Learning Global Features for Neuron Reconstruction in EM Images

Brian Matejek

May 4, 2016

Adviser: Tom Funkhouser

- There has been a substantial amount of research on predicting affinities between neighboring voxels.
- Most neuron reconstruction efforts focus on hierarchical agglomeration based on local features.
- Some algorithms focus solely on segmenting individual images in an image stack and producing geometrically consistent results afterwards.

SNEMI3D Challenge Dataset

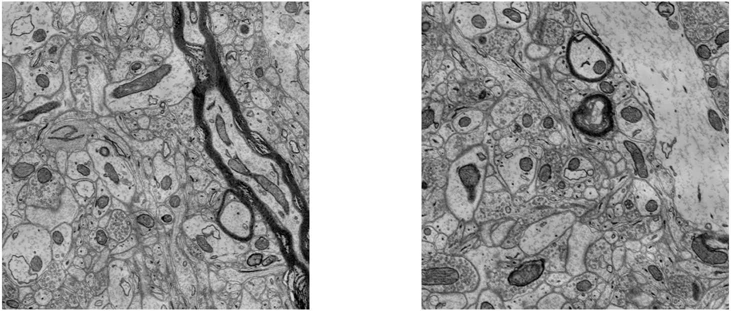

- Contains both training and testing datasets.
- Both datasets have a resolution of 6nm x 6nm x 30nm with 1024 x 1024 x 100 voxels.
- Each volume further divided into four components of size 512 x 512 x 100 voxels to create training, validation, and testing sets.

VD2D3D Convolutional Neural Network

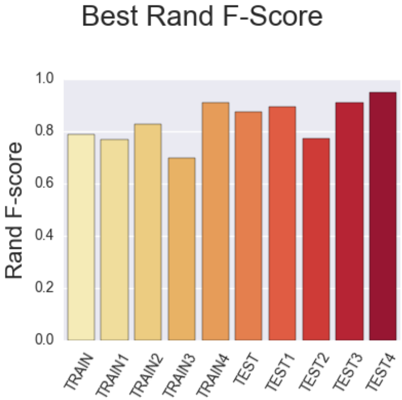

Dataset	Rand F-Score
Training	0.7882
Training 1	0.7691
Training 2	0.8265
Training 3	0.6986
Training 4	0.9111
Testing	0.8761
Testing 1	0.8926
Testing 2	0.7742
Testing 3	0.9091
Testing 4	0.9484

Entropy and Information Gain

- Convenient metric for determining how useful a feature is in a decision tree or random forest.
- Measures how well a given attribute "splits" the data into correctly labeled subsets.
- Entropy S of a dataset with labels $l \in \mathcal{L}$:

$$S = \sum_{l} -p_l \log p_l$$

- Information gain:

$$G = S(parent) - \mathrm{avg}_w[S(children)]$$

Boundary Features

- Watershed algorithm run on the affinity graph produces oversegmented supervoxels.
- Most current algorithms focus on features on boundaries between neighboring supervoxels.

Feature	Information Gain	Proportion Correct	Split Value
Boundary Mean	0.480152	0.902664	0.465558
Boundary Min	0.225903	0.684337	0.484120
Boundary Max	0.485295	0.905715	0.651979
Boundary Median	0.476195	0.897424	0.496180
Boundary Standard Deviation	0.059196	0.549403	0.007983
Boundary Ranking	0.110195	0.695079	15.000000
Boundary Scaled Ranking	0.087253	0.636174	0.598639
Boundary Voxels	0.085481	0.701021	6877.000000
Boundary Voxels Proportion	0.047943	0.659362	0.147625

Shortest Path Attributes

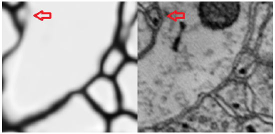

Metric	Oversegmented	Boundary
Boundary Max	0.807364	0.807397
Boundary Mean	0.771332	0.755981
Distance	1.170613	**1.596400**
Path Minimum	0.804500	**0.765497**
First Moment	0.822043	0.819276
Second Moment	0.006919	0.007619
Third Moment	0.400652	**-2.353312**

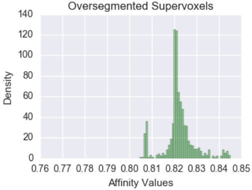
Oversegmented Supervoxels

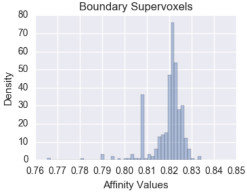
Boundary Supervoxels

Feature	Information Gain	Proportion Correct	Split Value
Minimum Affinity on Path	0.448539	0.865927	0.784019
Path Skew	0.111357	0.633956	-1.073322
Standard Deviation	0.390160	0.865749	0.011536
Distance	0.516898	0.919068	1.638740

Random Walk Features

- The shortest path is susceptible to holes in the membrane predictions.
- Generate randomized walks between supervoxels and determine the lowest affinity along the walk.
- Randomly select a voxel to start a random walk and count the number of times a given boundary is crossed. Repeat several hundred million times so each voxel starts on average of 5 times.

Feature	Information Gain	Proportion Correct	Split Value
Weighted Random Walk	0.503235	0.904358	0.432407
Unweighted Random Walk	0.501438	0.903375	0.432442
Monte Carlo Ratio	0.474363	0.898594	4550.148926

Merge Predictions

- Random forest trained on training sets 3 and 4 and validated on training sets 1 and 2.
- The random forest has 1000 trees, a max tree depth of 7, and $\sqrt{n}$ variables per split, where n is the number of features.

Datset	True Positives	False Positives	False Negatives	True Negatives	Proportion
Training 1	26573	954	566	16120	96.562
Training 2	21959	1176	662	18099	95.613
Training 3	31015	1473	1334	16320	94.402
Training 4	38843	1277	1028	15577	95.937
Testing 1	42920	967	1541	15649	95.894
Testing 2	41478	1343	1472	20229	95.637
Testing 3	40167	882	875	15907	96.962
Testing 4	45857	1019	1756	13088	95.504

Table: Results for the smoothing term random forest.

Merge Predictions

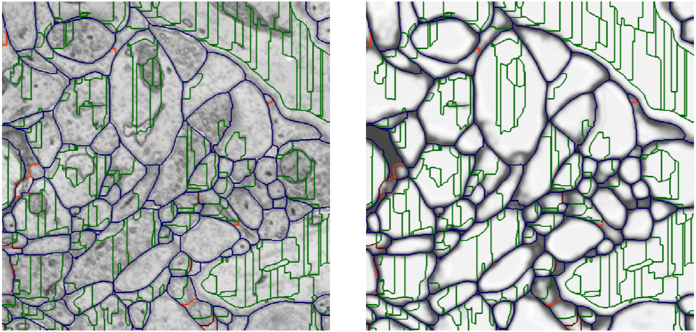

Figure: Green boundaries represent true positive merges, blue boundaries true negative rejections, red boundaries false positive merges, and orange boundaries false negative rejections. Image overlay (left), affinity overlay (right).

Data Predictions

- Random forest trained on training sets 3 and 4 and validated on training sets 1 and 2.
- Predicts if an interior supervoxel belongs to the same neuron as a volume boundary supervoxel.
- Same parameters as the merge random forest except there are 100 trees.

Datset	True Positives	False Positives	False Negatives	True Negatives	Proportion
Training 1	615,872	493,618	96,657	23,688,528	97.629
Training 2	591,114	428,579	120,914	24,473,833	97.855
Training 3	1,093,351	805,934	242,909	24,486,088	96.061
Training 4	4,579,180	1,639,462	310,407	29,193,114	94.542
Testing 1	11,755,055	7,594,526	406,510	42,256,582	87.098
Testing 2	6,151,521	6,880,161	447,788	51,177,505	88.666
Testing 3	9,432,835	3,843,551	400,256	33,610,444	91.025
Testing 4	9,794,379	21,794,848	875,410	22,065,211	58.426

Table: Results for the data term random forest.

K-Median Clustering

- Choose k centers out of all supervoxels to minimize the sum of distances from all other supervoxels to the nearest center.
- NP-Complete problem so approximate the optimal solution using an expectation maximization algorithm.
- How to choose a suitable value for k...

Varying Parameters

- Run k-medians algorithm for $2 \leq k \leq 1000$.

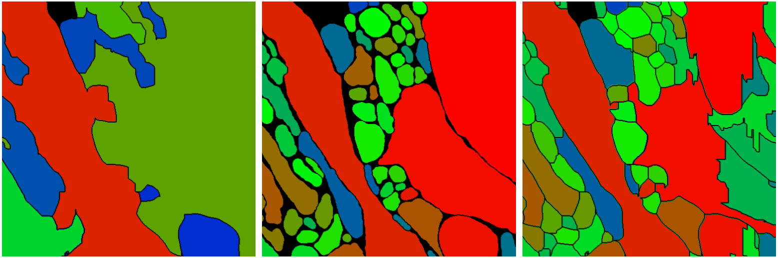

Figure: (left) $K = 25$, (center) Ground Truth, (right) $K = 600$

Facility Location Assignment

- Neurons are very long and the reconstructed volume is very small in diameter.
- Every supervoxel must belong to a neuron which exits the volume.
- Restrict k cluster medians to supervoxels on the volume boundary.

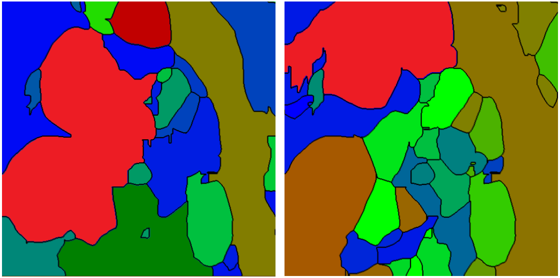

Figure: Differences for $k = 25$ between facility location (left) and k-median (right).

Creating Features from Global Algorithms

- Vary the global algorithm parameters so early iterations favor undersegmentation and the later iterattions tend to oversegment.
- Record the first parameter at which two supervoxels are predicted to belong to different neurons.
- Record the total proportion of iterations where the neighboring supervoxels were in different neurons.

Feature	Information Gain	Proportion Correct	Split Value
Facility Location First Split	0.272788	0.829087	34.000000
Facility Location Proportion	0.457039	0.905312	0.691995
K-Median First Split	0.161626	0.784442	11.000000
K-Median Proportion	0.216262	0.812225	0.876992

Graph Cut Approximation

- Minimize the following energy function where $\{\mathcal{S}\}$ is the set of supervoxels and $\{\mathcal{N}\}$ is the set of neighboring pairs of supervoxels:

$$Energy = \delta * \sum_{s \in \mathcal{S}} d_s(a) + \beta \sum_{s,t \in \mathcal{N}} v_{s,t}(a, b)$$

- $d_s(a)$ is the cost of assigning supervoxel s label a and $v_{s,t}(a, b)$ is the cost of assigning neighboring supervoxels s and t different labels a and b.
- Use the random forests described previously to generate the weights $d_s(a)$ and $v_{s,t}(a, b)$.

Graphcut Results

Feature	Information Gain	Proportion Correct	Split Value
Graph Cut First Split	0.608345	0.946476	1800.000000
Graph Cut Proportion	0.600311	0.944623	0.404255
Graph Cut $\delta = 360, \beta = 1080$	0.589250	0.941797	0.000000

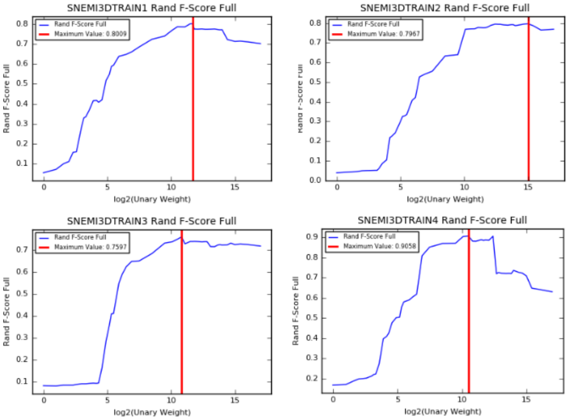

Global Merging

- Hierarchical merging is not robust to holes in the predicted membrane walls.
- When considering a merge, look at boundaries which will collapse. Only merge if a majority of the boundaries are predicted as oversegmented.

Figure: A membrane hole that results in an improper merge using traditional hierarchical clustering algorithms.

Global Merging (Results)

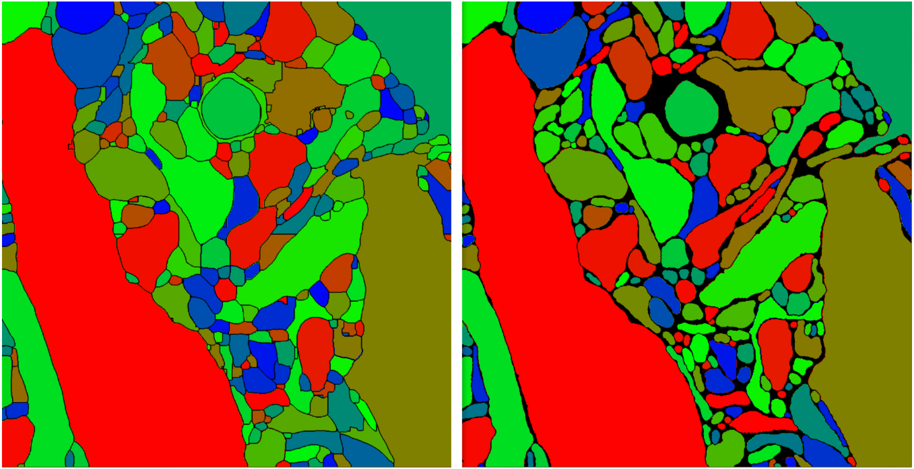

Figure: The predictions on the entire testing dataset using the global merging algorithm (left) next to the ground truth (right).

Global Merging (Results)

Dataset	Boundary Mean	Random Forest	±%
Training	0.7882	0.8729	+8.47%
Training 1	0.7691	0.8414	+7.23%
Training 2	0.8265	0.8455	+1.90%
Training 3	0.6986	0.8466	+14.80%
Training 4	0.9111	0.9513	+4.02%
Testing	0.8761	0.8841	+0.80%
Testing 1	0.8926	0.8946	+0.20%
Testing 2	0.7742	0.7794	+0.53%
Testing 3	0.9091	0.9235	+1.44%
Testing 4	0.9484	0.9512	+0.28%

Boundary Merging

- Similar to the global merging from above except merges start at the boundaries of the volume.
- Boundary merge candidates must belong to a supervoxel already connected to the boundary.
- Enforces the prior that all neurons must exit the volume.

Feature	Information Gain	Proportion Correct	Split Value
Global Merge First Split	0.631059	0.948675	59.000000
Global Merge Proportion	0.649119	0.954448	0.430000
Global Merge Threshold 0.80	0.658245	0.953728	0.000000
Boundary Merge First Split	0.513559	0.917720	60.000000
Boundary Merge Proportion	0.536999	0.922932	0.280000
Boundary Merge Threshold 0.80	0.547767	0.924298	0.000000

Global Approximation Results

Random Forest Improvements

- Random forest trained on training sets 3 and 4 and validated on training sets 1 and 2.
- The random forest has 1000 trees, a max tree depth of 7, and $\sqrt{n}$ variables per split, where n is the number of features.

Datset	Boundary	Shape (±%)	Combined (±%)	Global (±%)	Combined (±%)
Training 1	95.431	96.657 (+1.226)	96.562 (−0.095)	97.220 (+0.658)	97.306 (+1.875)
Training 2	94.081	95.422 (+1.341)	95.613 (+0.191)	96.097 (+0.484)	96.253 (+2.172)
Training 3	91.999	94.242 (+2.243)	94.402 (+0.160)	96.097 (+1.695)	96.071 (+4.072)
Training 4	94.160	95.372 (+1.212)	95.937 (+0.565)	97.174 (+1.237)	97.248 (+3.088)
Testing 1	94.996	95.627 (+0.631)	95.894 (+0.522)	96.372 (+0.478)	96.351 (+1.355)
Testing 2	94.114	95.318 (+1.204)	95.637 (+0.319)	96.238 (+0.601)	96.229 (+2.115)
Testing 3	96.038	96.659 (+0.621)	96.962 (+0.303)	97.531 (+0.569)	97.569 (+1.531)
Testing 4	93.897	95.403 (+1.506)	95.504 (+0.101)	96.234 (+0.730)	96.290 (+2.393)

Table: Results for the smoothing term random forest.

Random Forest Improvements

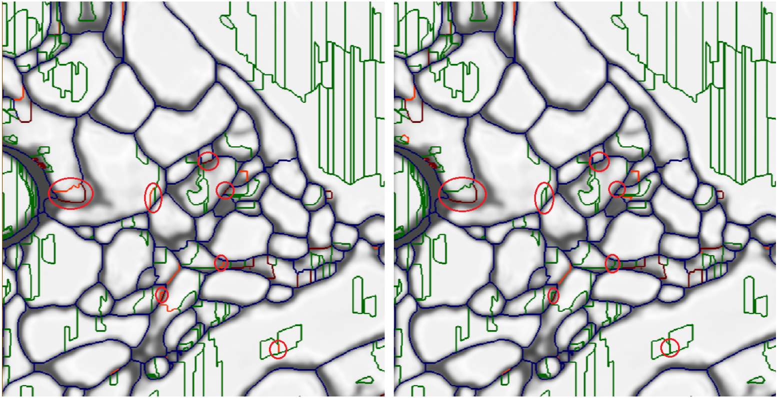

Figure: Improvements from the original random forest (left) when adding global features (right) are circled in red. Green boundaries represent true positive merges, blue boundaries true negative rejections, red boundaries false positive merges, and orange boundaries false negative rejections.

Hierarchical Agglomeration with Global Features

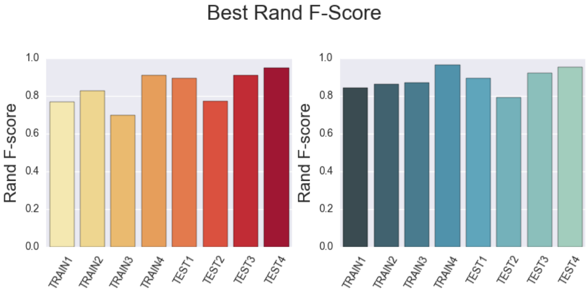

Dataset	Boundary Mean	Random Forest	±%
Training 1	0.7691	0.8431	+7.40%
Training 2	0.8265	0.8616	+3.51%
Training 3	0.6986	0.8711	+17.25%
Training 4	0.9111	0.9653	+5.42%
Testing 1	0.8926	0.8932	+0.06%
Testing 2	0.7742	0.7919	+1.77%
Testing 3	0.9091	0.9202	+1.11%
Testing 4	0.9484	0.9534	+0.50%

Conclusion and Future Work

- Global features and other shape descriptors improve merge prediction results by offering robust solutions to errors in the affinity graphs.
- In the future, I want to explore additional shape features that are robust to different values in the affinities themselves.
- Augment other existing agglomerations algorithms with these features to improve on performance.