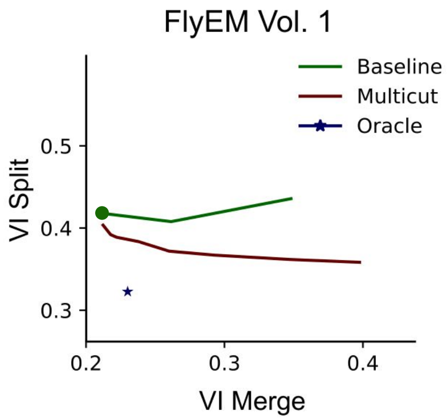

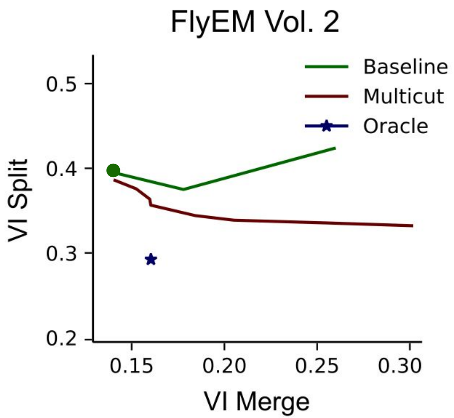

Total VI Improvement:

10.4% 8.9% 5.4%

Qualitative Results

Qualitative Results

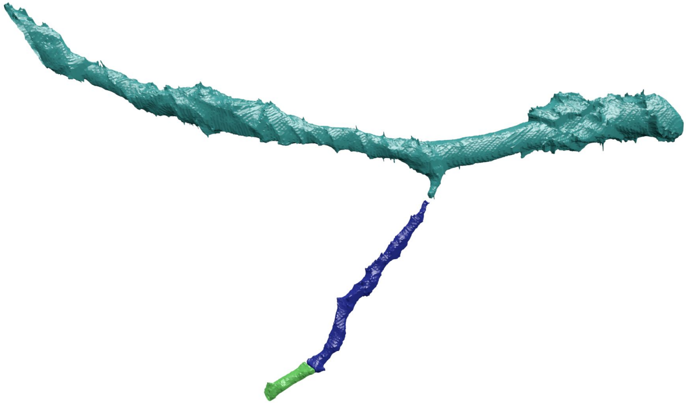

Qualitative Results

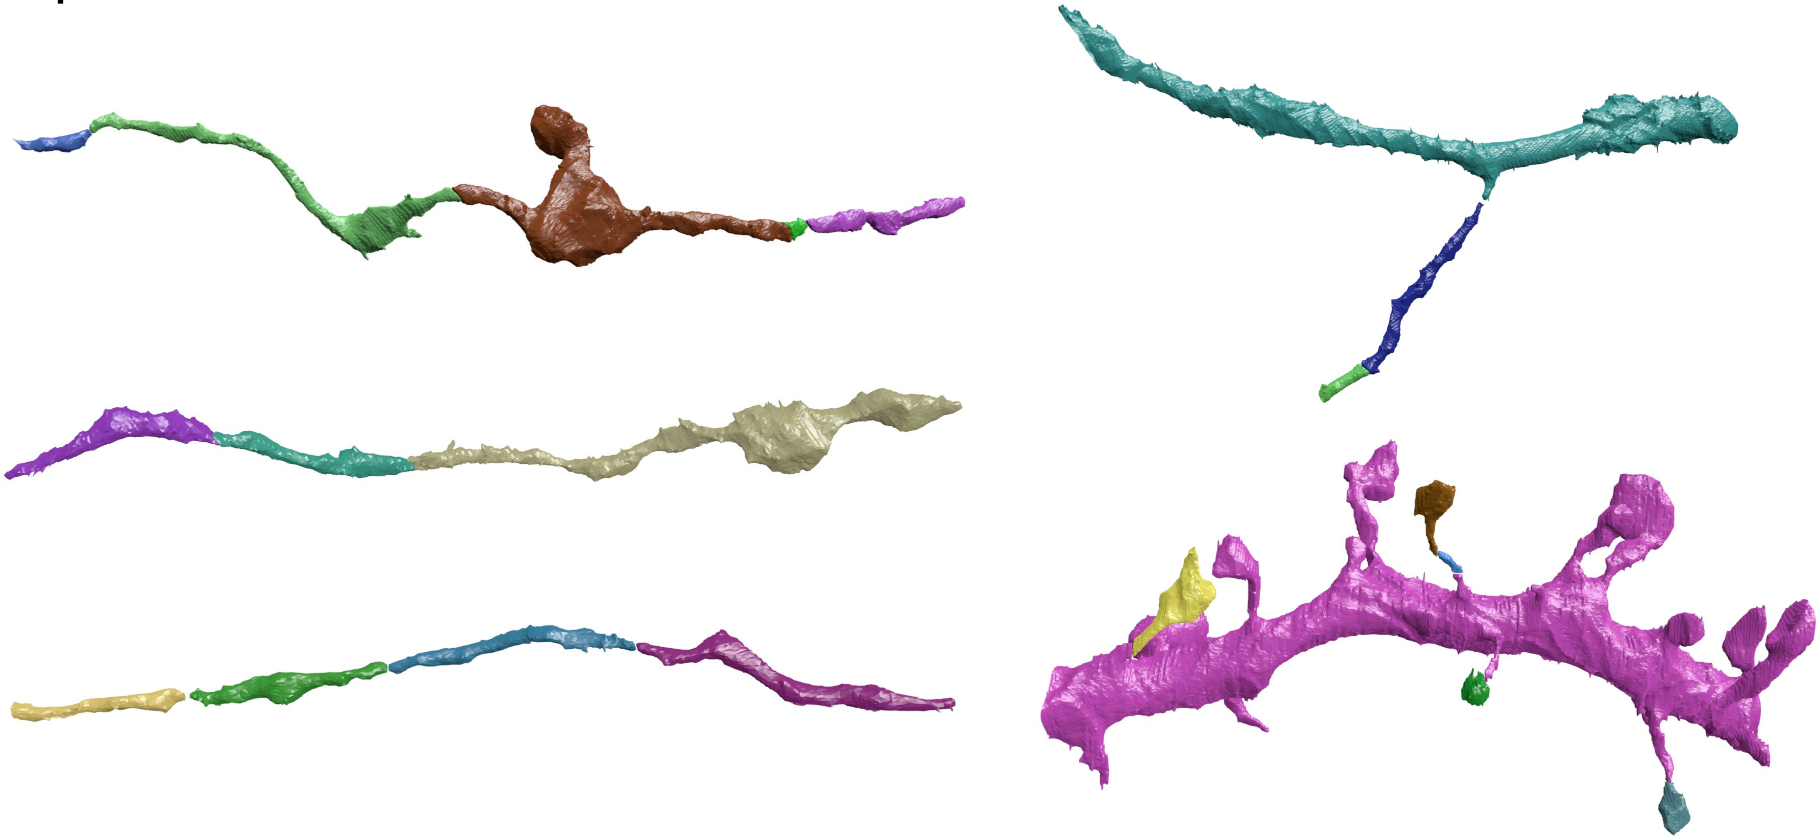

Failure Cases

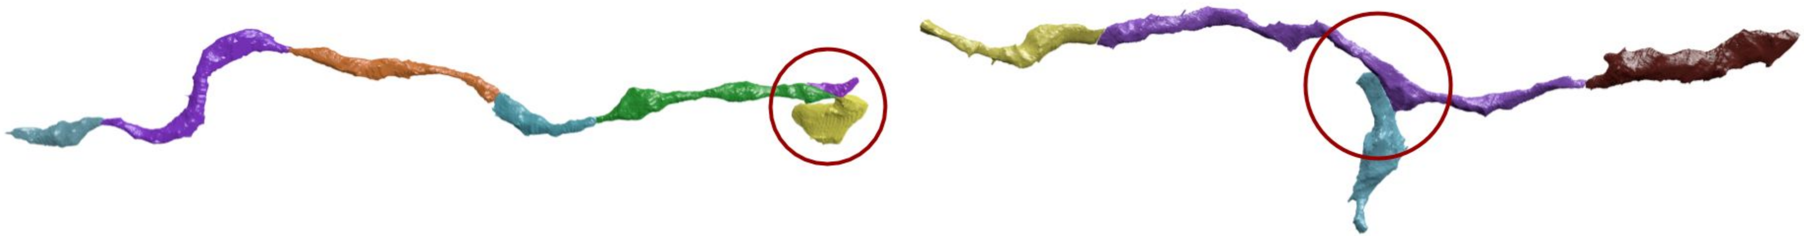

Failure Cases

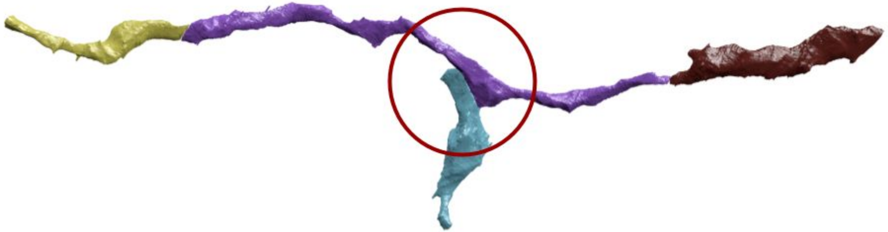

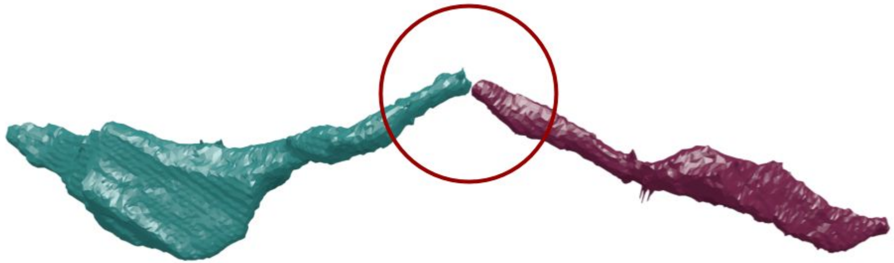

Failure Cases

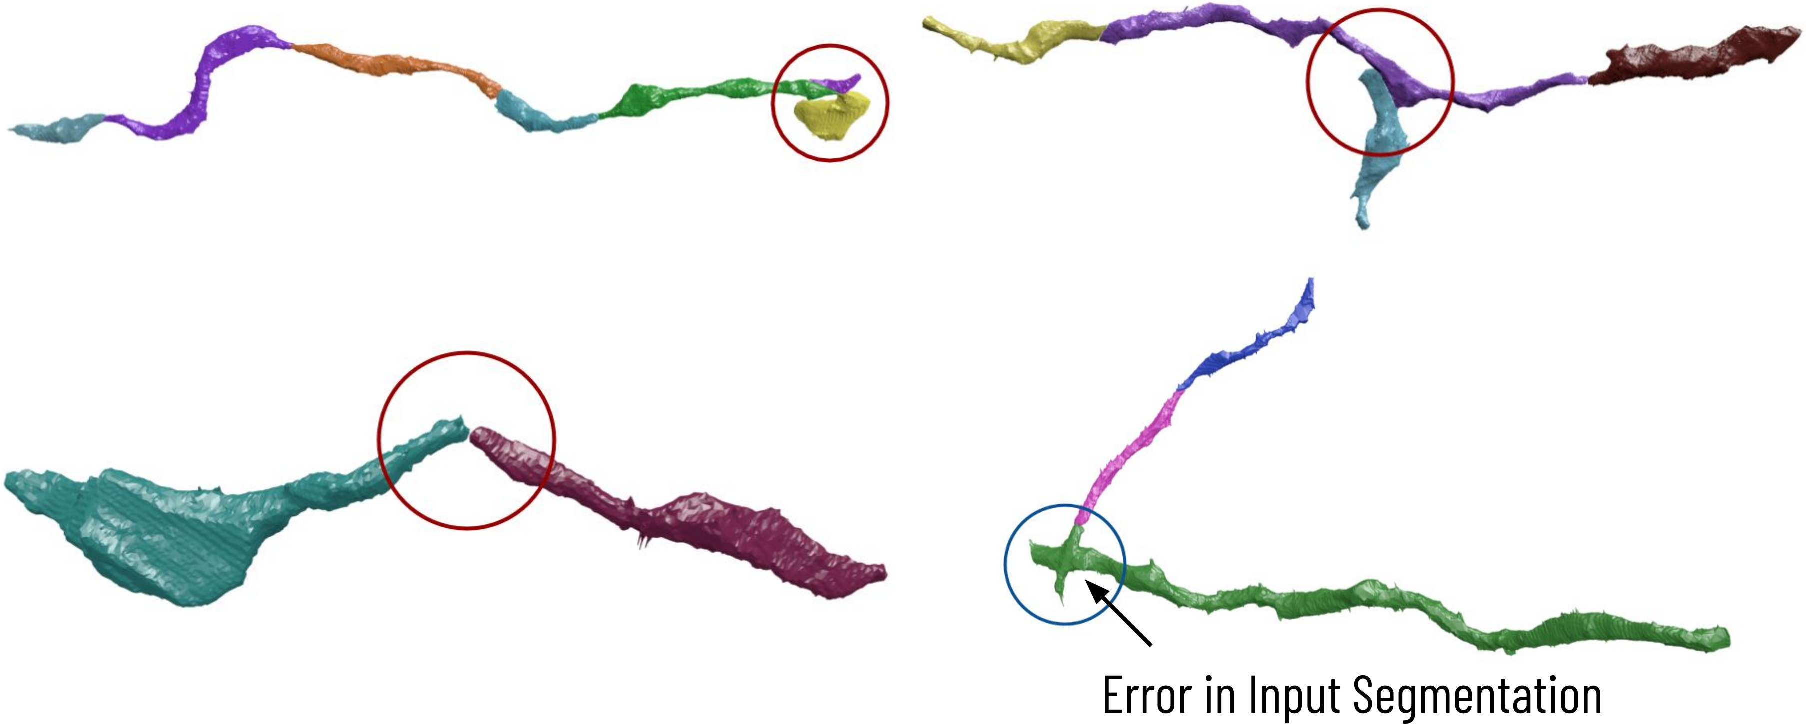

Graph Pruning

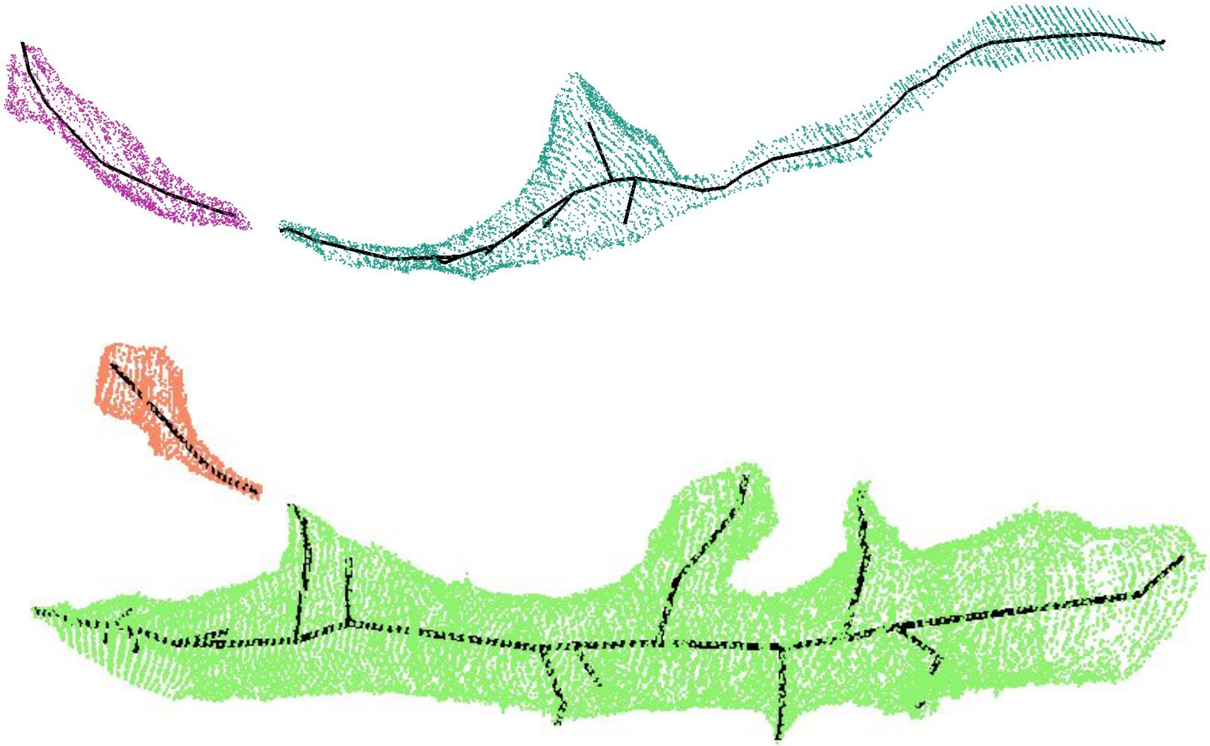

Failure Cases

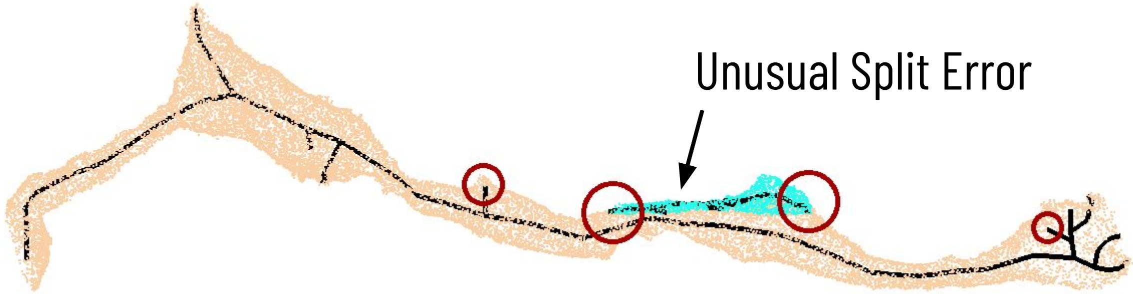

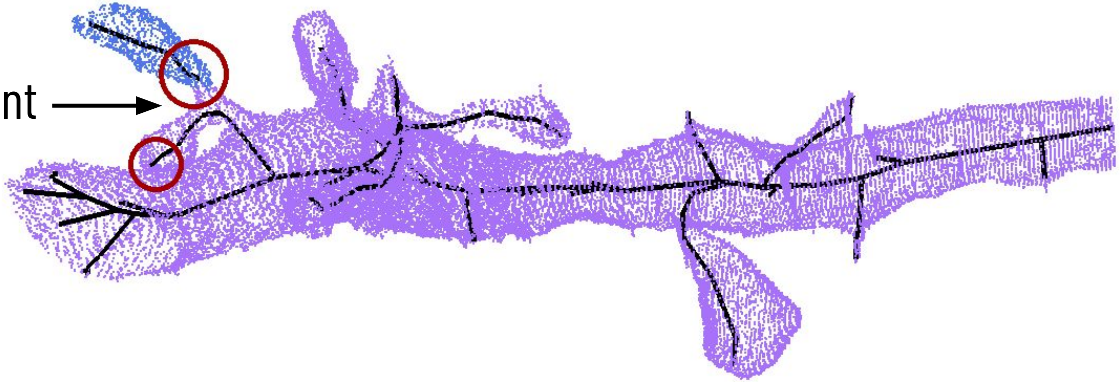

Graph Extraction Results

Table 1: The results of our graph pruning approach compared to the baseline graph with all adjacent regions. We show the number of true merge locations (e.g., 974) compared to total number of edges in the graph (e.g., 25,798) for each case. The number of missed splits corresponds to the number of split errors that our method misses compared to an adjacency matrix.

Dataset	Segment Adjacency	Skeleton Pruning	Missed Splits	Gained Edges
Kasthuri	974 / 25,798	764 / 6,218	307	97
FlyEM Vol. 1	304 / 15,949	212 / 4,578	105	13
FlyEM Vol. 2	298 / 17,614	197 / 4,366	120	19

CNN Results

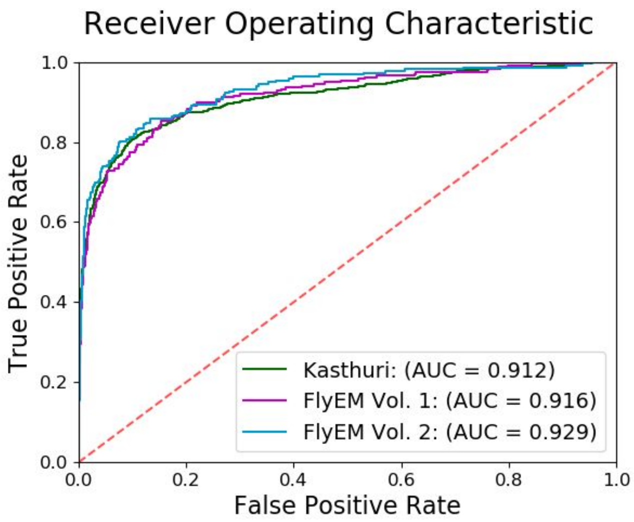

Accuracies:

Kasthuri	90.4%
FlyEM Vol. 1	94.4%
FlyEM Vol. 2	95.2%

Multicut Results

Table 2: Precision, recall, and accuracy changes between CNN only and CNN paired with graph-optimized reconstructions for the training and three test datasets. The combined method results in better precision and accuracy. The lifted multicut extension provides very slight improvements in recall and accuracy over these three datasets.

	Multicut			Lifted Multicut		
Dataset	**Δ Precision**	**Δ Recall**	**Δ Accuracy**	**Δ Precision**	**Δ Recall**	**Δ Accuracy**
Kasthuri	31.94%	-36.24%	0.71%	-1.01%	0.60%	0.02%
FlyEM Vol. 1	40.87%	-42.37%	1.26%	0.35%	0.85%	0.04%
FlyEM Vol. 2	27.80%	-44.95%	0.33%	0.54%	0.92%	0.04%

Running Times

Skeletonization: 0.56 seconds per segment on average

Graph Extraction: 31 seconds

CNN Inference: 124 seconds

Multicut: 37 seconds

~45 minutes

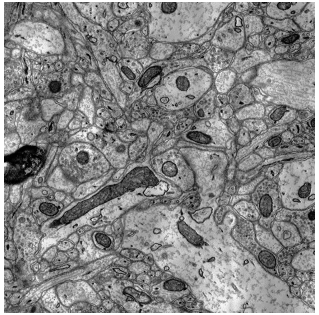

Kasthuri Vol. 1

Mouse

$6 \times 6 \times 30\ nm^3$ / vx

1335 x 1809 x 338

8.01μm x 10.85μm x 10.14μm

Connectomics Pipeline

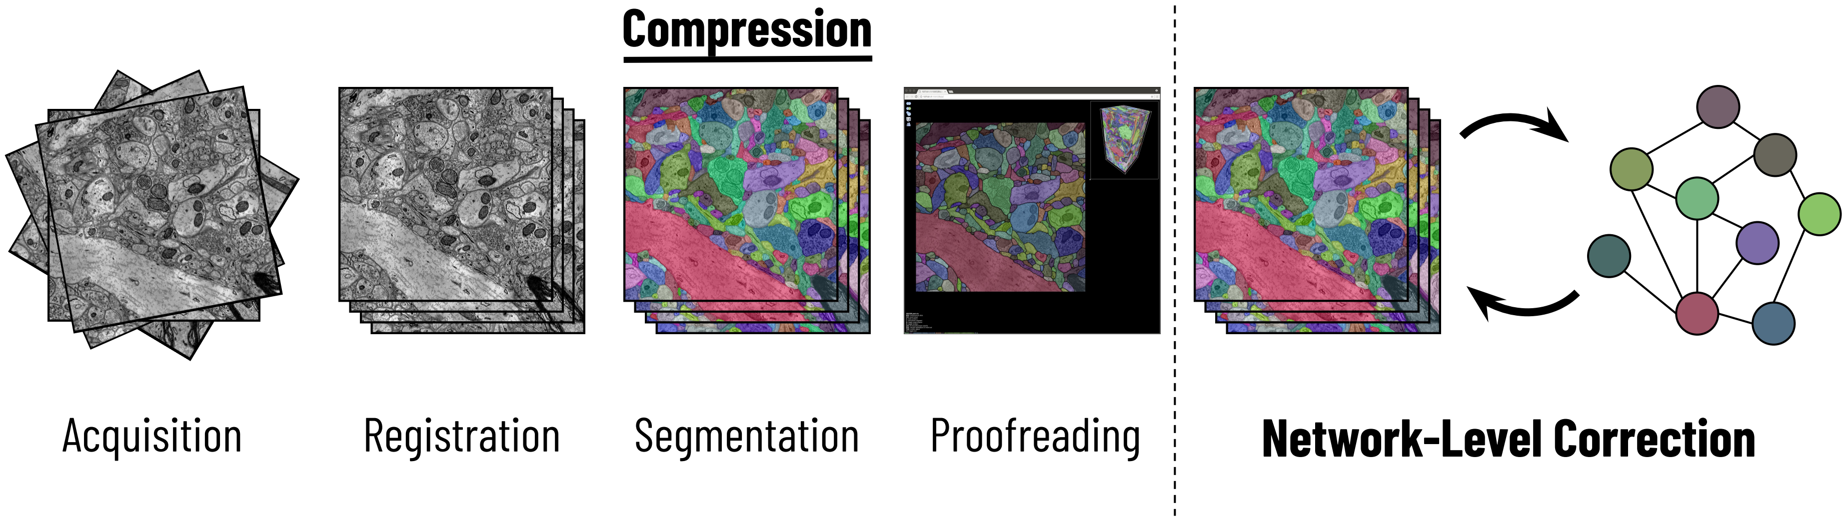

Connectomics Pipeline

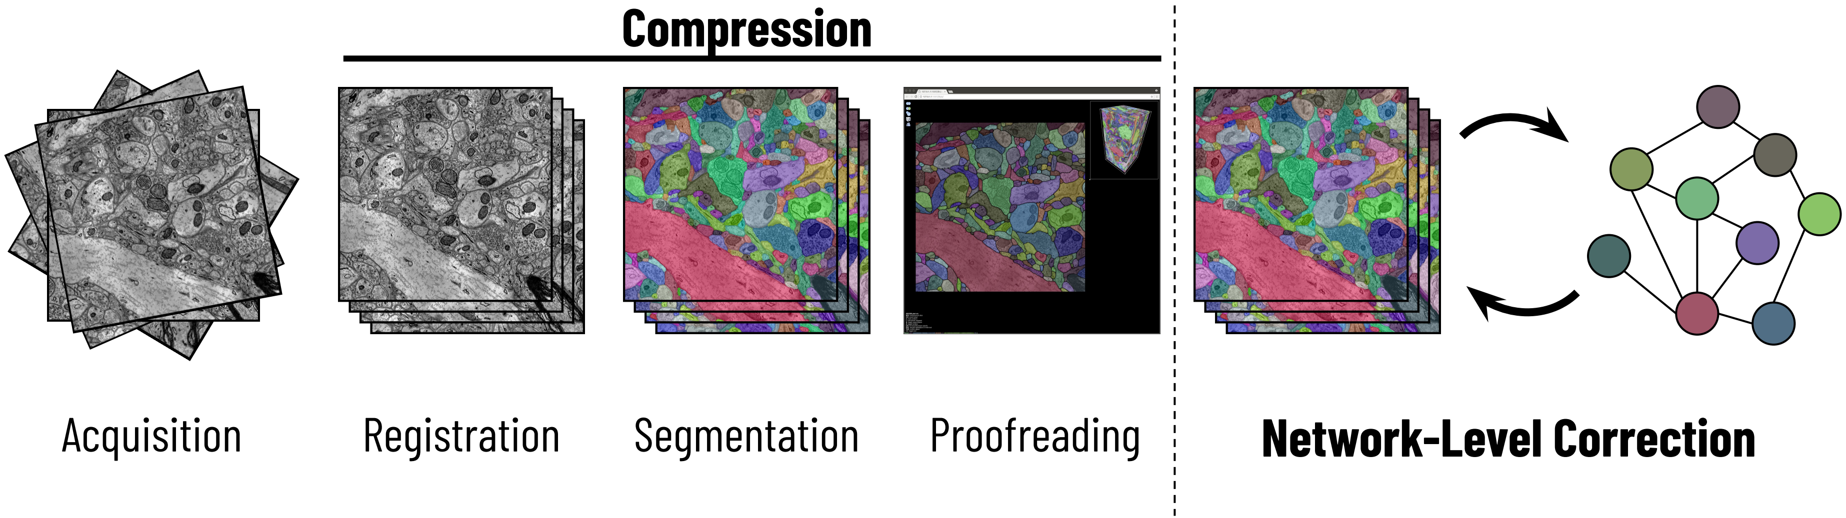

Compression Future Work

Specialized compression techniques for raw images (currently use JPEG 2000)

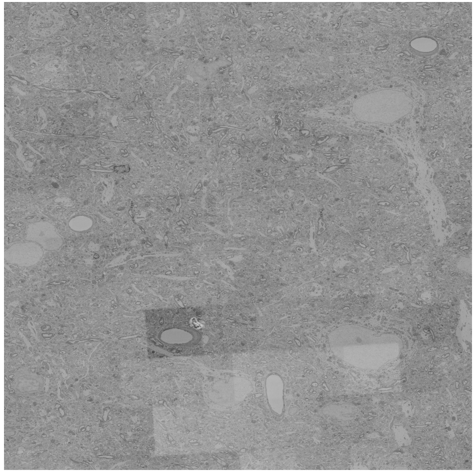

Compression Future Work

Specialized compression techniques for raw images (currently use JPEG 2000)

Use convolutional neural networks to improve compression of images

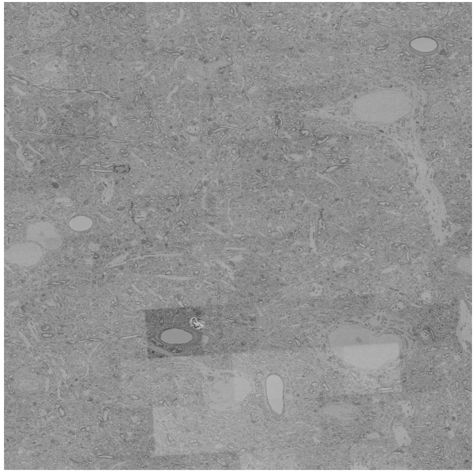

Compression Future Work

Specialized compression techniques for raw images (currently use JPEG 2000)

Use convolutional neural networks to improve compression of images

Add random access to Compresso for smoother real-time visual analysis of large datasets

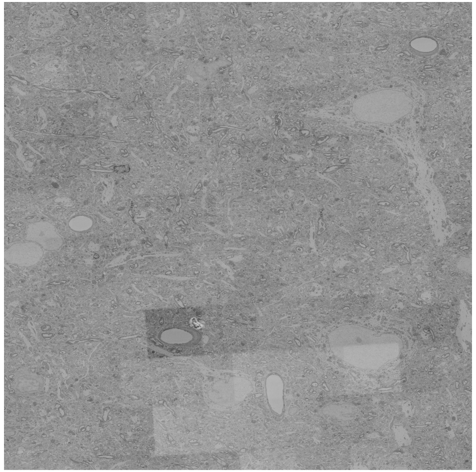

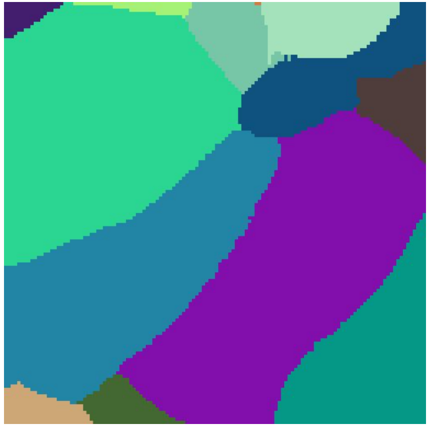

Connectomics Pipeline

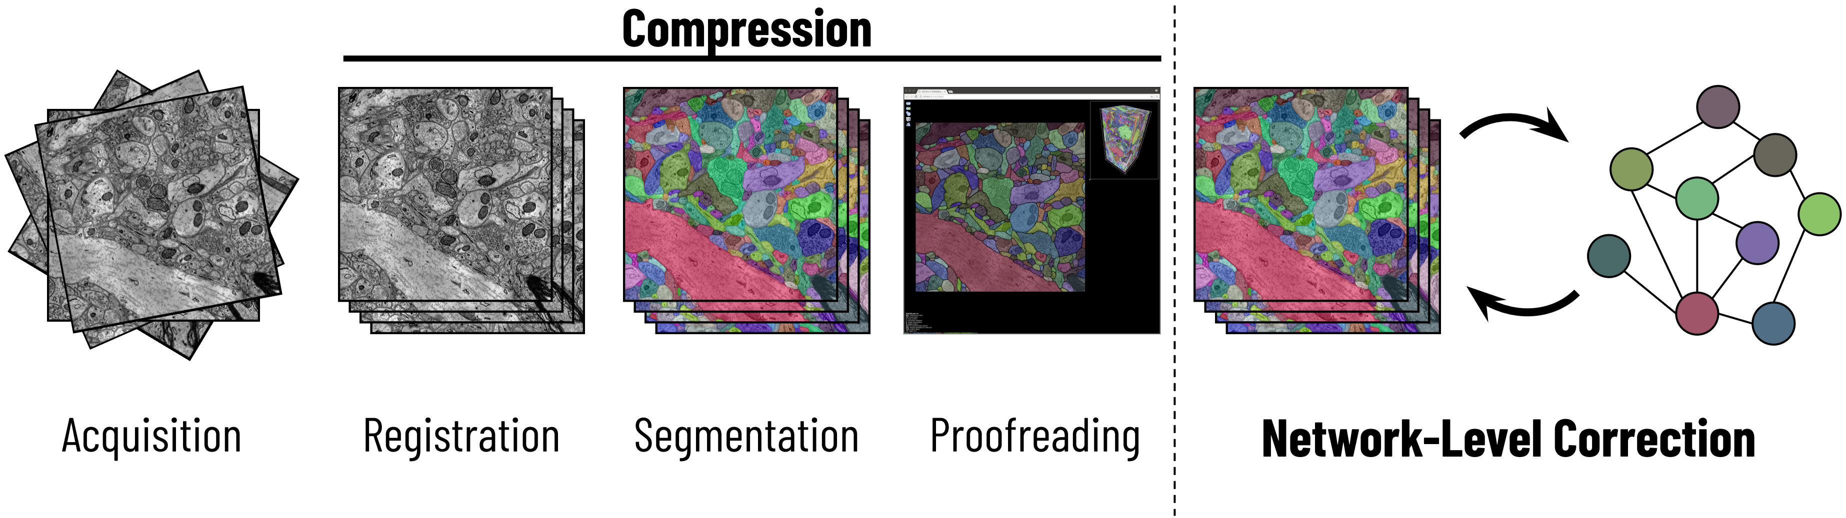

Additional Biological-Constraints

Use synaptic information to prevent dendrites and axons from merging

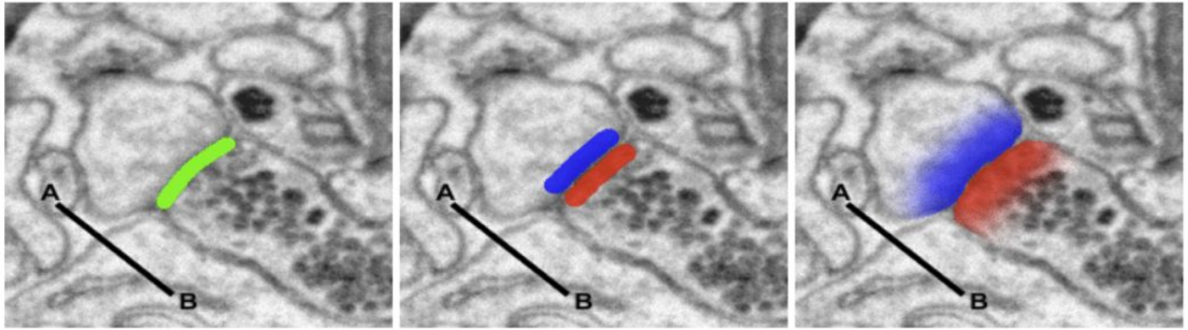

Additional Biological-Constraints

Use synaptic information to prevent dendrites and axons from merging

Classify neuron types to prevent inhibitory and excitatory neurons from merging

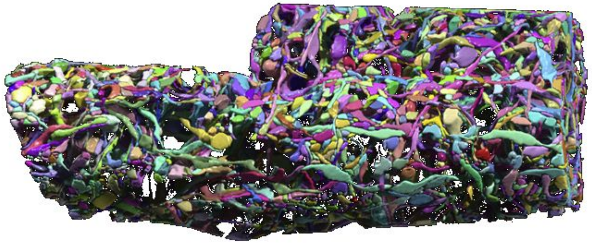

Excitatory Axons

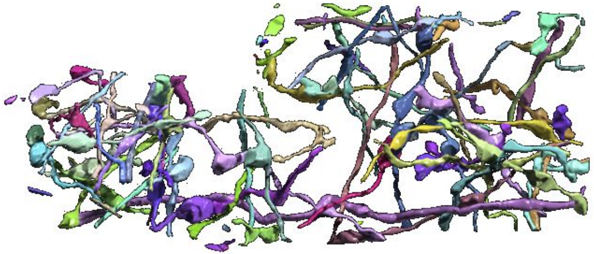

Inhibitory Axons

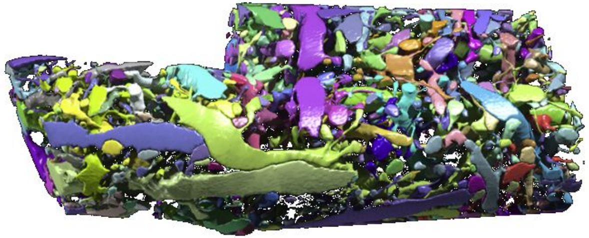

Dendrites

Address Merge Errors

Currently difficult because the number of potential split candidates grows quickly

Address Merge Errors

Currently difficult because the number of potential split candidates grows quickly

Use skeletons to quickly locate potential merge errors

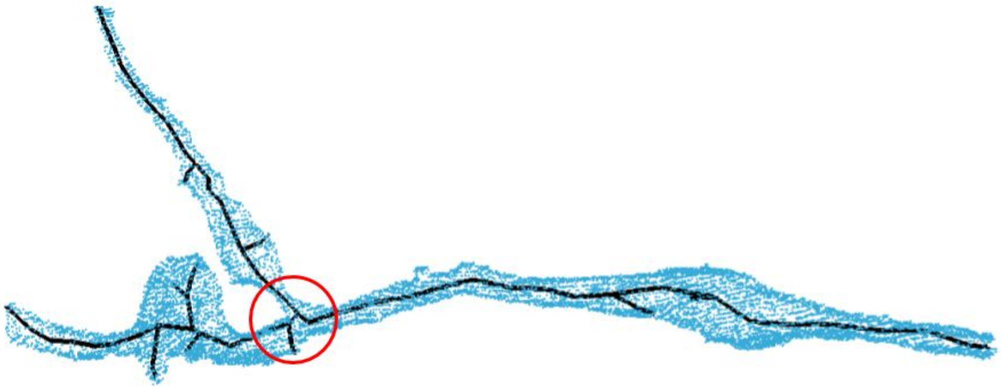

Address Merge Errors

Currently difficult because the number of potential split candidates grows quickly

Use skeletons to quickly locate potential merge errors

Divide segments with a watershed algorithm and use existing CNN

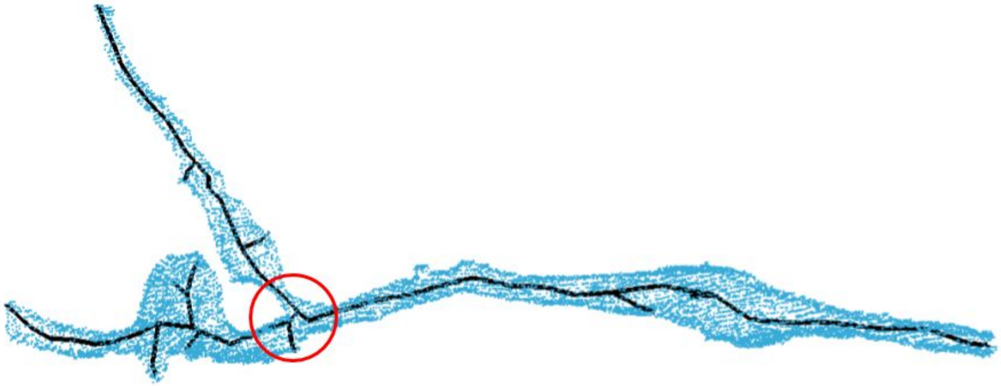

Thank you!

Questions?