CSE 40171:
Artificial Intelligence

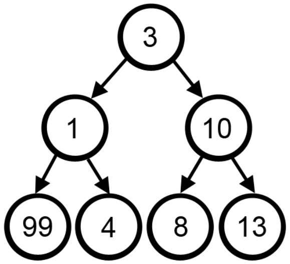

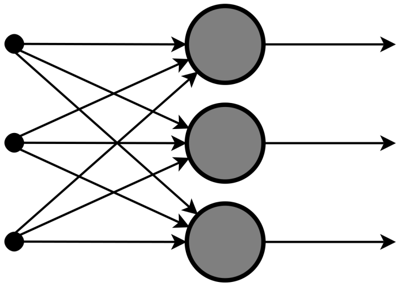

Connectomics: Classical Approaches to the Segmentation of Neural Volumes

Homework #5 has been released

It is due at 11:59PM on 11/13

C. elegans Connectome

Start Small: 302 neurons

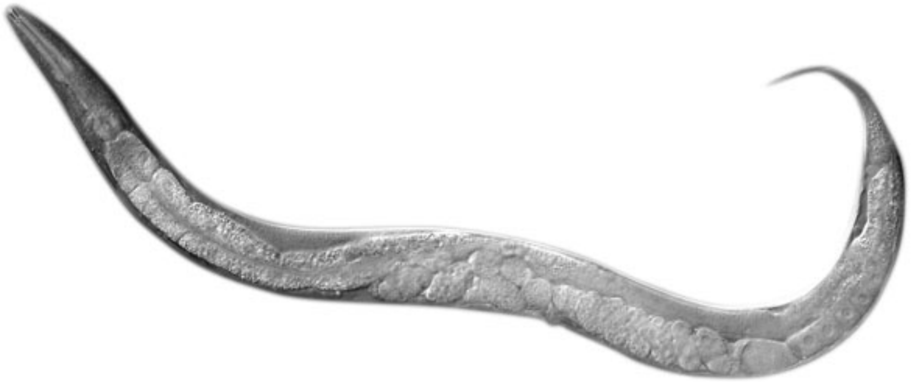

Cut 50nm cross-sections of the worm, and trace the synapses

C. elegans Connectome

12+ year manual effort (1970s - 1980s)
*Computer Vision was in its infancy

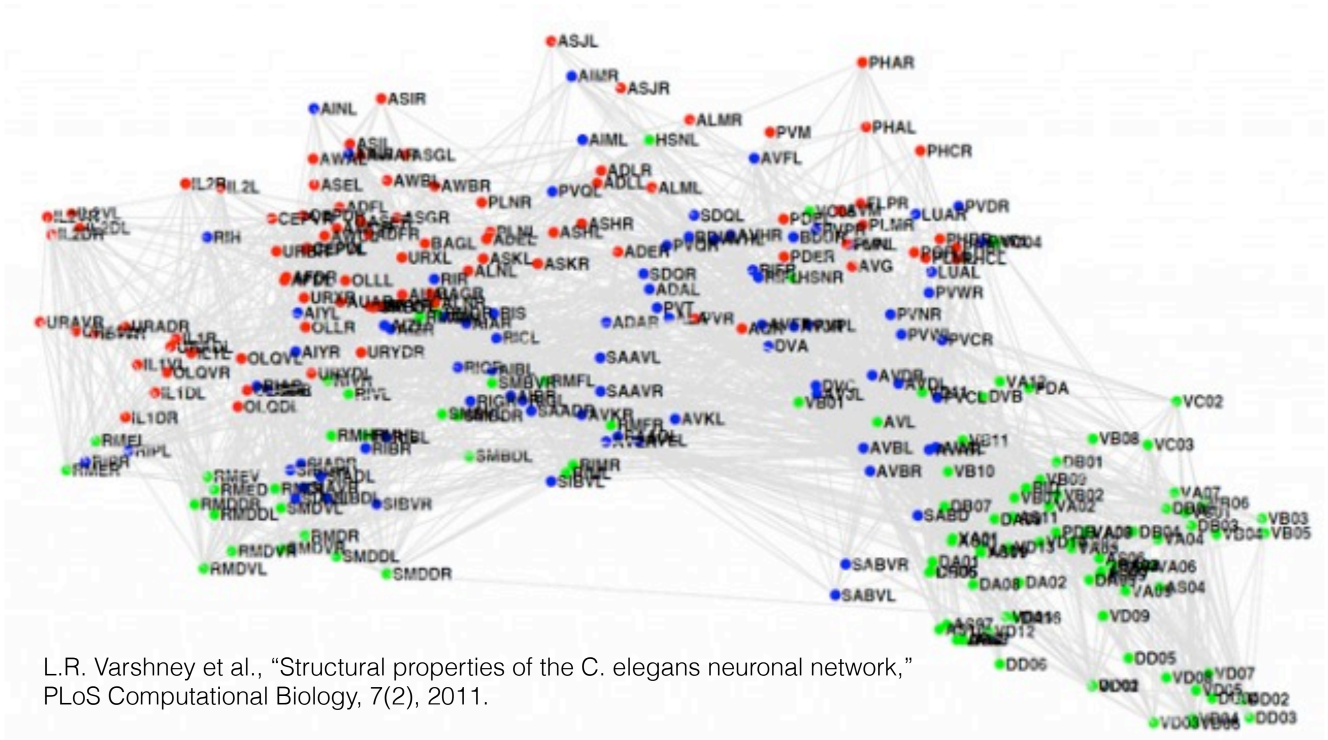

L.R. Varshney et al., “Structural properties of the C. elegans neuronal network,”
PLoS Computational Biology, 7(2), 2011.

Drosophila Connectome

- Approximately 100,000 Neurons
- Approximately 10,000,000 Synapses

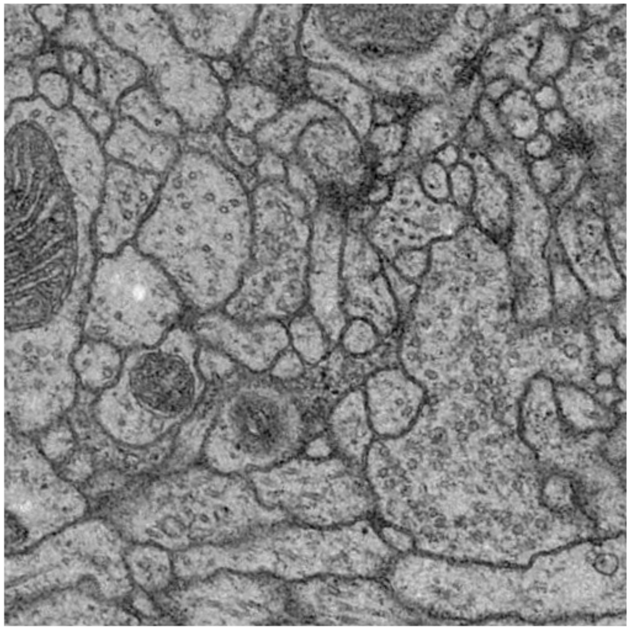

serial section Transmission Electron Microscopy (ssTEM) data set of the Drosophila first instar larva ventral nerve cord (VNC)

Albert Cardona, Stephan Saalfeld, Stephan Preibisch, Benjamin Schmid, Anchi Cheng, Jim Pulokas, Pavel Tomancak and Volker Hartenstein (10, 2010), "An Integrated Micro- and Macroarchitectural Analysis of the Drosophila Brain by Computer-Assisted Serial Section Electron Microscopy", PLoS Biol (Public Library of Science) 8 (10)

Rat Connectome

- Tens of millions of neurons and billions of connections between them
- Petabytes of data
- Cannot do this by hand: we need computer vision

Dense Segmentation and Reconstruction

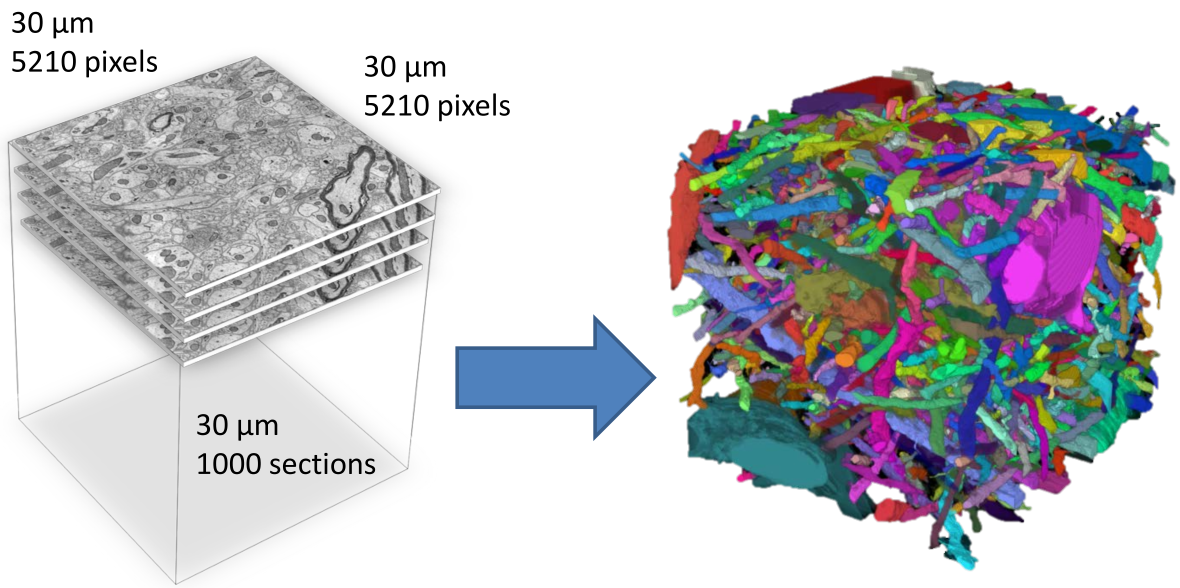

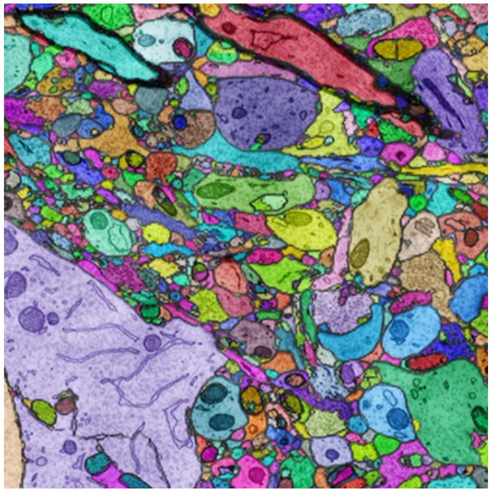

Software tools for dense reconstruction

Package	**Method**	**Link**
Randomer Forests	Decision Forests	http://ttomita.github.io/RandomerForest
Gala	Active Learning	https://github.com/janelia-flyem/gala
VESICLE	Deep Learning	http://openconnecto.me/vesicle
Synapse Segmenter	Context Features + Adaboost	http://cvlab.epfl.ch/software/synapse
ATMA	3D Pixel Features + Random Forests	https://github.com/RWalecki/ATMA
ZNN	CNN	https://github.com/seung-lab/znn-release
PRIM	CRF	http://github.com/funkey/prim
ilastik	Random Forests	http://ilastik.org/
Rhoana	CNN	https://github.com/Rhoana
FFN	CNN	https://github.com/google/ffn

Challenges for Computer Vision

- Serial-section electron microscopy (ssEM)
  - 2D physical resolution of the x-y dimensions is one order of magnitude finer than the z dimension
    - irregular spatial discontinuities
  - Neurons branch, merge, originate or terminate anywhere
  - Spines can be as thin as the distance between sections
  - Inconsistencies in staining and cutting of tissue

Standard Datasets

ISBI 2012 Challenge (http://brainiac2.mit.edu/isbi_challenge/)

- A full stack of Drosophila EM slices is used to train machine learning algorithms for the purpose of automatic segmentation of neural structures

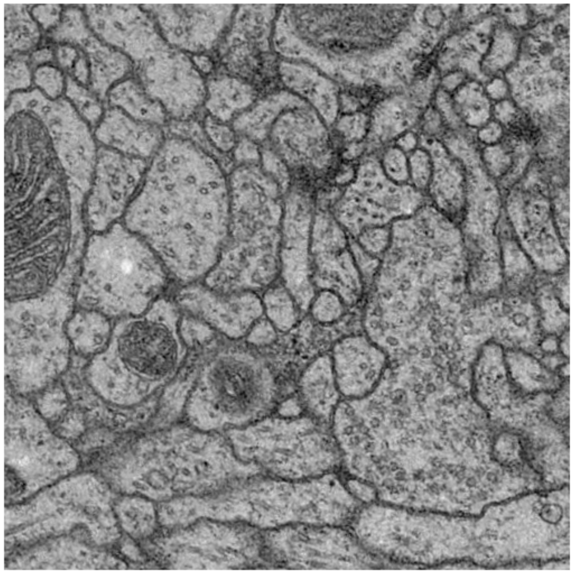

ssTEM image

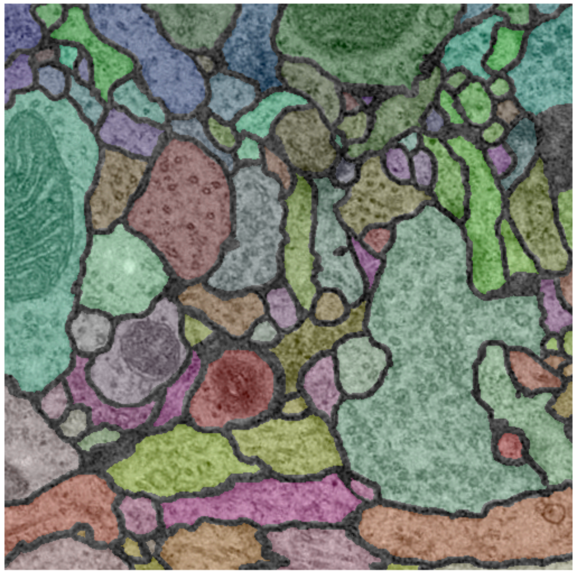

corresponding segmentation

The microcube measures 2 x 2 x 1.5 microns approx.,
with a resolution of 4x4x50 nm/pixel.

ISBI 2012

30 Training and 30 testing Samples

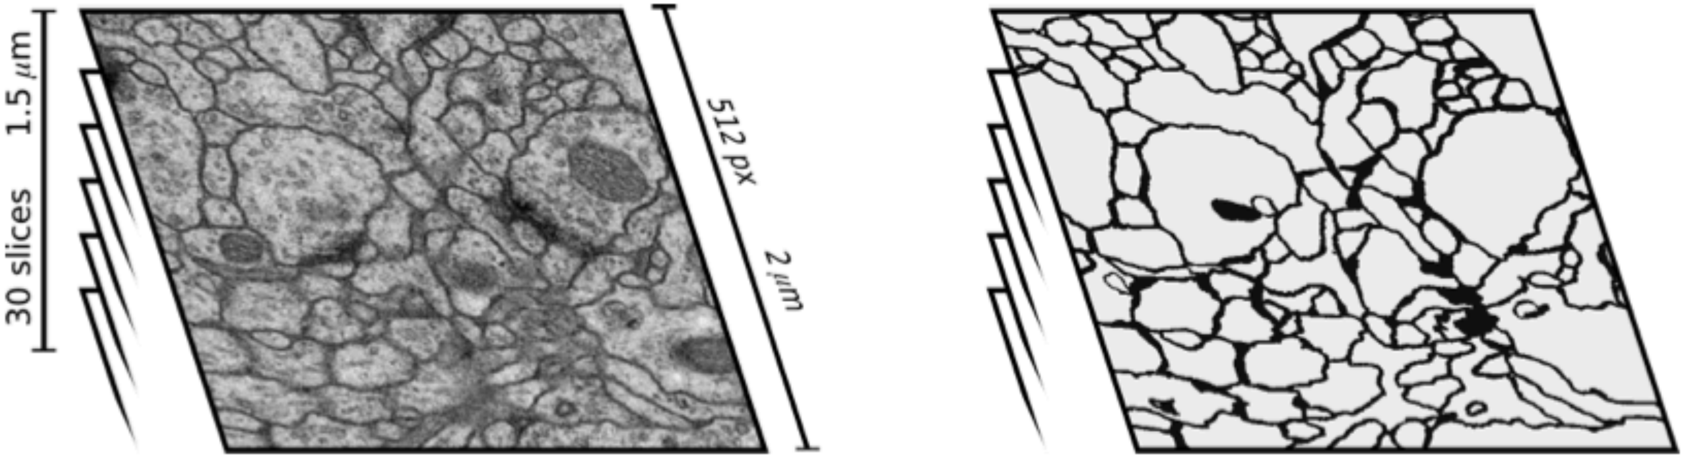

ISBI 2012

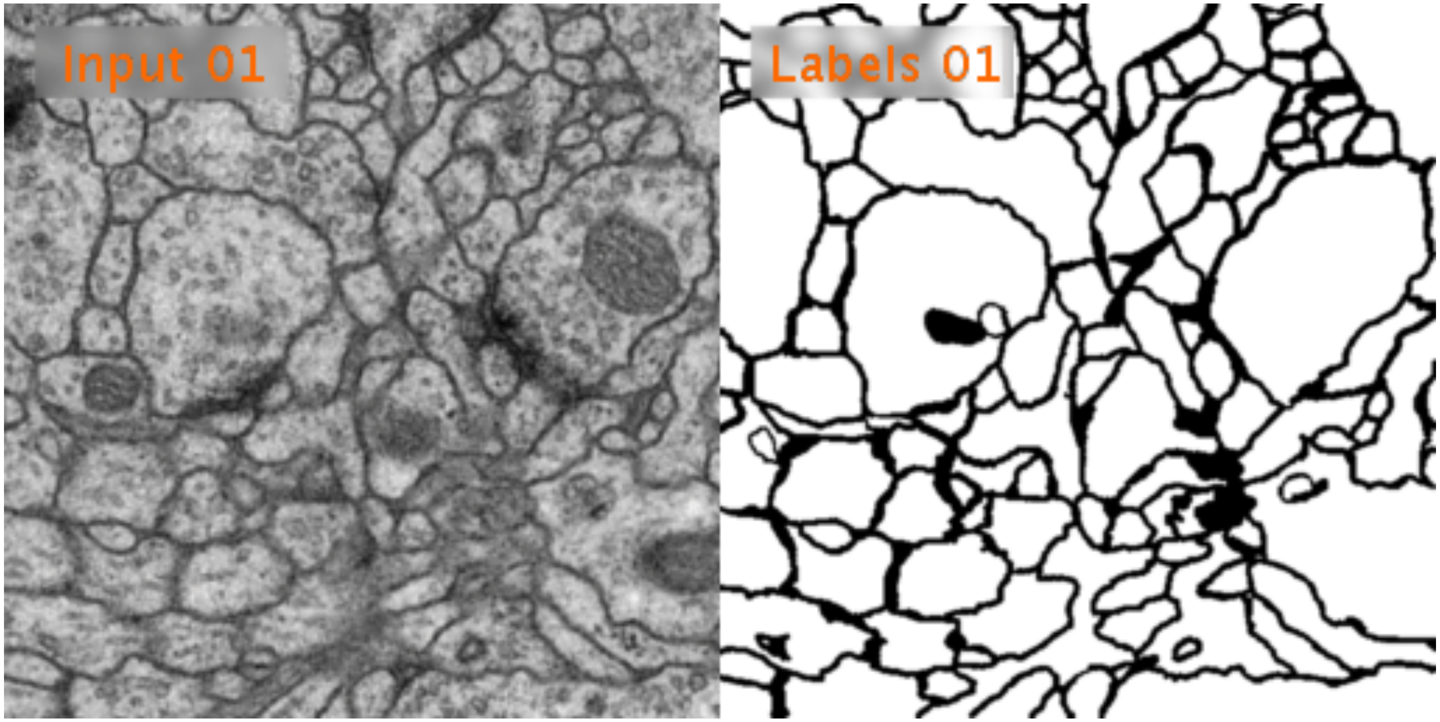

Split and Merge Errors

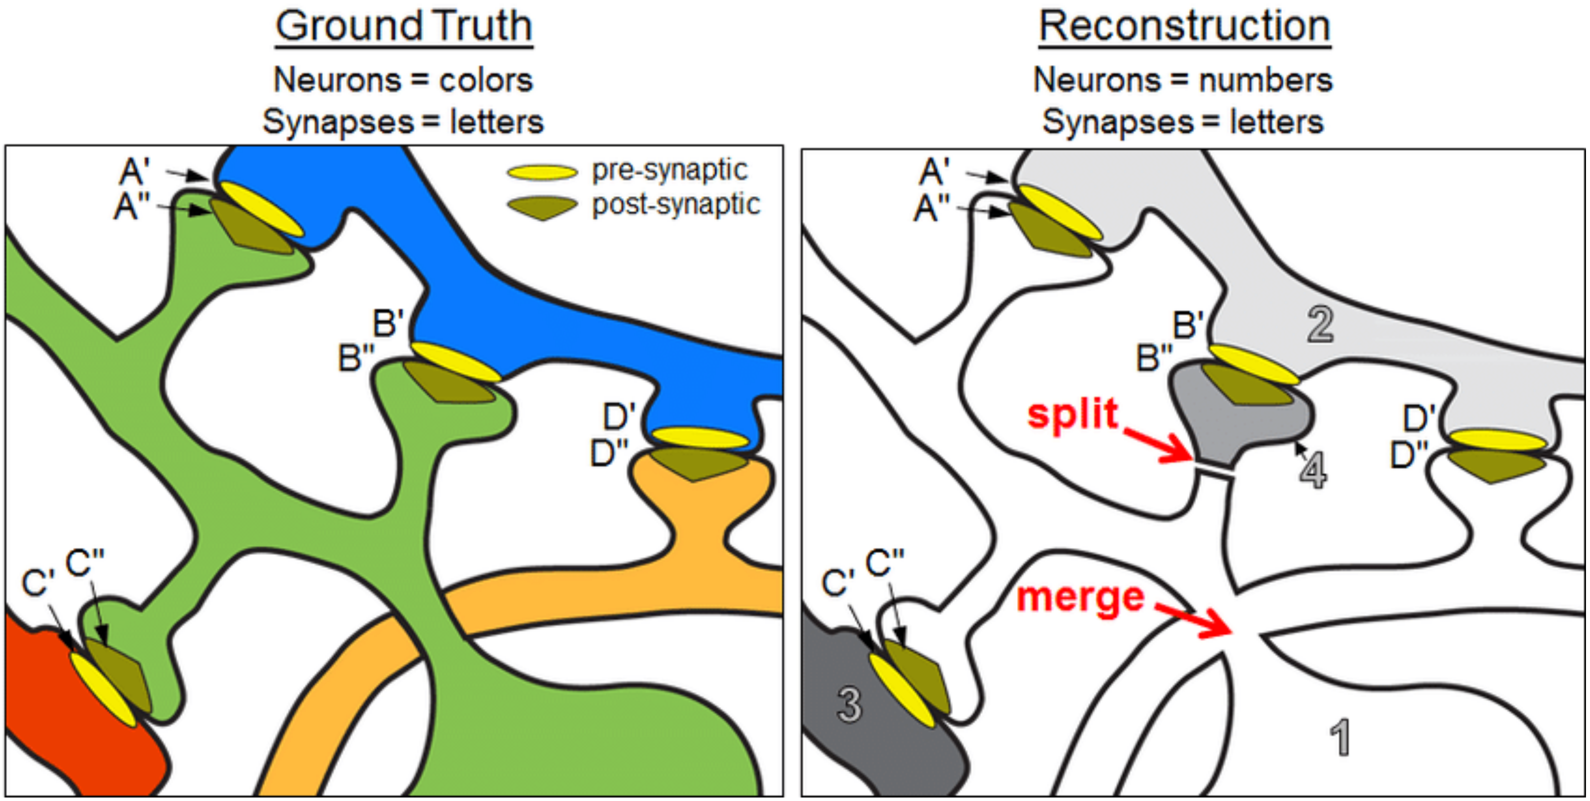

E. Reilly et al. Frontiers in Neuroinformatics 2017

Split and Merge Errors

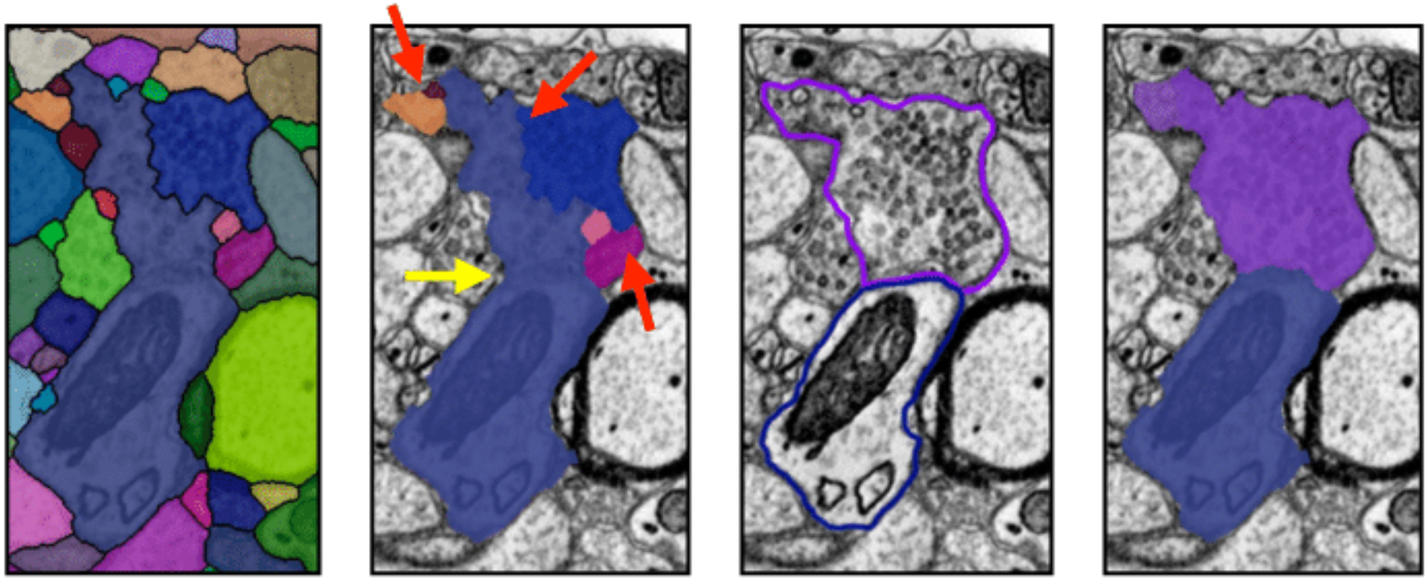

Initial Segmentation

Merge- and Split Errors

Correct Borders

Fixed Segmentation

D. Haehn et al. IEEE/CVF CVPR 2018

ISBI 2012 Evaluation Metrics

Foreground-restricted Rand Scoring:

prob. random pixel belongs to segment i in prediction and segment j in ground-truth

$$V_{\alpha}^{\mathrm{Rand}} = \frac{\sum_{ij} p_{ij}^2}{\alpha \sum_k s_k^2 + (1-\alpha) \sum_k t_k^2}$$

prob. random pixel belongs to segment k in prediction

prob. random pixel belongs to segment k in ground-truth

Rand F-score defined at $\alpha = 0.5$, which weights split and merge errors equally

ISBI 2012 Evaluation Metrics

Information Theoretic Scoring:

Mutual Information

$$V_{\alpha}^{\text{info}} = \frac{I(S;T)}{(1-\alpha)H(S)+\alpha H(T)}$$

Entropy

$\alpha = 0.5$ is the information theoretic F-score

ISBI 2012

From an algorithms perspective, why is this problem hard?

- Trouble exploiting context
- Lack of useful progress in recognition
- Cycles required to handle sufficiently large amounts of training data

Most basic approach: Thresholding (Otsu's Method)

1. Compute histogram and probabilities of each intensity level
2. Set up initial weights $\omega_i(0)$ and class means $\mu_i(0)$
3. Step through all possible thresholds $t = 1, \ldots \textit{maximum intensity}$
    1. Update ω_i and μ_i
    2. Compute variance $\sigma_b^2(t)$
4. Desired threshold corresponds to the maximum $\sigma_b^2(t)$

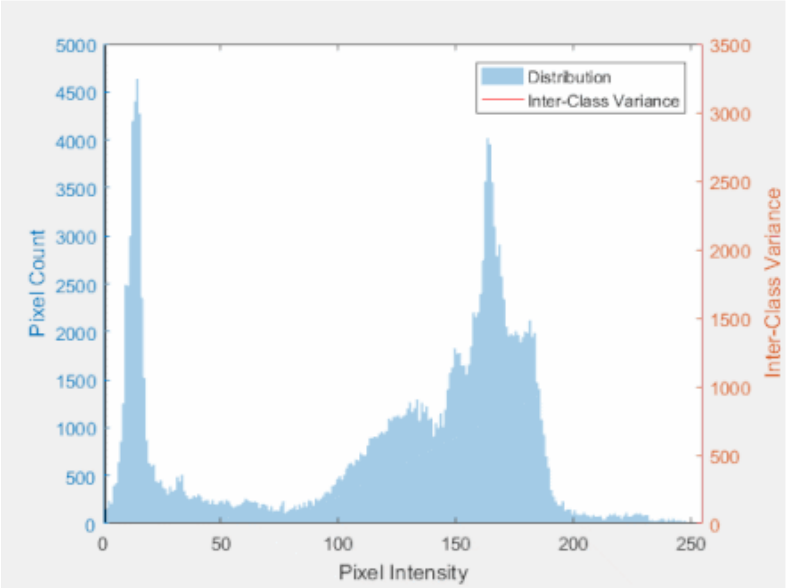

Otsu's Method Visualization 🅭 BY-SA 4.0 Lucas(CA)

Most basic approach: Thresholding (Otsu's Method)

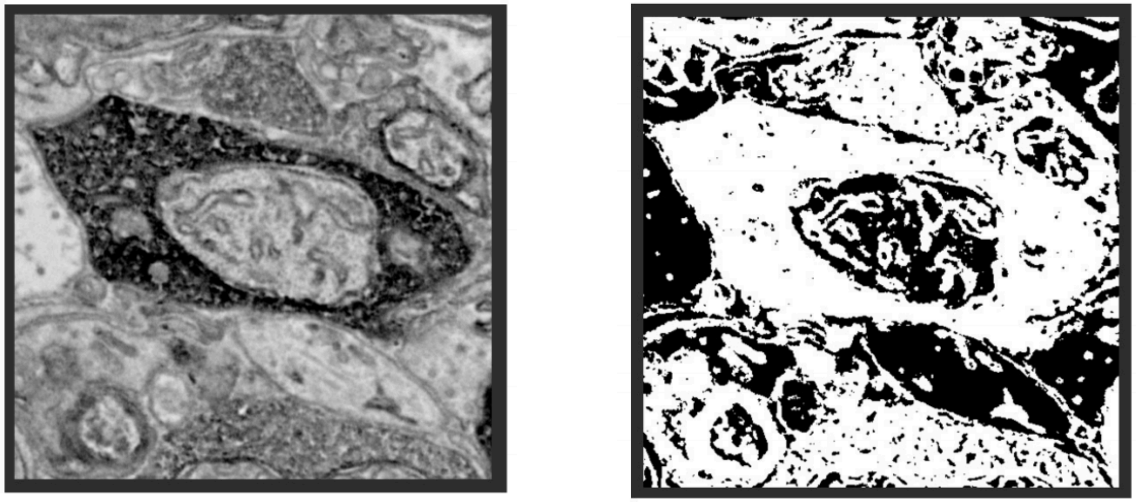

Original EM
(APEX Labeled Process)

Otsu's Method

A. Shahbazi et al. Scientific Reports 2018

ISBI 2012 Scores:

0.724521829 (Rand Score)

0.817598215 (Information Score)

Flexible Learning-Free Segmentation and Reconstruction (FLoRIN)

- Shahbazi et al. Scientific Reports 2018
- Code: https://github.com/jeffkinnison/florin
- Volumetric Learning-Free Segmentation
- Designed for sparse segmentation problems, but could be applied to the dense reconstruction problem

FLoRIN Steps

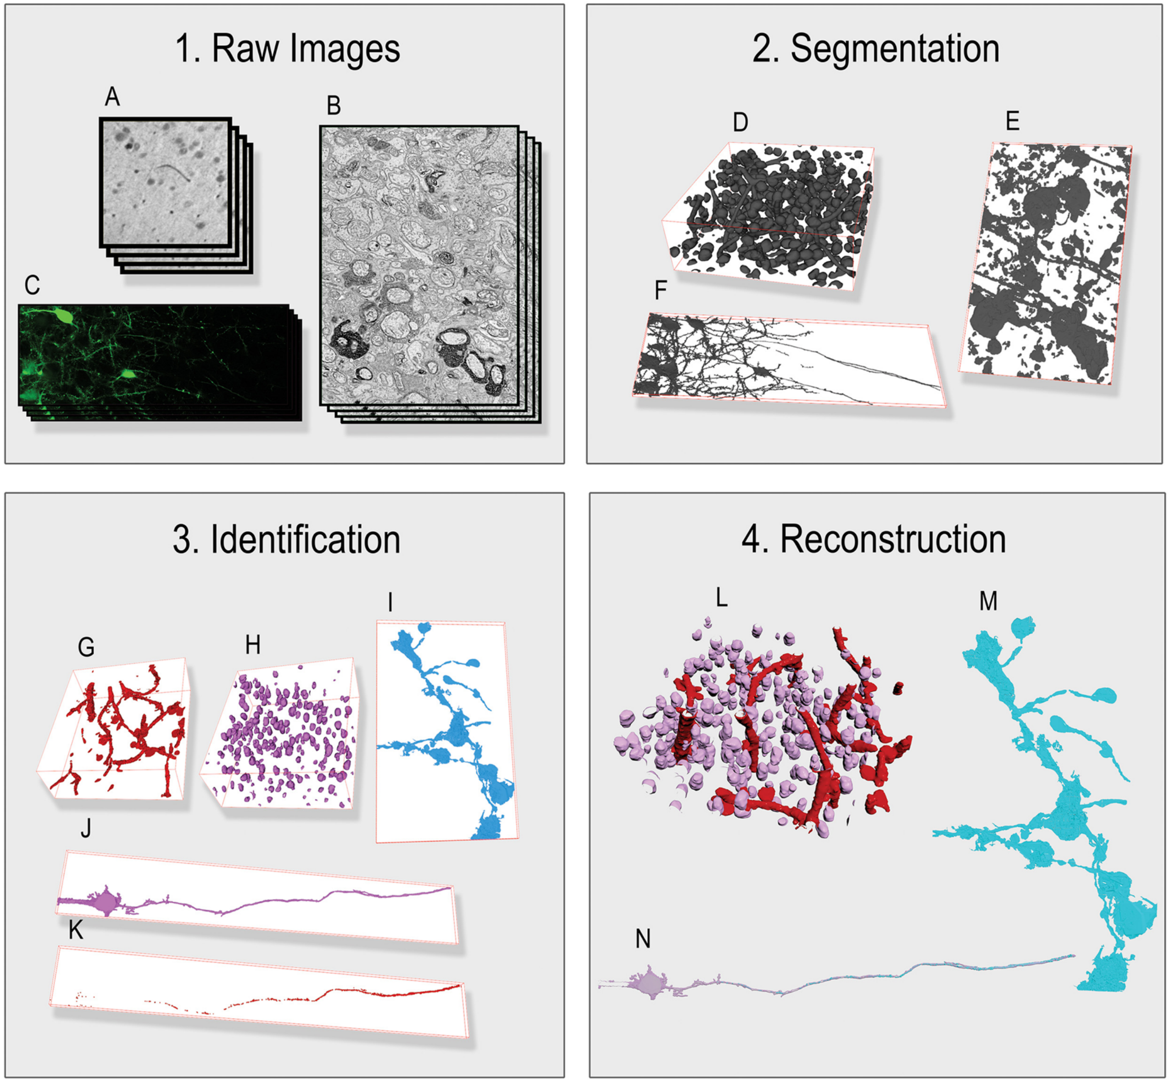

Segmentation

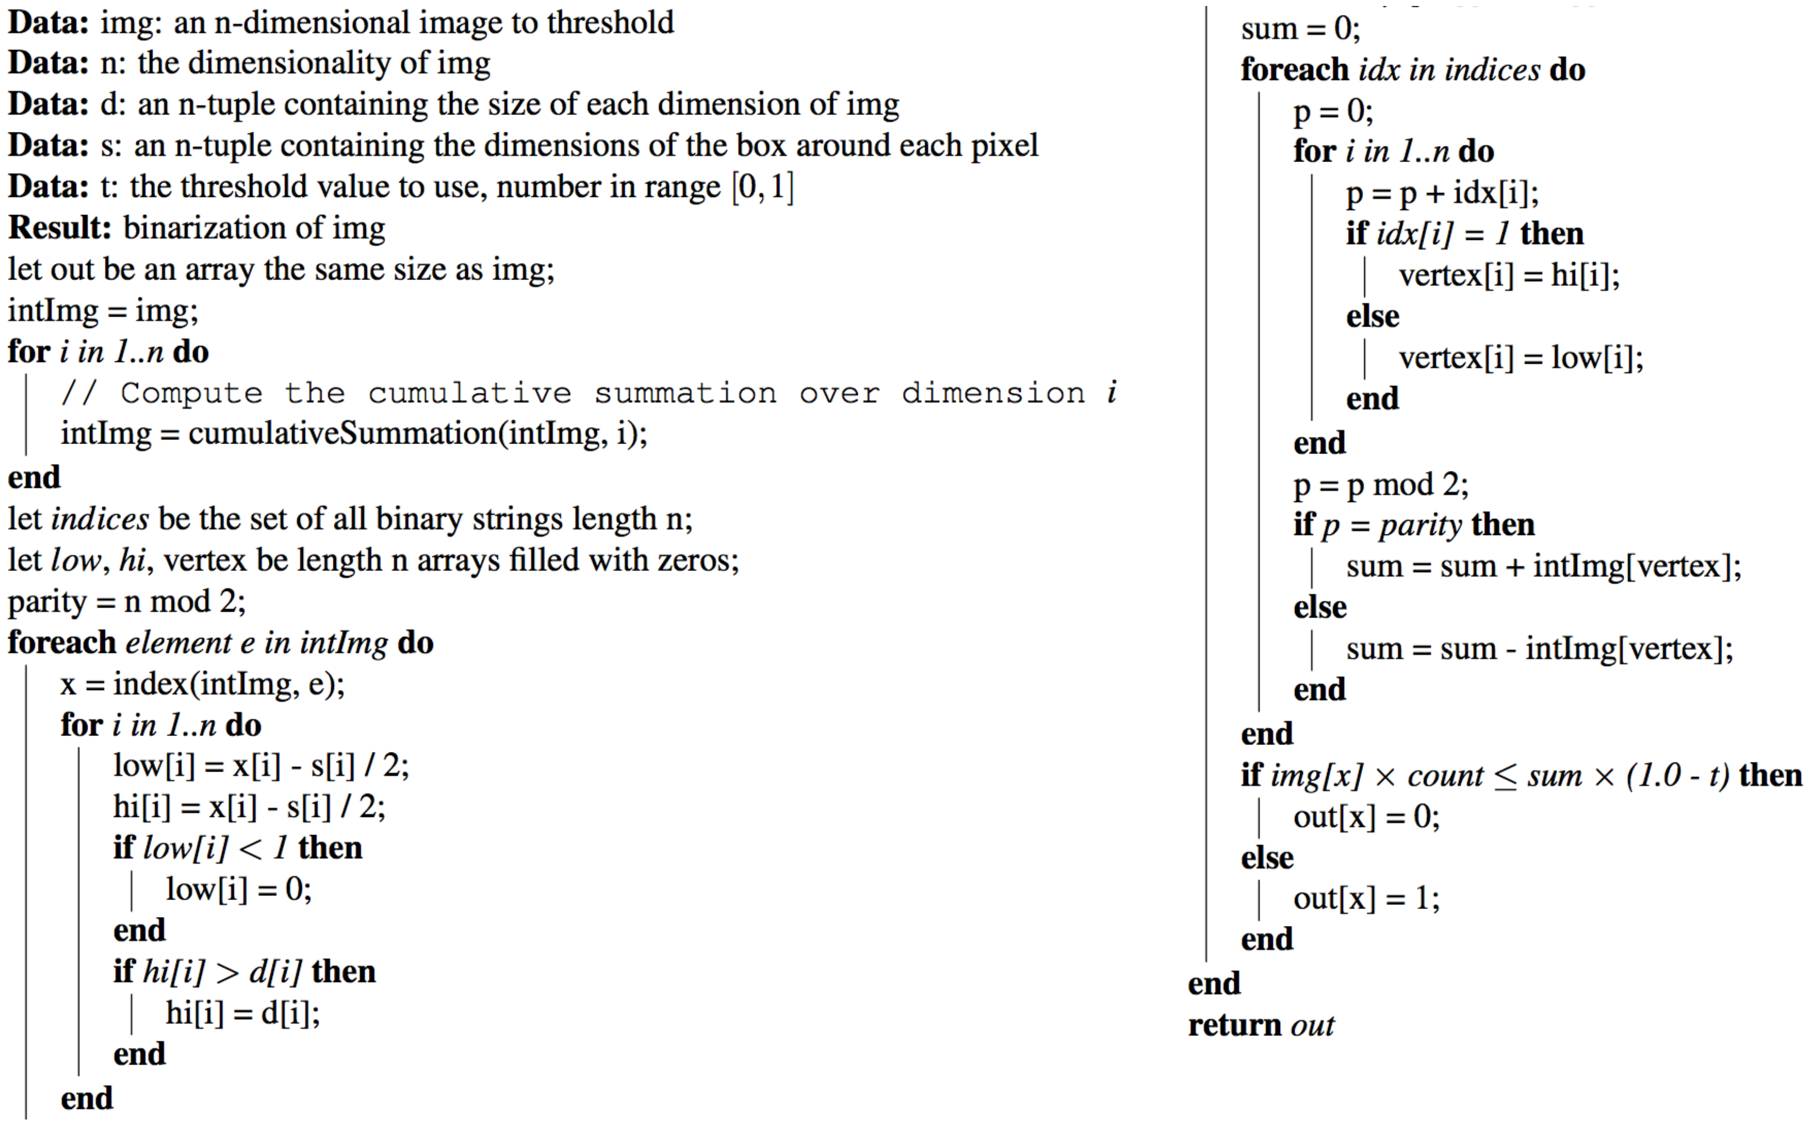

```
Data: img: an n-dimensional image to threshold
Data: n: the dimensionality of img
Data: d: an n-tuple containing the size of each dimension of img
Data: s: an n-tuple containing the dimensions of the box around each pixel
Data: t: the threshold value to use, number in range $[0, 1]$
Result: binarization of img
let out be an array the same size as img;
intImg = img;
for i in 1..n do
    // Compute the cumulative summation over dimension i
    intImg = cumulativeSummation(intImg, i);
end
let indices be the set of all binary strings length n;
let low, hi, vertex be length n arrays filled with zeros;
parity = n mod 2;
foreach element e in intImg do
    x = index(intImg, e);
    for i in 1..n do
        low[i] = x[i] - s[i] / 2;
        hi[i] = x[i] - s[i] / 2;
        if $low[i] < 1$ then
            low[i] = 0;
        end
        if $hi[i] > d[i]$ then
            hi[i] = d[i];
        end
    end
    sum = 0;
    foreach idx in indices do
        p = 0;
        for i in 1..n do
            p = p + idx[i];
            if $idx[i] = 1$ then
                vertex[i] = hi[i];
            else
                vertex[i] = low[i];
            end
        end
        p = p mod 2;
        if $p = parity$ then
            sum = sum + intImg[vertex];
        else
            sum = sum - intImg[vertex];
        end
    end
    if $img[x] \times count \leq sum \times (1.0 - t)$ then
        out[x] = 0;
    else
        out[x] = 1;
    end
end
return out
```

Identification

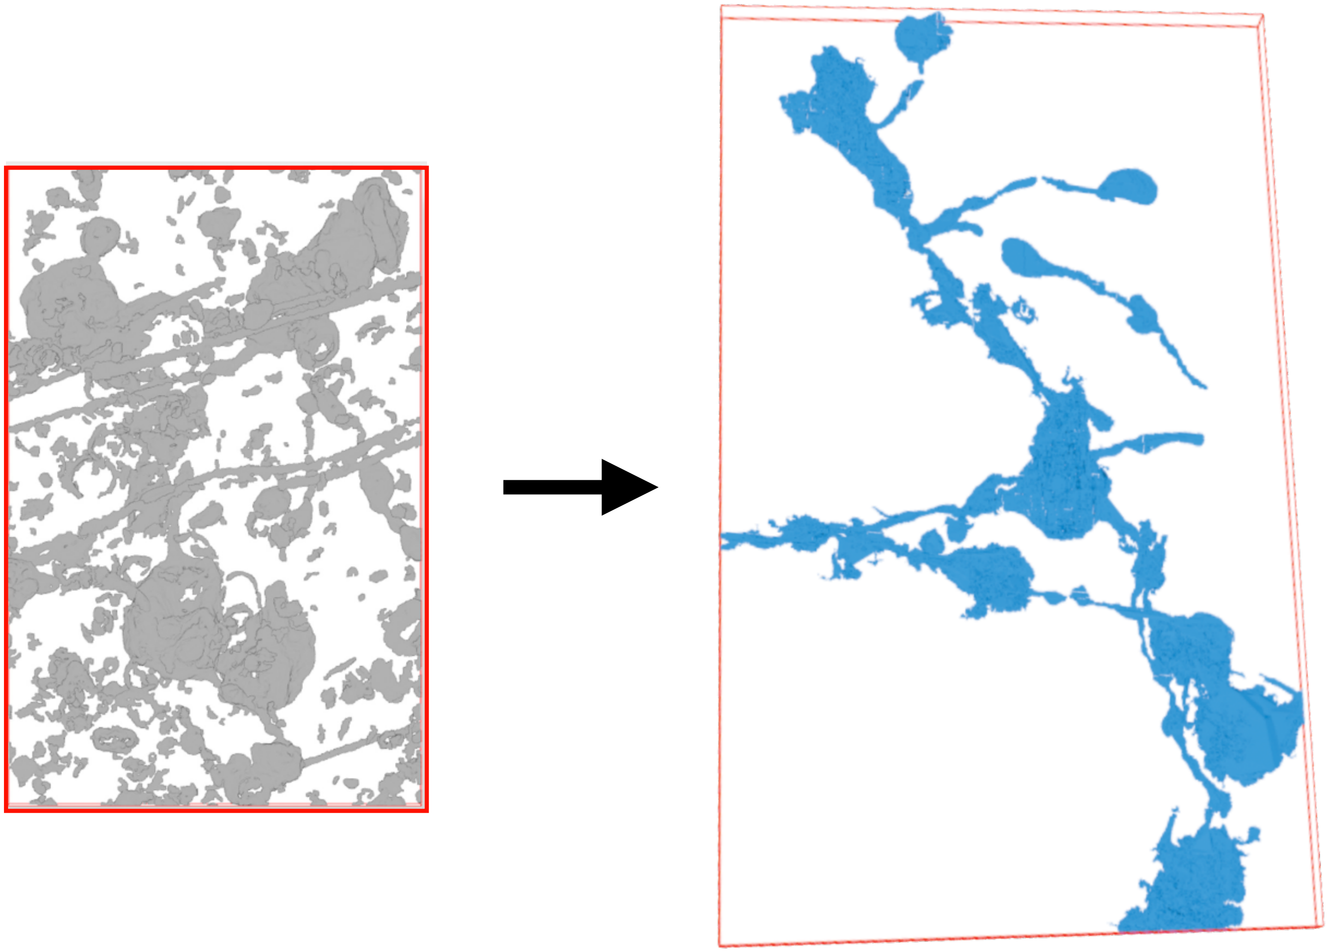

- Perform morphological operations
- Grouping connected components

Reconstruction

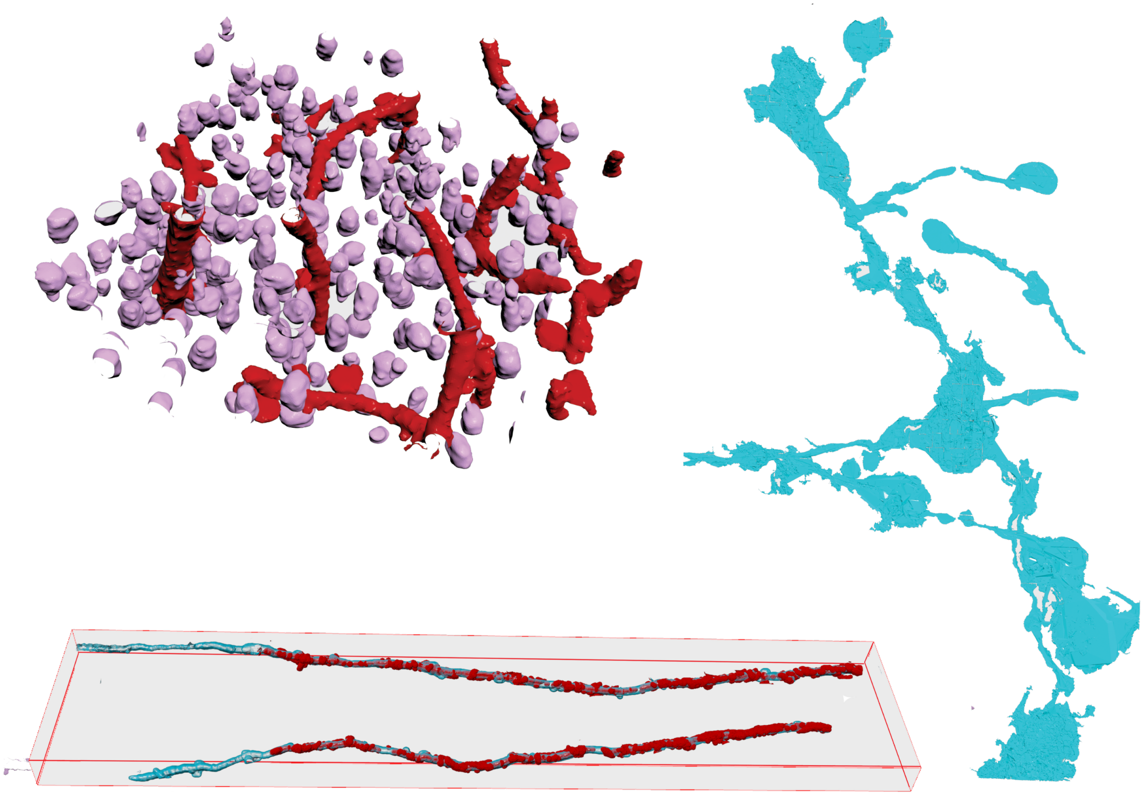

- Convert back to 3D if desired
- Compute and save statistics related to microstructures

FLoRIN Segmentation Performance

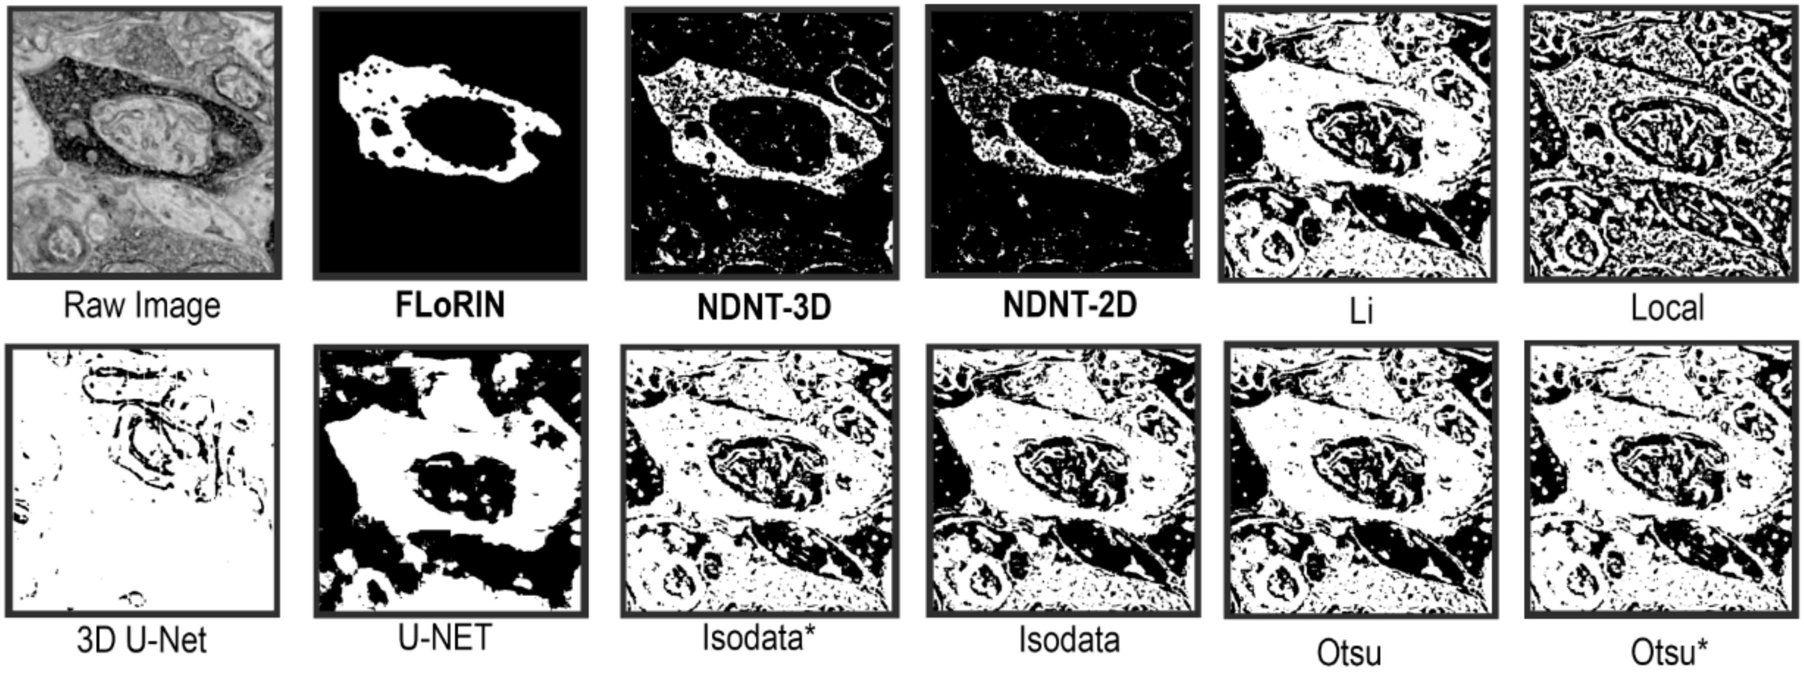

ISBI 2012 Scores:

0.896659612 (Rand Score)

0.952527835 (Information Score)

FLoRIN Segmentation Performance

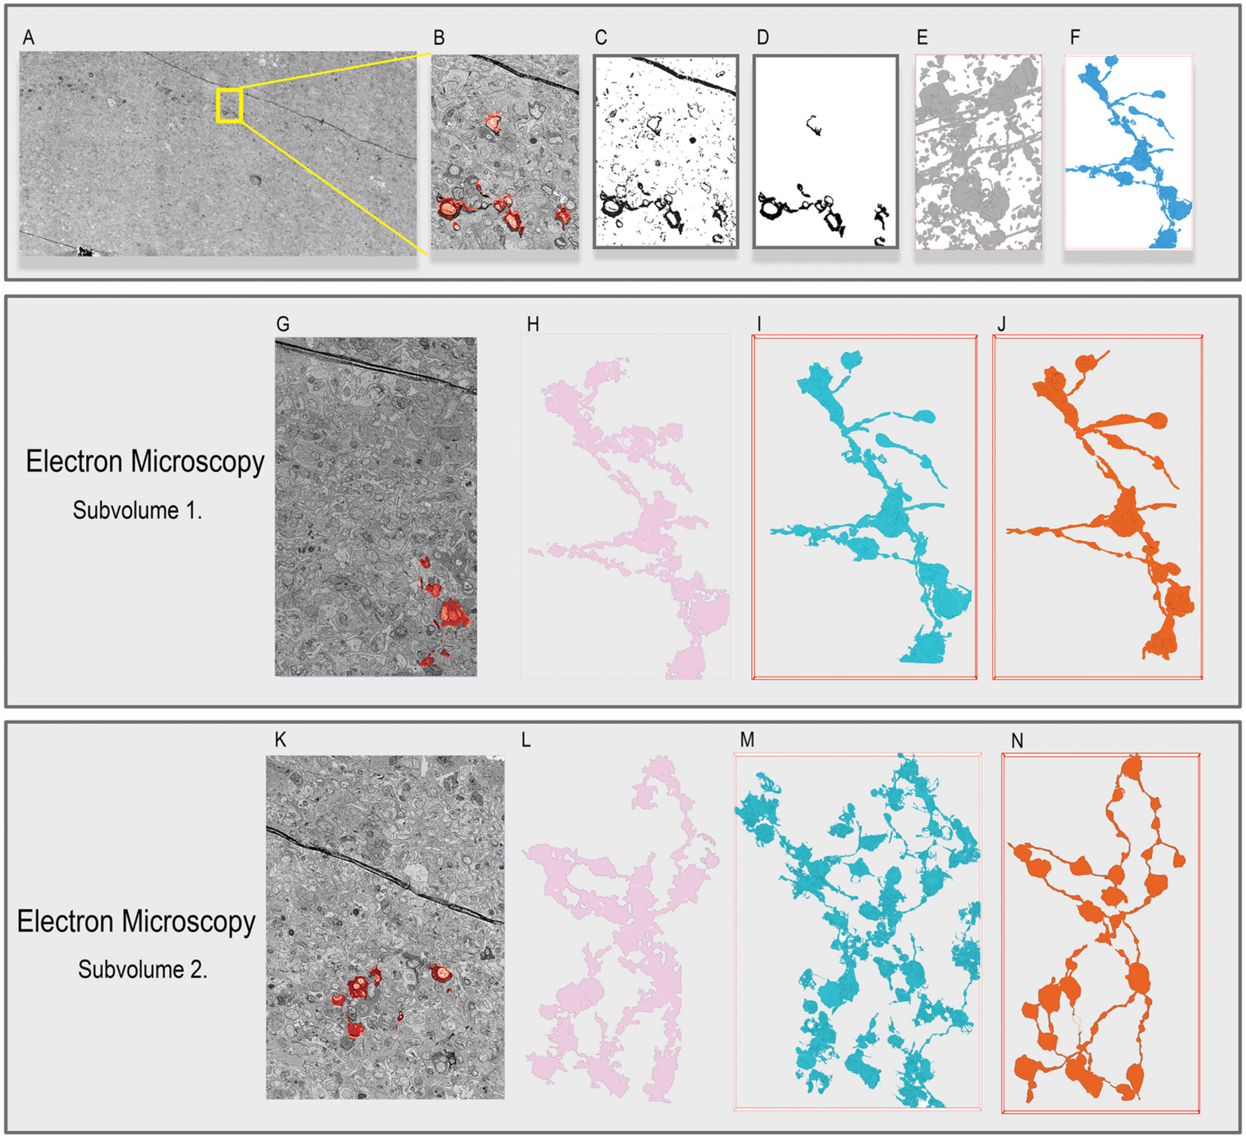

Machine Learning

- Given the variation observed in EM images, finding good universal thresholds is likely impossible
- Better strategy: train a classifier that can determine whether a candidate pixel is membrane or non-membrane
- What is biggest challenge for ML-based connectomics segmentation?

Strategy: Classifier + Context

- A. Vazquez-Reina et al. ICCV 2011
- Leverage information from **the entire volume** to obtain a globally optimal 3D segmentation
- Formulate segmentation as the solution to a fusion problem
  - Enumerate multiple possible 2D segmentations for each section in the stack
  - Enumerate a set of 3D links that may connect segments across consecutive sections
  - Identify the fusion of segments and links that provide the most globally consistent segmentation of the stack

Overview of segmentation fusion

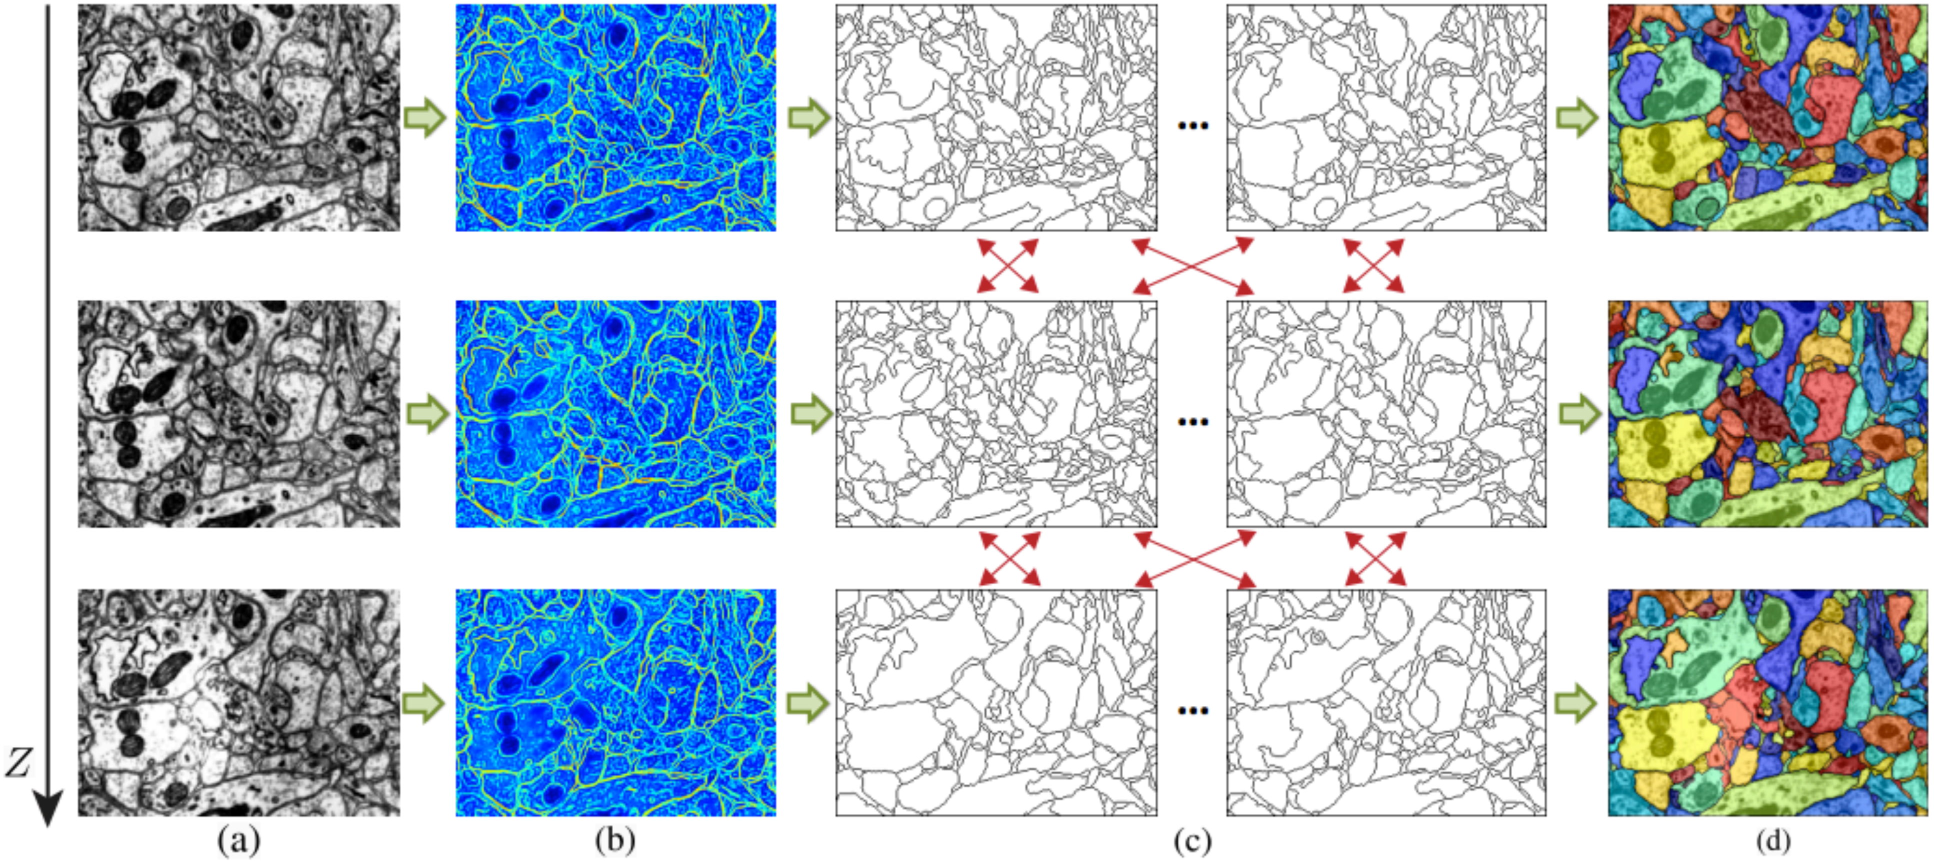

Compute probability of a pixel belonging to cell boundary **(random forest)**

Watershed transforms at different heights

Red arrows: possible 3D links

Segmentation Fusion

Image Credit: A. Vazquez-Reina et al. ICCV 2011

Average percentage of merge and split errors in four ssEM stacks

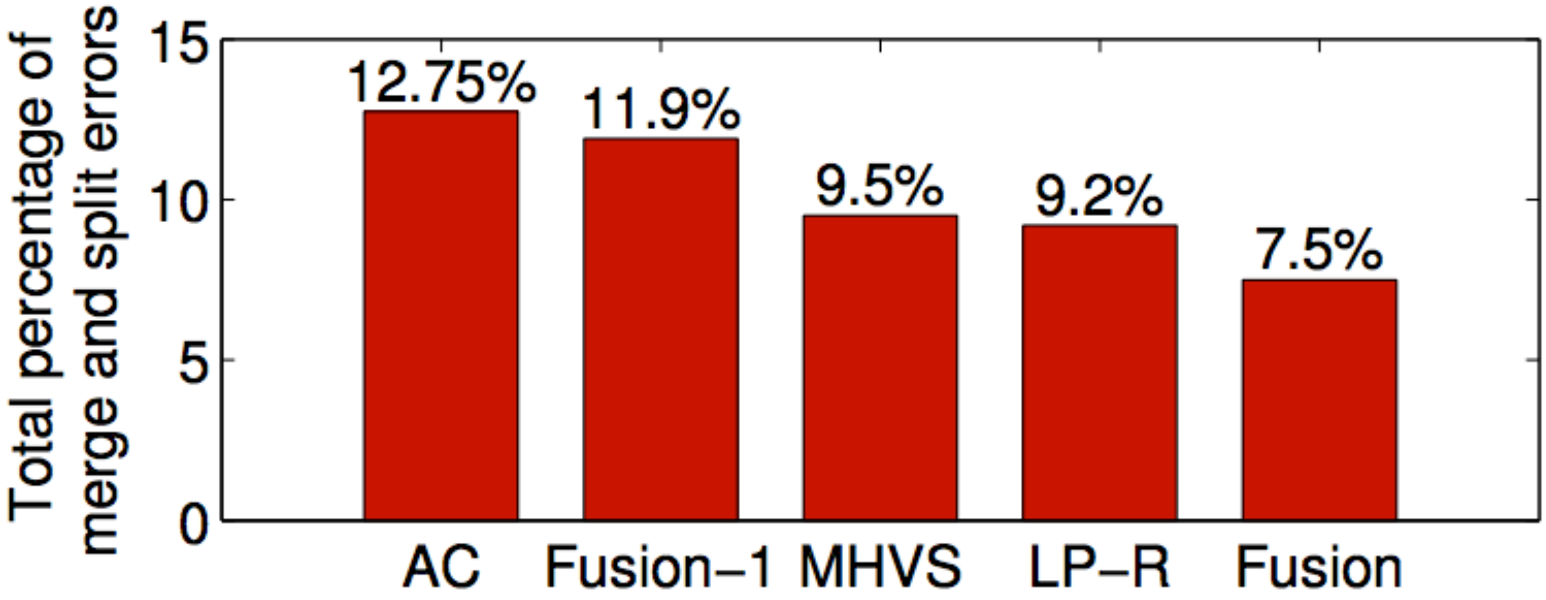

(1 2D seg.)

Most fusion errors come from:

- false splits along the z-axis
- creation of spurious small segments in areas between cells

Image Credit: A. Vazquez-Reina et al. ICCV 2011

ilastik

https://www.ilastik.org

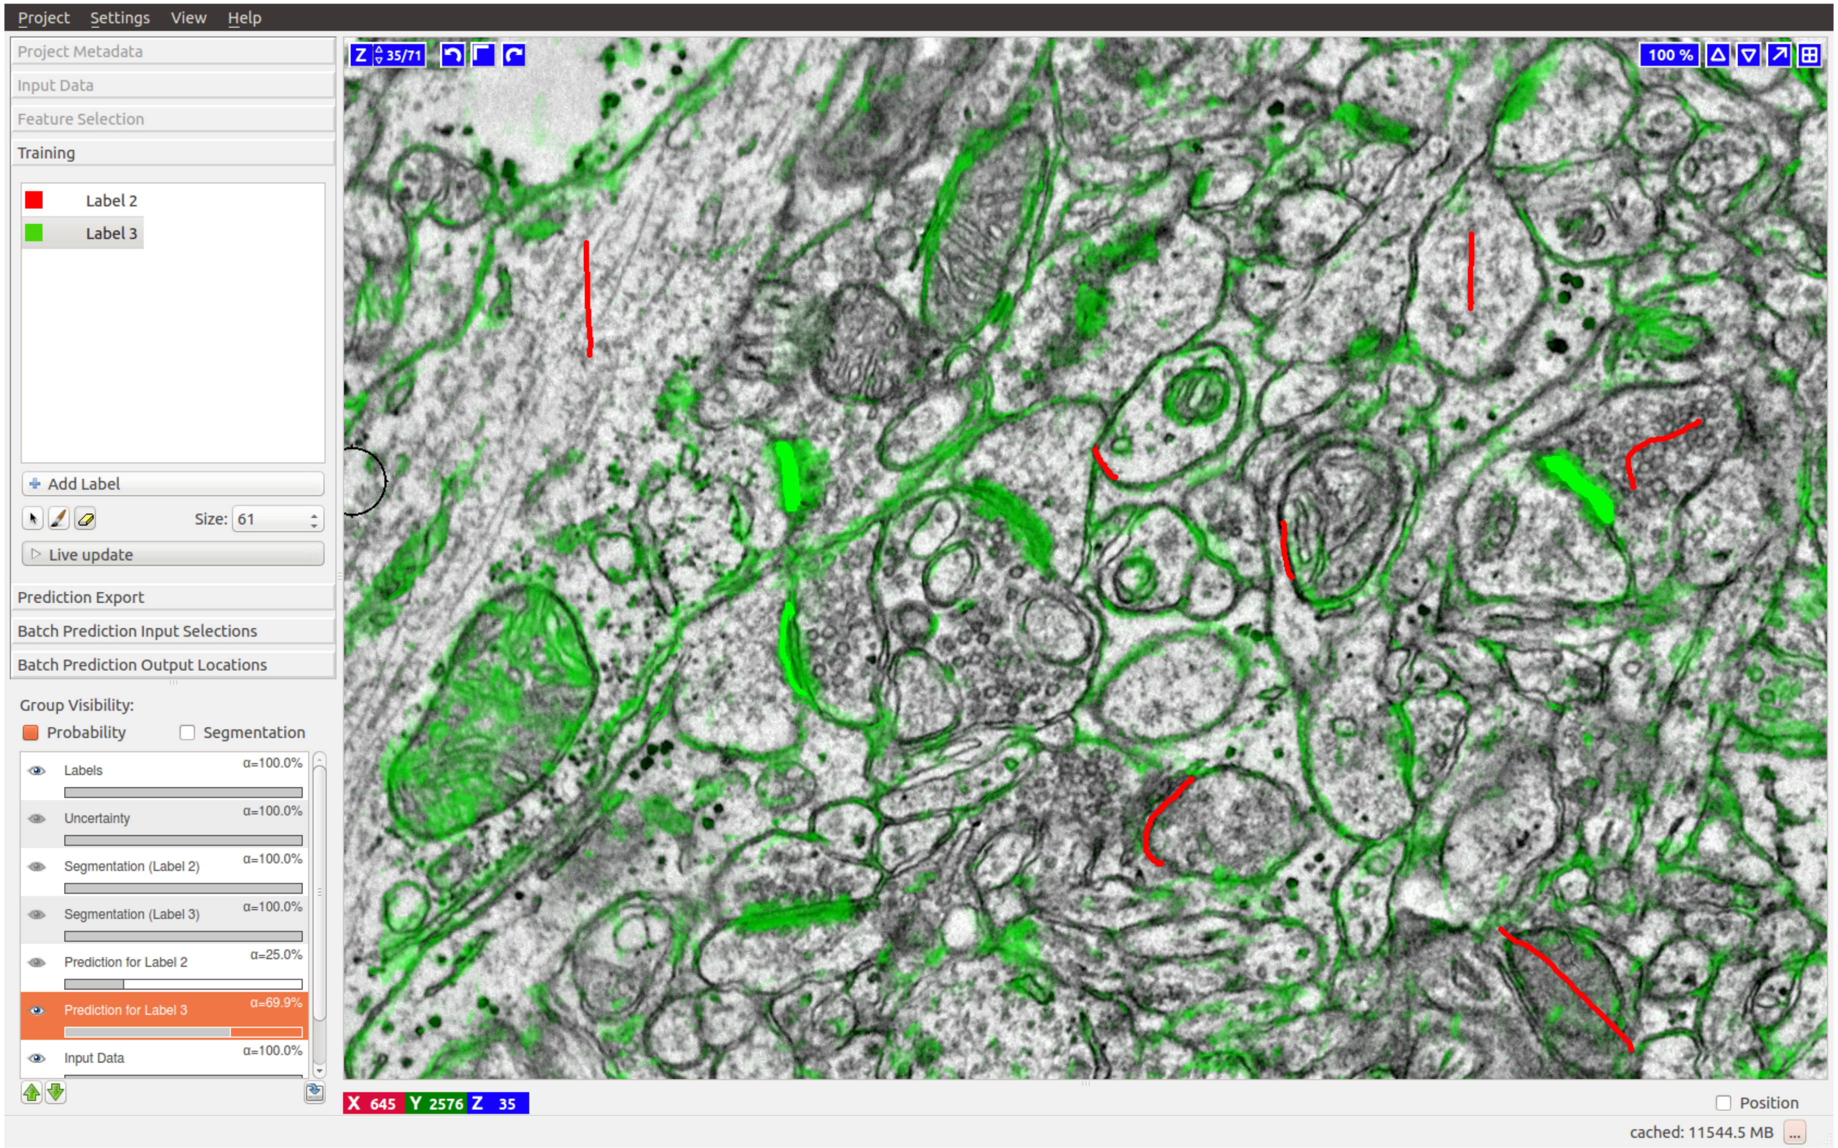

Size: 61

☑ Probability ☐ Segmentation

Labels α=100.0%

Uncertainty α=100.0%

Segmentation (Label 2) α=100.0%

Segmentation (Label 3) α=100.0%

Prediction for Label 2 α=25.0%

Prediction for Label 3 α=69.9%

Input Data α=100.0%

ilastik Pixel Classification Workflow

1. Select features: smoothed pixel intensity, edge filters or texture descriptors

2. Train random forest classifier

3. Perform semantic segmentation, then convert to objects:

    3a. Return probability map of each class

    3b. Perform thresholding and connected component analysis

Current ISBI Leaderboard

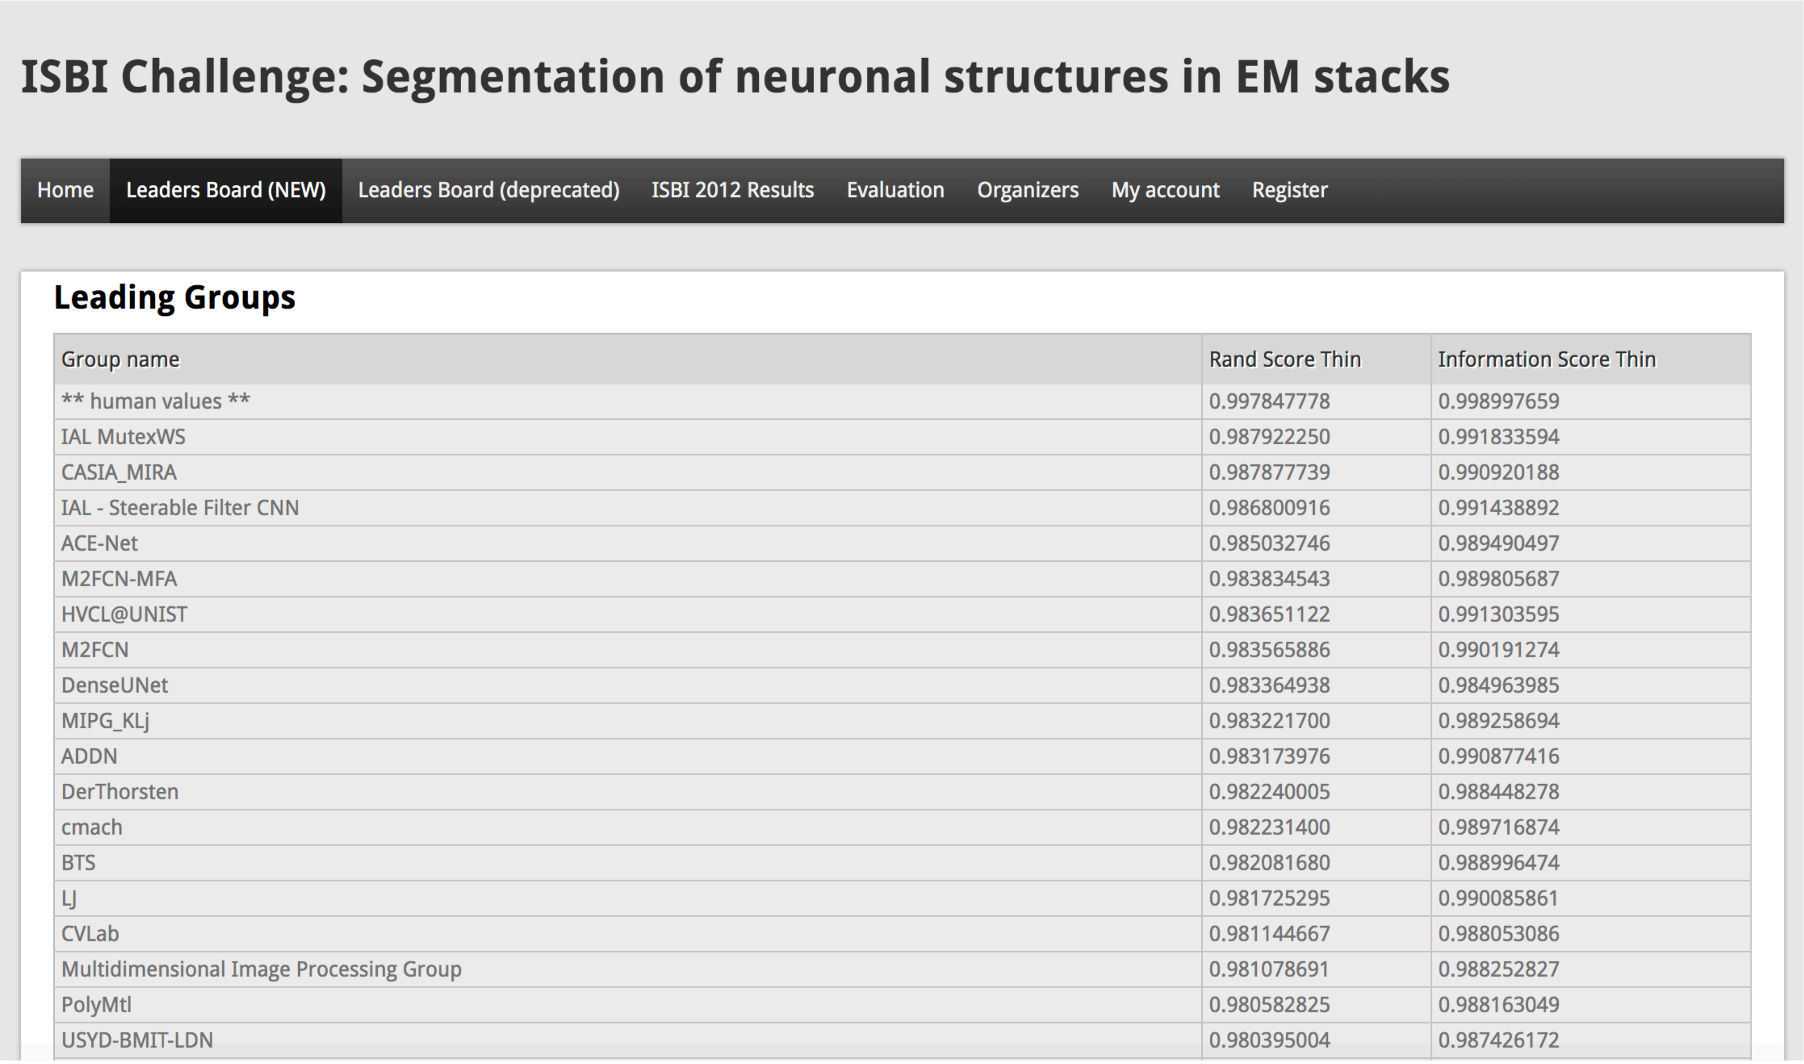

ISBI Challenge: Segmentation of neuronal structures in EM stacks

Leading Groups

Group name	Rand Score Thin	Information Score Thin
** human values **	0.997847778	0.998997659
IAL MutexWS	0.987922250	0.991833594
CASIA_MIRA	0.987877739	0.990920188
IAL - Steerable Filter CNN	0.986800916	0.991438892
ACE-Net	0.985032746	0.989490497
M2FCN-MFA	0.983834543	0.989805687
HVCL@UNIST	0.983651122	0.991303595
M2FCN	0.983565886	0.990191274
DenseUNet	0.983364938	0.984963985
MIPG_KLj	0.983221700	0.989258694
ADDN	0.983173976	0.990877416
DerThorsten	0.982240005	0.988448278
cmach	0.982231400	0.989716874
BTS	0.982081680	0.988996474
LJ	0.981725295	0.990085861
CVLab	0.981144667	0.988053086
Multidimensional Image Processing Group	0.981078691	0.988252827
PolyMtl	0.980582825	0.988163049
USYD-BMIT-LDN	0.980395004	0.987426172

http://brainiac2.mit.edu/isbi_challenge/leaders-board-new