

RESEARCH ARTICLE

A Context-Aware Delayed Agglomeration Framework for Electron Microscopy Segmentation

Toufiq Parag^{1*}, Anirban Chakraborty², Stephen Plaza¹, Louis Scheffer¹

1 Janelia Research Campus, HHMI, Ashburn, VA, USA, **2** Department of Diagnostic Radiology, National University of Singapore, Singapore

* paragt@janelia.hhmi.org



Abstract

Electron Microscopy (EM) image (or volume) segmentation has become significantly important in recent years as an instrument for connectomics. This paper proposes a novel agglomerative framework for EM segmentation. In particular, given an over-segmented image or volume, we propose a novel framework for accurately clustering regions of the same neuron. Unlike existing agglomerative methods, the proposed context-aware algorithm divides superpixels (over-segmented regions) of different biological entities into different subsets and agglomerates them separately. In addition, this paper describes a “delayed” scheme for agglomerative clustering that postpones some of the merge decisions, pertaining to newly formed bodies, in order to generate a more confident boundary prediction. We report significant improvements attained by the proposed approach in segmentation accuracy over existing standard methods on 2D and 3D datasets.

OPEN ACCESS

Citation: Parag T, Chakraborty A, Plaza S, Scheffer L (2015) A Context-Aware Delayed Agglomeration Framework for Electron Microscopy Segmentation. PLoS ONE 10(5): e0125825. doi:10.1371/journal.pone.0125825

Academic Editor: Stefan Strack, University of Iowa, UNITED STATES

Received: September 18, 2014

Accepted: March 26, 2015

Published: May 27, 2015

Copyright: © 2015 Parag et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data is in the paper and available from GitHub (https://github.com/janelia-flyem/neuroproof_examples).

Funding: The authors have no support or funding to report.

Competing Interests: The authors have declared that no competing interests exist.

1 Introduction

Extracting the network structure among neurons in animal brain has gained prominence lately in the field of neuroscience. Rapid advances in imaging technology, in particular Electron Microscopy (EM) techniques, have enabled us to trace neural bodies in unprecedented level of details. However, recording in such high resolution (at nanometer scale) generates massive amount of data that is too large to annotate manually. Automated region labeling or segmentation is considered to be the most viable strategy for generating a dense reconstruction of neural anatomy. Some recent efforts of such reconstruction yielded impressive results utilizing machine learning/computer vision tools such as image segmentation, and offered valuable biological insights to the neuroscience community [1][2].

Image segmentation for natural scenes has a long history in computer vision literature [3][4][5]. In recent years, there have also been many fruitful attempts to automatically identify meaningful regions in EM images using segmentation techniques [6][7][8][9][10][11]. Most of these studies initially apply a pixelwise (we denote locations on both 2D and 3D EM data as pixels in this paper) classifier to determine whether or not any particular pixel belongs to the

cell boundary. The quantified confidence values of the pixelwise classifier are utilized to produce an initial (over-) segmentation through methods such as Watershed [12].

Different approaches resort to different methods to generate the final segmentation by merging or clustering the over-segmented bodies to corresponding neuronal cells. Andres et al. [7] addresses this problem by searching the optimal subset of superpixel borders that form closed surfaces. Several studies work with the watershed merge tree in order to identify the regions to be combined for the final segmentation [13][14]. Some approaches applied agglomerative or hierarchical clustering [8][10][15] for this purpose. For anisotropic datasets, where the depth resolution (z -dimension) is coarser than the planar resolution (x, y dimensions), segmented bodies on one section overlap with multiple regions in the adjacent sections. Therefore, a complete 3D reconstruction needs to establish the correct correspondence, through an alignment or co-segmentation technique [9][11], among segmented regions across multiple planes. In this study, we restrict ourselves only to segmentation on images for anisotropic data—the subsequent alignment is out of the scope of this paper. For both isotropic and anisotropic reconstruction, the outputs of segmentation algorithms need to be corrected afterwards, either manually [2] or by combining with a manually traced skeletonized representation [1].

Biologically, the interior of a neuron cell comprises several distinct sub-structures (or sub-categories) such as cytoplasm, mitochondria, vesicles etc. An ideal binary pixel classifier—which assigns a pixel to one of the two categories: cell boundary and cell interior—should label all locations within these sub-categories to cell-interior class. Several past studies [16][17][18][19] recommend increasingly complex pixelwise detector models to attain a binary prediction output. In contrast, some recent works [15][11] represent the sub-categories (cytoplasm, mitochondria etc.) of cell body by multiple classes and apply relatively simpler classifiers (in terms of model size, learning time and convenience) for this multiclass classification problem. The results of [15][11] suggest that using prior domain knowledge to divide a problem into multiple components can achieve high segmentation quality with simpler classifier models requiring less computation.

However, we believe the methods of [15][11] do not exploit the full benefit of multiclass predictions on EM data. Regardless of the quality of multiclass pixel classifier output, the algorithms in [15][11] do not distinguish between regions of one sub-structure (e.g., cytoplasm) from those of another (e.g., mitochondria) at the superpixel level. That is, the classification is divided into multiple sub-classes in pixel-level, but the subsequent fusion or superpixel clustering step does not utilize this additional information to compute the final segmentation. This often leads to sub-optimal performances by these methods. For example, Fig 1 shows an EM image (plane in a 3D volume) and the corresponding pixel predictions for the mitochondria sub-class. By not using this sub-class prediction explicitly in the clustering step, the final output of [15] failed to merge many regions into the correct cell (marked by ‘S’) and connected some of them to wrong cells (marked by ‘M’).

This paper introduces a context-aware scheme for combining over-segmented regions by utilizing the prior knowledge of sub-classes. We adopt an agglomerative or hierarchical clustering framework [8][10] due to its advantages such as low space, time complexity and flexibility to tune for over/under segmentation. We develop a two-pass agglomeration policy where the (estimated) cytoplasm regions are grouped together in the first phase and then the remaining mitochondria bodies are absorbed into the cell cytoplasm. In these two stages, the superpixels are agglomerated based on different merge criteria that are defined by different contexts, which is why we call it context-aware agglomeration. Our proposed context aware approach significantly reduces the false split and merge errors (example shown in Fig 1) provided fairly accurate sub-structure detection. In addition, this strategy substantially reduces the training data requirement, as well as the predictor model complexity, which in turn offers significant

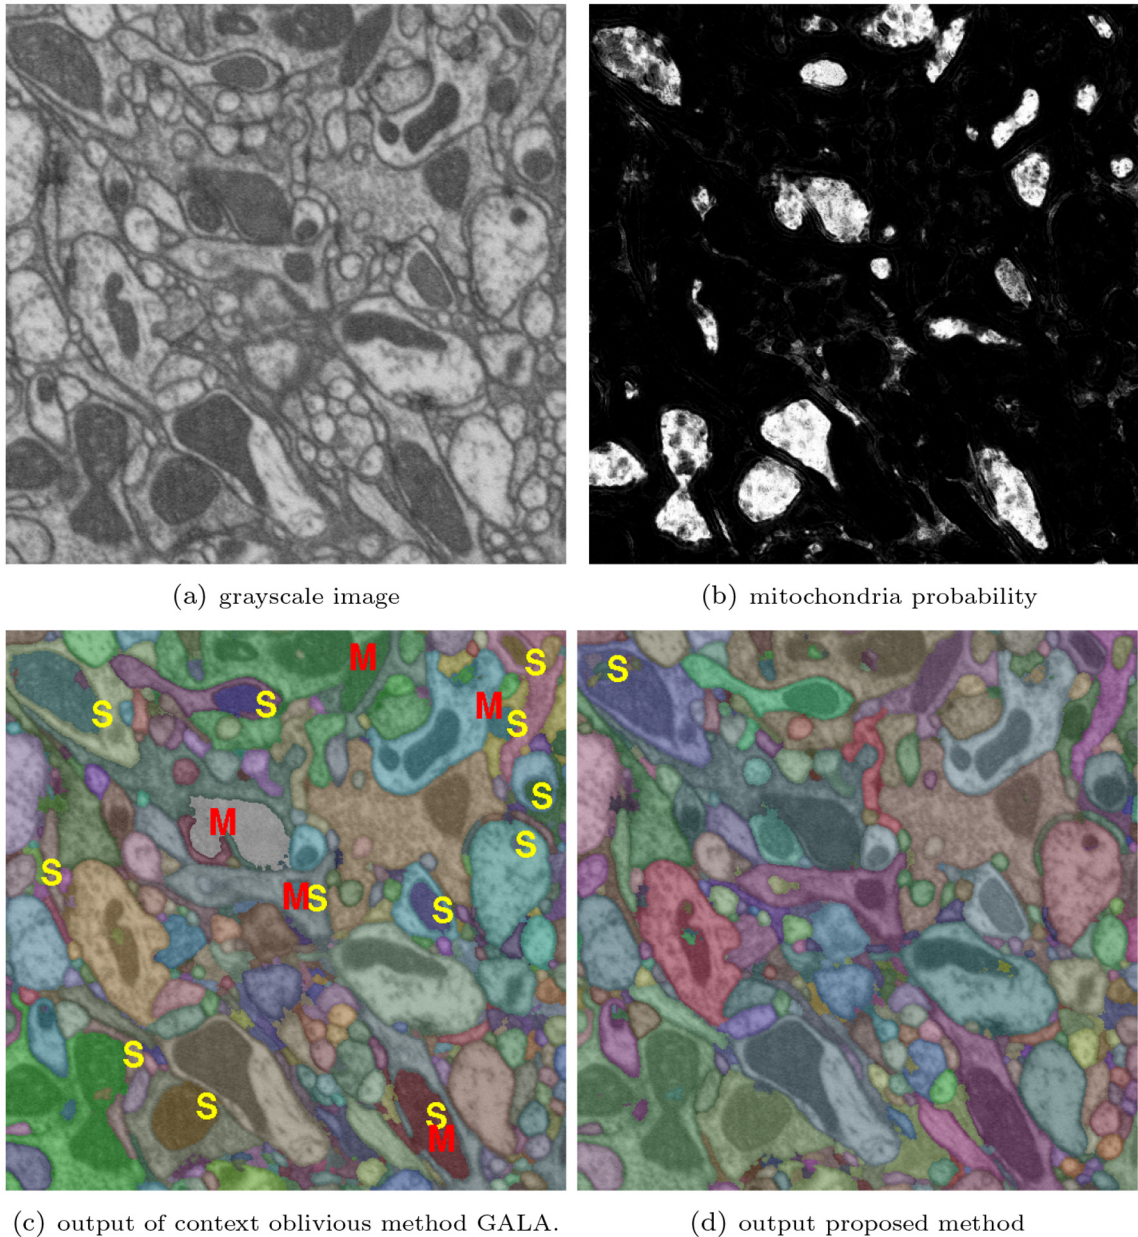


Fig 1. Improved segmentation accuracy by context-aware approach on FIBSEM data. (a) one plane of input volume, (b) mitochondria detection on that plane, (c) the output of GALA [15] (context oblivious), and (d) the output of proposed context aware method. The segmented region labels are overlaid on the image using random artificial colors. S and M on images indicate locations of false split and merge respectively

doi:10.1371/journal.pone.0125825.g001

increase in learning speed. The findings of this study further inspired us to design an interactive training algorithm [20] for region boundary predictor that does not require exhaustively labeled groundtruth. Generating such an exhaustive annotation is considered to be a bottleneck for neural reconstruction [21].

We also propose a modified version of the hierarchical clustering algorithm to cluster the superpixels in both phases of the context-aware framework. The proposed clustering method emphasizes on minimizing under-segmentation errors since these errors are conventionally

costlier to correct than the over-segmentation errors [8]. In order to minimize the number of false merges, we ‘delay’ the merge decisions on a certain type of boundaries to be resolved at a later time. Compared to the traditional agglomerative scheme of [10], the proposed modification reduces the number of false merges significantly. We also attempt to analyze why our agglomeration approach performs better than the Global multicut scheme [7] on the dataset used for our experiments.

The paper is organized as follows. We define the problem in Section 2 and briefly describe the existing clustering segmentation algorithms in Section 2.1. Then we explain the proposed delayed agglomeration scheme in Section 2.2. This delayed strategy is employed in both the stages of our context-aware algorithm discussed in Section 2.3. Section 3 reports our experimental setup, both quantitative and qualitative results and their analyses. We conclude and discuss our findings further in Section 4.

2 Methods

A formal definition of the problem we are addressing assumes an initial over-segmentation, comprising N superpixels $\{S_1, S_2, \dots, S_N\} \subseteq \mathcal{S}$, of an EM image or volume with M neurites (neuronal regions) where $N \gg M$. Let $L(S)$ be the neurite region that S actually belongs to. Our goal is to correctly assign these N superpixels such that each S_i , $i = 1, 2, \dots, N$ is assigned to its corresponding $L(S_i)$.

We denote a boundary between two superpixels (i.e., oversegmented regions) by a pair of regions $e \triangleq \{S_i, S_j\}$ and the set of all such boundaries by E . In a graph representation, each of the regions S_i is considered to be a node and the boundary or face between two regions is regarded as an edge—a notation we will be using throughout the paper. Also, let the boundary label map $B: \mathcal{S} \times \mathcal{S} \rightarrow \{0, 1\}$ assign a 1 to a boundary that actually separates one neurite region from another and a 0 to the boundary incorrectly generated due to over-segmentation. The problem of correctly merging S_i to its corresponding $L(S_i)$ is similar to a clustering problem where the number of clusters cannot be computed a priori. Following [8][10][15], we adopt an agglomerative approach for superpixel clustering.

In our context-aware scheme, the set of superpixels is divided into two subsets: 1) the \mathcal{S}_c of potential cytoplasm superpixels, and 2) the set \mathcal{S}_m of potential mitochondria superpixels. The set of cytoplasm superpixels is clustered first with the proposed delayed agglomeration algorithm. Agglomeration of the mitochondria superpixels is also performed by the proposed delayed method, but with a different merge criterion. In order to assist the reader to comprehend the novelty of the proposed approach, we introduce the prior studies on agglomerative clustering for EM segmentation [8][10][15] in Section 2.1. Afterwards, Sections 2.2 and 2.3 discuss the delayed agglomeration and the context-aware framework respectively.

2.1 Prior Works on Agglomerative Clustering for EM Segmentation

Several existing EM segmentation approaches [8][10][15] tackled the problem of superpixel clustering by agglomerative hierarchical clustering, as described in Table 1 Algorithm 1. These methods assume a superpixel boundary estimator $h: \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ that assigns real valued confidences to all edges in E . This boundary estimator may represent the real-valued prediction of a classifier distinguishing true boundaries from the false ones [10], or compute the mean value of boundary pixel probabilities [8] or return the overlap percentage between borders of two adjacent superpixels. The value of $h(\{S_i, S_j\}) \in [0, 1]$ indicates how confident the estimator is about the existence of a true boundary between S_i and S_j ; a large $h(\{S_i, S_j\})$ implies the estimator is very confident that the boundary $\{S_i, S_j\}$ is correct while a small value implies the boundary was probably generated as an artifact of over-segmentation and therefore is false. Given such a

Table 1. Algorithm 1: Existing Agglomerative Segmentation.

Input: S_1, S_2, \dots, S_N and confidence function h .

Output: R_1, R_2, \dots, R_N

1 for all the i do $R_i = S_i$;

2 Repeat

3 $e^* \equiv \{R_i, R_j\} = \min_{e \in E} h(e)$;

4 Merge R_j to R_i and update E ;

5 for all the $R_b \in Nbr(R_i)$ do

6 Recompute $h(\{R_i, R_b\})$;

7 until $h(e) \leq \delta$;

doi:10.1371/journal.pone.0125825.t001

function h , the hierarchical clustering algorithm iteratively merges cell boundaries in the increasing order of confidence values $h(e)$ (Line 1 in Table 1 Algorithm 1) until a stopping criterion is satisfied, e.g., $h(e) > \delta$ where δ is a pre-defined threshold. After each merge, it updates the neighborhood structure of the merged superpixel, i.e., the neighbors of the absorbed region become the neighbors of the (newly) merged cell.

Each time a superpixel border is dissolved in standard agglomerative clustering, it modifies the characteristic representations of the pixels within the superpixels and on the boundary. This demands the confidences of the estimator function $h(e)$ on these boundaries be recomputed (Line 1 in Table 1 Algorithm 1). The edges, for which $h(e)$ decreases due to a merge, receive higher priority to be dissolved than it had before. The proposed delayed agglomeration strategy modifies this step and postpones merging these edges for a later time.

2.2 Proposed Delayed Agglomerative Clustering

Our adaption of segmentation commences with the boundary with lowest estimator confidence and repeatedly dissolve edges with in ascending order of $h(e)$. Recall that, $h(e)$ may measure the prediction of a superpixel boundary classifier, or the mean probability values on boundary pixels, or the fraction of overlap between the borders of two superpixels. After two regions have been joined due to a merge, the boundaries of the combined region is updated and the estimator function h is applied to recompute the new confidences. The edges for which $h(e)$ decreases are set aside to be considered at a later stage. They are reexamined after all the borders, initially generated by over-segmentation process, have been checked.

This method is described in Table 2 Algorithm 2. After region R_j is absorbed into R_i , we do not immediately consider all the new boundaries $\{R_i, R_b\}$ between the recently merged R_i and its updated neighbors R_b . We maintain a set of edges W and insert the new edge $\{R_i, R_b\}$ only if its confidence increases from that of $\{R_j, R_b\}$ after R_j is absorbed into R_i (Line 2 in Table 2 Algorithm 2). The faces, for which $h(\{R_i, R_b\})$ decreases from previous value, are kept aside until there are no members left in W and the modified confidence on $\{R_i, R_b\}$ is less than the agglomeration threshold (Line 2 in Table 2 Algorithm 2). Once all $e \in W$ have been considered for merge and W is empty, these boundaries repopulate the list W (Line 2 in Table 2 Algorithm 2) and renew the agglomeration process which continues until there exists no e such that $h(e) \leq \delta$.

Effectively, the proposed strategy ‘delays’ the merging of new edges $\{R_i, R_b\}$ resulting from a merge: either due to an increase in $h(\{R_i, R_b\})$ or deliberately if $h(\{R_i, R_b\})$ decreases. To avoid propagating wrong decisions made on smaller superpixels to the larger ones, this design postpones the merge decisions on the newly formed bodies for a later time. Our analyses support

Table 2. Algorithm 2: Delayed Agglomerative Segmentation.

```

Input:  $S_1, S_2, \dots, S_N$  and confidence function  $h$ .
Output:  $R_1, R_2, \dots, R_N$ 
1 forall the  $i$ , edge  $e$  do  $R_i = S_i$ ,  $\text{Flag}(e) = \text{ACTIVE}$ ;
2 Repeat
3  $W = \{e \in E \mid \text{Flag}(e) = \text{ACTIVE}\}$ ;
4 if  $W$  empty then
5 forall the  $e: \text{Flag}(e) = \text{DELAY}$  do
6 if  $h(e)$  Within range then  $\text{Flag}(e) = \text{ACTIVE}$ ;
7  $e^* \equiv \{R_i, R_j\} = \min_{e \in W} h(e)$ ;
8 Merge  $R_i$  to  $R_j$ , i.e.,  $R_i = \{R_i \cup R_j\}$ , and update  $W$ ;
9 forall the  $R_b \in \text{Nbr}(R_i)$  do
10 if  $h(\{R_i, R_b\}) > h(\{R_j, R_b\})$  then
11  $\text{Flag}(\{R_i, R_b\}) = \text{ACTIVE}$ ;
12 else
13  $\text{Flag}(\{R_i, R_b\}) = \text{DELAY}$ ;
14 until  $h(e) \leq \delta$ ;

```

doi:10.1371/journal.pone.0125825.t002

that deferring decisions on these edges significantly reduces false merges during agglomeration.

2.2.1 Time Complexity. Asymptotically, the running time of the delayed algorithm remains the same as the traditional agglomerative clustering in the worst case. Instead of adding the adjacent boundaries to the priority queue, the delayed algorithm stores them in a separate list. Later, building a queue from this list would require $O(n_1)$ time where the length n_1 of new list must be smaller than that of the previous one (which contains all edges): $n > n_1$.

Our implementation is tuned to reduce the running time of delayed agglomeration. Notice that, a subset of adjacent boundaries is not pushed back or updated into the queue (Line 13 of Table 2 Algorithm 2). We may as well apply a simple trick to avoid updates at each merge altogether: instead of increasing key of the edges with increased h value (Line 11 of Table 2 Algorithm 2), we can postpone the check and increase the key until it becomes a candidate for merge (Line 7 of Table 2 Algorithm 2) or in Line 6 when it being considered to be inserted into W . Thus, we can reduce the computation by $O(dn \log n)$ where d is the degree of S_2 and n is the queue size.

2.3 Proposed Context-aware Segmentation

The proposed context-aware agglomeration is composed of two different phases. We separate the set S_m of potential mitochondria superpixels from the set S_c of potential cytoplasm superpixels assuming the existence of an effective mitochondria superpixel detector (e.g., [22]). The regions in S_c are agglomerated first by the proposed delayed policy. Motivated by [6][10][15], a Random Forest (RF) [23] classifier h_c is trained to act as the boundary predictor function for clustering the set S_c of cytoplasm superpixels. During h_c training, mitochondria-cytoplasm borders are treated the same way as cell membrane.

In the second step, the mitochondria-cytoplasm edges are merged in the same delayed scheme as explained in Section 2.2, but with a different estimator function h_m . In order to absorb mitochondria into corresponding cells, we apply the delayed-agglomeration algorithm

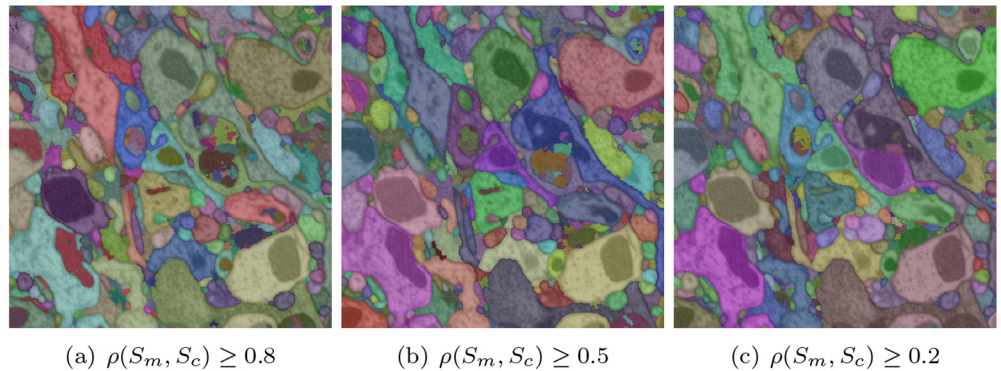


Fig 2. Merging mitochondria into cytoplasm. The figure shows mitochondria superpixels absorbed into cytoplasm superpixels up to different values of overlap ratio $\rho(S_m, S_c)$

doi:10.1371/journal.pone.0125825.g002

with a small alteration. The set of candidate edges W only contains the edges between mitochondria and cytoplasm, that is, $W = \{ \{S_c, S_m\} \mid \text{type}(S_c) = \text{Cyto}, \text{type}(S_m) = \text{Mito}, \text{Flag}(\{S_c, S_m\}) = \text{ACTIVE} \}$; mitochondria-mitochondria edges are not considered for agglomeration. Biologically, each mitochondrion should reside within a cell body. Therefore, boundary confidence for mitochondria merging should reflect how much a mitochondrion is contained within a cytoplasm. In order to quantify this, we define the overlap ratio $\rho(\{S_m, S_c\})$ to be the fraction of the total boundary of S_m which separates S_m from S_c : $\rho(\{S_m, S_c\}) = \frac{\text{length}(\{S_m, S_c\})}{\sum_i \text{length}(\{S_m, S_i\})}$. For any edge $\{S_m, S_c\}$ with a mitochondria superpixel S_m and a cytoplasm superpixel S_c , the confidence is defined as $h_m(\{S_m, S_c\}) = 1 - \rho(\{S_m, S_c\})$.

In effect, the mitochondria superpixels are combined with the cytoplasm superpixels in the descending order of the overlap ratio between these two types of regions. That is, a mitochondria superpixel is merged into the adjacent cytoplasm region with the largest overlap between their boundaries. The combined cytoplasm-mitochondria superpixel created by such merge then identifies the next mitochondria superpixel with the largest overlap to absorb in the next step. We show snapshots of this process, at different values of $\rho(S_m, S_c)$ in Fig 2. It is worth noting that, for 3D segmentation, the overlap is computed across many different planes on which the two cells are neighbors to each other.

3 Results

We have applied the proposed method to EM images of two different modalities: isotropic Focused Ion Beam Scanning Electron Microscope (FIBSEM) data and anisotropic serial section Transmission Electron Microscopy (ssTEM) data. For both types of input data, the image (volume for the isotropic data) is first over-segmented for the agglomeration to be applied on. In the following sections, we explain our over-segmentation process and the error measures used to evaluate segmentation performance before reporting the results on FIBSEM and ssTEM data in Sections 3.3 and 3.4 respectively.

3.1 Over-segmentation and training

We learn a classifier to assign each individual pixel into multiple categories, such as cell boundary, cytoplasm, mitochondria and mitochondria boundary, using the interactive tool Ilastik

[24]. Our pixelwise detector is a Random Forest (RF) classifier [23] trained on a few sparse samples from the dataset. The locations with lowest pixelwise cell boundary prediction are utilized as markers for the Watershed algorithm [12] to produce an over-segmentation of the image/volume. Unless otherwise specified, the same pixel prediction and watershed regions are provided as input to all (competing) methods.

The set \mathcal{S}_m of probable mitochondria superpixels is populated with all regions possessing mean mitochondria probability (estimated by our pixelwise RF classifier trained by Ilastik) above a certain threshold. The rest of the superpixels constitute the set \mathcal{S}_c of possible cytoplasm regions. The training set for superpixel boundary classifier h_c consists of all boundaries among members of \mathcal{S}_c as well as the mitochondria-cytoplasm borders. Similar to [6][15], each superpixel edge is represented by the statistical properties of the multiclass probabilities estimated by Ilastik. The statistical properties include mean, standard deviation, 4 quartiles of the predictions generated for the data locations on the boundary, two regions it separates as well as the differences of these region statistics. All of these features can be updated in constant time after a merge—a property which improves the efficiency of the segmentation algorithm substantially. The code and example dataset are publicly available at <https://github.com/janelia-flyem/NeuroProof.git>.

3.2 Segmentation error measures

We report segmentation error of both types, namely under- and over-segmentation, separately because one of these errors (under-segmentation) is costlier than the other. Split versions of variance of information (VI) [25] and Rand Error (RE) [10] were selected to evaluate segmentation errors. Given a groundtruth (GT), $GT = \{g_1, \dots, g_M\}$, and a segmentation (SG), $SG = \{r_1, \dots, r_P\}$, we compute the over-segmentation (OE) and under-segmentation (UE) errors by splitting the terms in VI and RE. For split-VI, the over and under-segmentation are quantified as follows.

$$VI_{OE} = H(GT | SG) = - \sum_{i,j} \frac{|g_i \cap r_j|}{Z} \log \frac{|g_i \cap r_j|}{|g_i|}. \tag{1}$$

$$VI_{UE} = H(SG | GT) = - \sum_{i,j} \frac{|r_i \cap g_j|}{Z} \log \frac{|r_i \cap g_j|}{|r_i|}. \tag{2}$$

In these equations, $|\cdot|$ denotes the size, \cap denotes the intersection between two regions and Z is a normalizing constant. From information theoretic perspective, these two terms are conditional entropies defined over a set GT given SG, and vice versa.

We also quantify segmentation error by average percentage ($\times 10^{-5}$) of pairs of voxels falsely merged and split by any method. Formally, the over-segmentation (OE) and under-segmentation (UE) is computed based on the following formula.

$$RE_{OE} = \% \text{ pixel pairs within same cluster in GT but different cluster in SG.} \tag{3}$$

$$RE_{UE} = \% \text{ pixel pairs within same cluster in SG but different cluster in GT.} \tag{4}$$

3.3 Segmentation Performances-FIBSEM data

Dataset: The first set of experiments was conducted on isotropic datasets from fruit fly visual system imaged at 10 nm isotropic resolution using FIBSEM technology. This data is segmented

as a volume (i.e., 3D segmentation) and both the voxelwise multi-class predictor and the super-voxel boundary classifier are learned on one 250^3 volume and applied on two 520^3 test volumes.

Competing methods: We have compared the following algorithms in this study: 1) LASH: Standard agglomeration with an RF supervoxel classifier learned based on the iterative procedure of [10]. 2) LASH-D: LASH classifier with delayed agglomeration (proposed extension). 3) GALA [15]: an agglomerative method with repetitive learning phases like LASH, except it accumulates the training sets of multiple phases. 4) CADA-F: Proposed two stage delayed agglomeration with standard RF learned using training set accumulation similar to GALA. 5) CADA-L: Proposed delayed agglomeration with a depth-limited RF (depth = 20) learned *without* training set accumulation. 6) Global Multicut: the optimization framework for finding a closed-surface segmentation proposed in [7]. For [7], the boundary confidences were generated by the CADA-L predictor.

Performance evaluation: In order to compare different supervoxel clustering schemes, we trained (on one 250^3 volume) and segmented two 520^3 volumes 5 times and averaged their scores. We plot the average VI_{UE} and VI_{OE} respectively on x and y-axis respectively in plots on the left column of Fig 3 and for test Volumes 1 and 2. Similarly, we show the average RE_{UE} and RE_{OE} errors on x and y-axis respectively on right columns of Fig 3. In these figures, an ideal algorithm should achieve a zero value for both over and under-segmentation. For all algorithms except the Global multicut method, each point in a plot refers to the boundary confidence threshold $\delta_c \in [0.1, 0.2]$ which was used as stopping criterion for cytoplasm merging. For [7], we instead changed the value of the bias parameter in weight calculation within the range $[0.2, 0.9]$.

As the plots show, both the delayed agglomeration and two-phase segmentation process attained significant improvement over past methods: compare the performance of LASH (red +) with LASH-D (black x) and that of GALA (cyan *) with CADA variants (green square and blue circle). Compared to the rest of the techniques, the two variants of proposed methods, namely CADA-L and CADA-F, appear to achieve the most favorable segmentations by reducing the over-segmentation steeply without increasing the false merge numbers much. During segmentation, the delayed version decreases the time needed for segmentation approximately 5 times among the agglomerative approaches.

It is also worth mentioning that, in a two stage segmentation scheme, the performance of a depth limited RF (i.e., CADA-L, green square), learned without accumulating training set over multiple passes, is very similar to that of the standard RF (CADA-F, blue circle) trained over cumulative learning passes. Training full-depth RF (CADA-F) with multiple passes needed several hours whereas training a depth limited single iteration (CADA-L) required ≤ 5 minutes.

Fig 4 shows three sample planes from the test volume 1. In the following plot in Fig 5, we show example outputs of the methods LASH-D, GALA [15], Global multicut [7] and CADA-L on these planes. Three columns correspond to three planes, and each row presents the outputs of the aforementioned methods. The segmentation labeling is overlaid with artificial (randomly selected) colors. We have selected the parameter that results in the lowest false merges (under-segmentation) with a false split (over-segmentation) error below 0.7 for all except the proposed CADA-L for which we selected the lowest over-segmentation (error value approx 0.56). The results are largely compatible with the quantitative ones. All the three methods, especially Global Multicut, leave many false boundaries intact. The false-splits are not limited to cytoplasm mitochondria borders, both Global and GALA over-segmented some cytoplasm regions as well. By separating these two sub-classes within cell bodies, the proposed method CADA-L was able to eliminate the false merges between them.

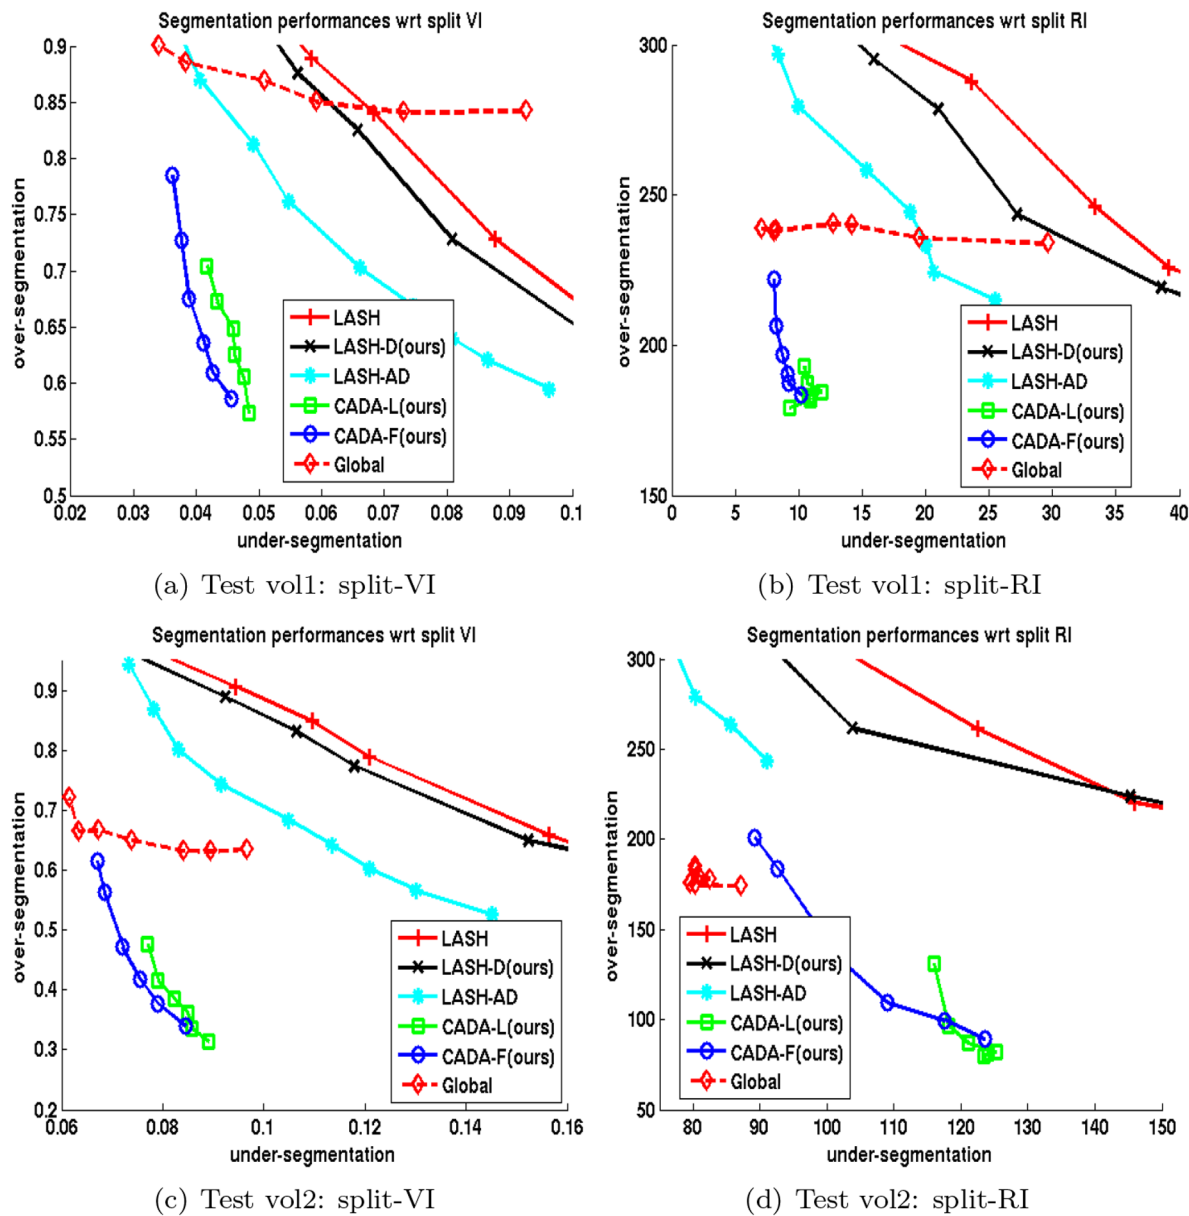


Fig 3. Segmentation error in terms of split-VI and split-RE on two FIBSEM volumes. Top: Test volume 1 and bottom: Test volume 2. Left column shows split-VI error: V_{UE} in x-axis, V_{OE} in y-axis; right column shows split-RE: RE_{UE} in x-axis, RE_{OE} in y-axis. Each curve is the average of results in 5 trials. Each point represents either a stopping point for clustering or bias parameter for [7].

doi:10.1371/journal.pone.0125825.g003

Runtime comparison: In Table 3, we report the running times of the context oblivious standard agglomerative clustering used in LASH [10]; the context-oblivious delayed agglomerated clustering used in LASH-D, GALA [15]; the context-aware delayed agglomeration used in CADA-L, CADA-F; and the Global Multicut [7] method. Both the standard and delayed agglomeration were executed up to the same threshold $\delta = 0.2$. The context-aware method executes two phases of agglomeration, which is why CADA required more time than LASH-D. The Global multicut algorithm utilizes the solution of an optimization problem (requires

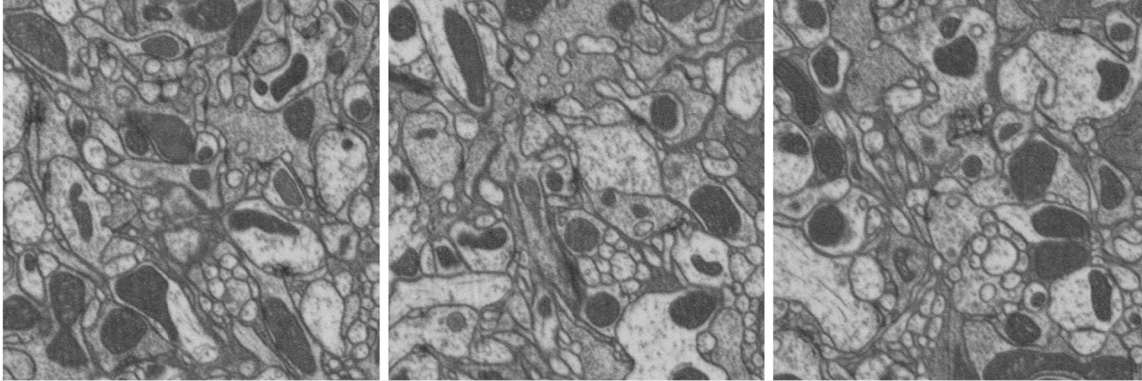


Fig 4. Sample planes from the test volume 1.

doi:10.1371/journal.pone.0125825.g004

optimization packages like CPLEX or Gurobi) in order to find the edges to merge for producing the final segmentation.

All the algorithms, except CADA-L and Global Multicut, perform standard agglomeration multiple times (we repeated 5 times) in order to obtain extensive training sets for superpixel boundary learning. Both CADA-L and Global method exploited the same classifier learned from the initial set of boundaries existed in the over-segmented data (without training set augmentation).

In the following subsections, we analyze why the proposed strategies improve the segmentation performance over the existing approaches.

3.3.1 Context-aware vs Context-oblivious agglomeration. It is perhaps intuitive that traditional context-oblivious agglomeration will result in higher degree of over-segmentation than the context-aware method. The mitochondria-cytoplasm borders indeed have strong feature similarity with cell membranes and consequently superpixel boundary predictors cannot distinguish between these two types of borders perfectly. Recall that, for segmentation, we need to dissolve the mitochondria-cytoplasm border but retain the cell boundaries. In order to substantiate our claim, we trained a superpixel boundary classifier in context-oblivious fashion (0: false cell membrane, 1:true cell membrane) and computed its confidences on these two types of boundaries. [Fig 6](#) shows the histogram of confidence levels for actual cell boundaries and mitochondria-cytoplasm borders in red and blue respectively. If we wish to minimize false merges among neurons, we have to stop agglomeration at a lower value ($\delta \leq 0.3$). The overlap between the two distributions in the range $0.1 \sim 0.5$ suggests that many of the mitochondria borders will not be merged and will lead to over-segmentation.

In addition, due to appearance dissimilarity, the distribution of same features computed on cytoplasm and mitochondria will be substantially different from each other. Combining these two types of feature value distribution will impede the identification of false boundaries between cytoplasm superpixels such as the one in the lower left corner of the output of GALA in [Fig 1](#).

In practice, mitochondria from two different cells could also lead to false merges. Often the mitochondria regions from two cells are closely located to the cell membrane, or other mitochondria regions from neighboring cells, blurring the boundary. [Figs 1](#) and [5](#) show several such locations where the existing techniques failed to avoid false merge.

3.3.2 Global multicut vs Proposed. The split-VI plot in [Fig 7](#) show that both variants of the proposed CADA algorithm generates significantly low under and over-segmentation errors

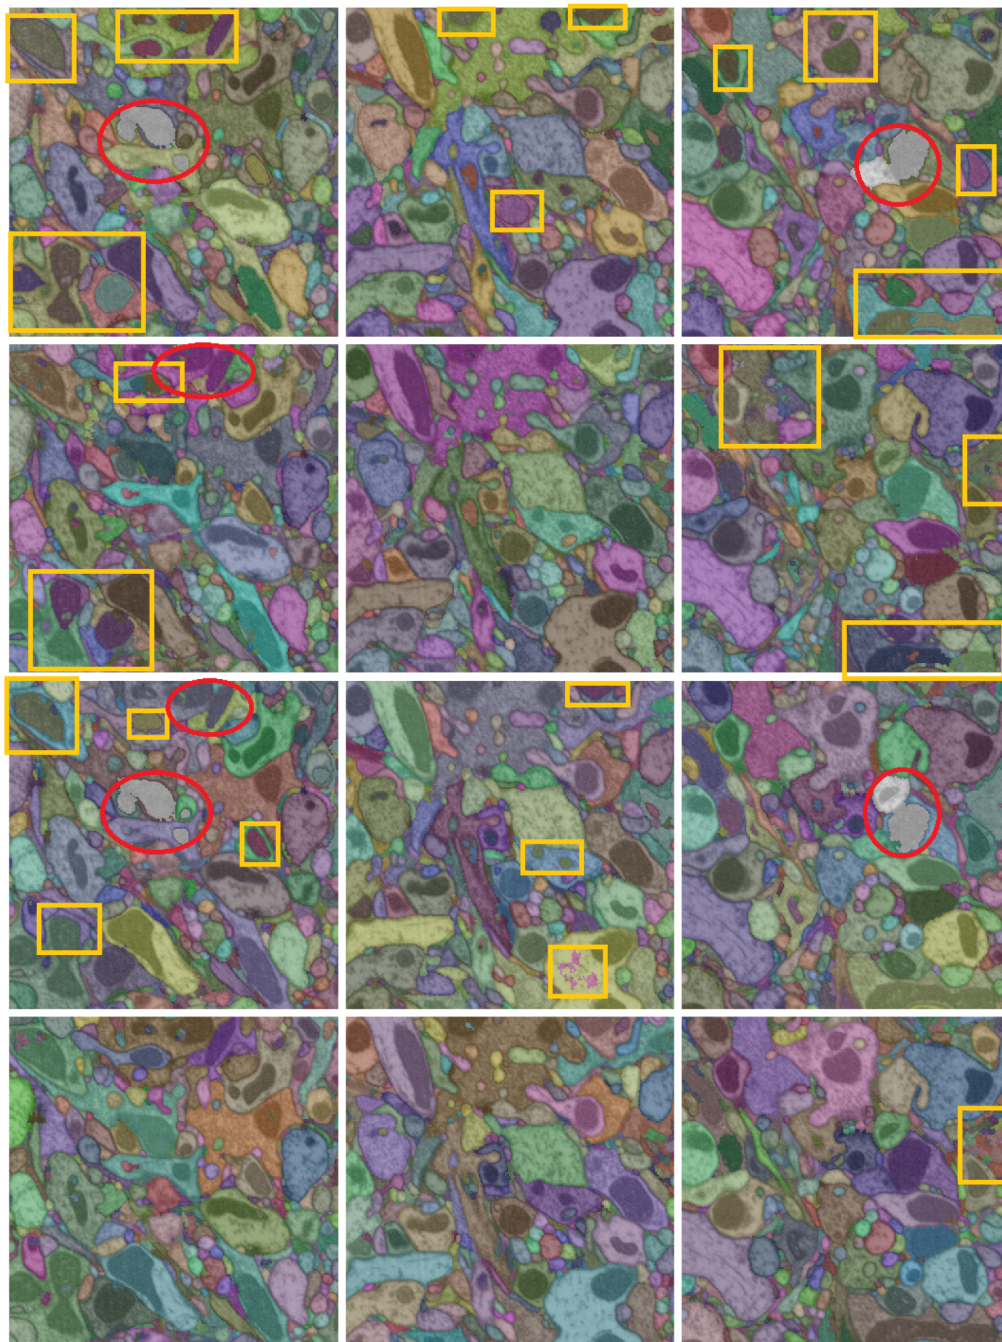


Fig 5. Qualitative results on test volume 1. Three columns show segmentation outputs overlaid with random colors on three planes of the FIBSEM volume. The rows, from top to bottom, show the output of LASH-D, GLObal multicut [7], GALA [15] and the proposed method CADA-L. Some significant over-segmentation errors and under-segmentation errors are marked in yellow rectangles and red ellipses respectively.

doi:10.1371/journal.pone.0125825.g005

Table 3. Runtime for different algorithms.

Method	Run time (min)
LASH: Standard Agglomeration, context-oblivious	5.35 ± 0.2
LASH-D: Delayed Agglomeration, context-oblivious	2.72 ± 0.06
CADA: Delayed Agglomeration, context-aware	4.69 ± 0.02
Global multicut	7.13 ± 1.1

doi:10.1371/journal.pone.0125825.t003

than those of Global [7] method in clustering *cytoplasm regions only (mitochondria not merged)*. In order to analyze why this happens, we save the initial confidences (predictor confidence at the beginning of agglomeration) of $h_c(e)$ on all e that

- were incorrectly split (over-segmented) by the Global method,
- were correctly merged by proposed algorithm.

These boundary predictions were plotted on x-axis of Fig 8. The y-axis of Fig 8 corresponds to confidences $h_c(e)$ at the time e was correctly merged by the proposed method. The threshold on boundary confidences to stop agglomeration was $\delta_c = 0.2$.

Note that, the agglomerative process correctly reduced the confidences of many false boundaries that received a high score by the predictor at the beginning (high x value but low y value). This refinement is possible through the evolution of the superpixels in the agglomerative process—an advantage the Global method of [7] cannot benefit from. The Global method [7], in comparison, generated many more false positive boundaries as depicted by the

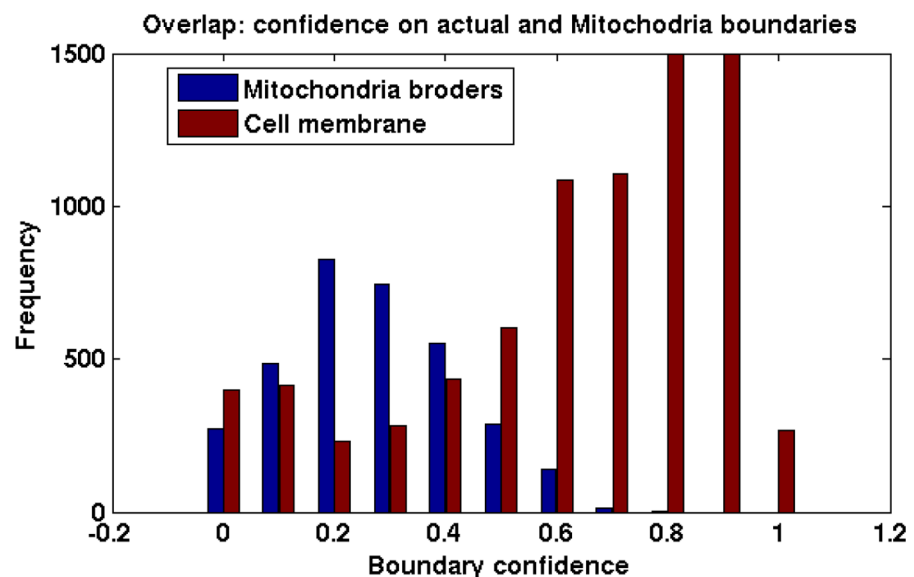


Fig 6. Distribution of predicted boundary confidences on cytoplasm-mitochondria borders (blue) and correct cell boundaries (red). The plot is clipped at $y = 1500$ for better visualization. Notice the overlap between these two distributions within confidence range $[0, 0.6]$.

doi:10.1371/journal.pone.0125825.g006

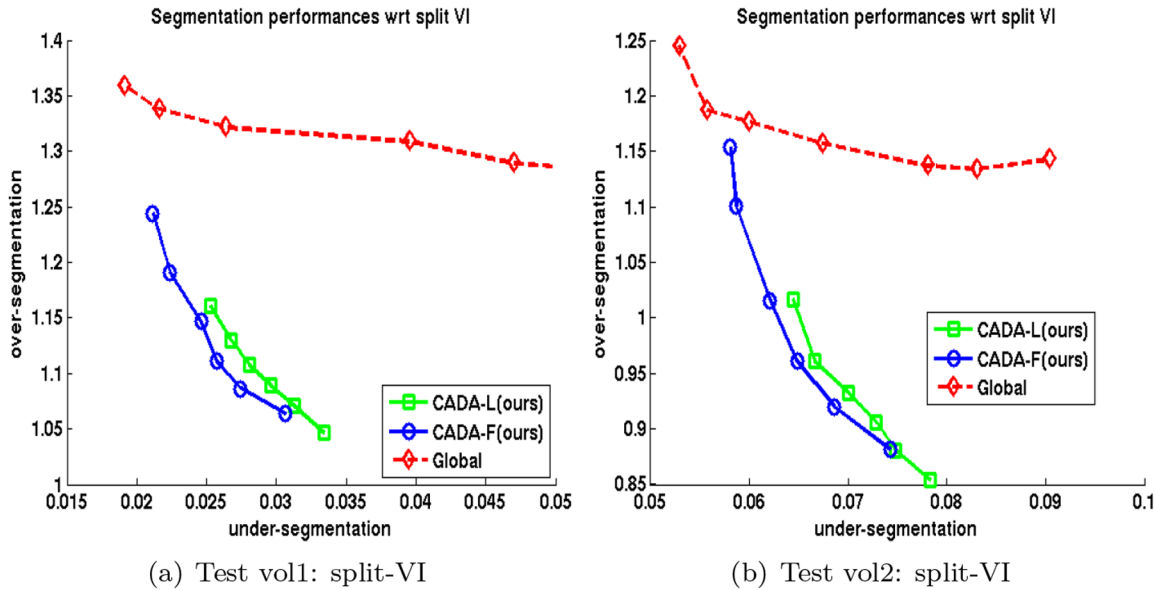


Fig 7. Split-VI of cytoplasm segmentation of two FIBSEM volumes. Left column: test volume 1, right column: test volume 2. Each curve is the average of results in 5 trials. Each point represents either a stopping point for clustering or bias parameter.

doi:10.1371/journal.pone.0125825.g007

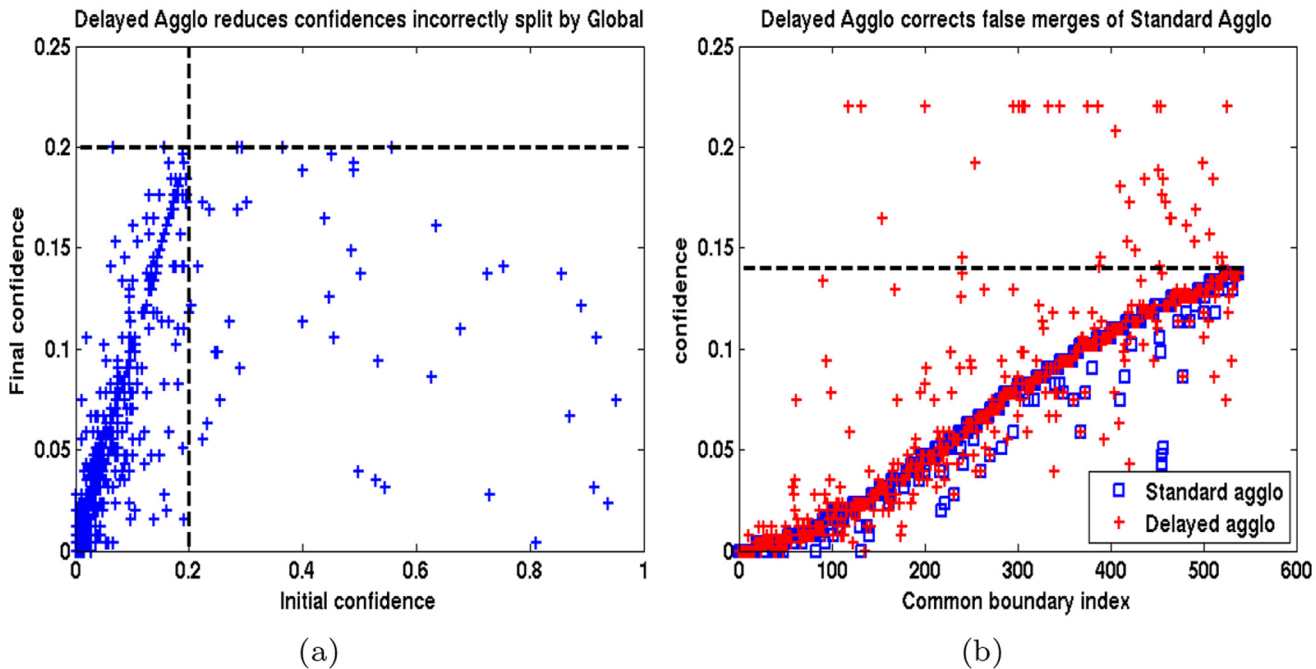


Fig 8. Left: False splits (over-segmentation) of Global method corrected by proposed CADA-L. Each point corresponds to a false boundary that Global method failed to dissolve. The x-axis labels indicate the predictor confidence at the beginning of the proposed agglomeration and y-axis plots the predictor confidence at the point it was merged accurately by the agglomeration. Right: False merges (under-segmentation) of Standard agglomeration corrected by delayed method—x-axis: boundary indices, y-axis: predictor confidence. The confidences computed for the same correct edge in traditional agglomeration and in the proposed delayed version is plotted in blue square and red '+'. The confidences on many true boundaries were increased by the delayed approach.

doi:10.1371/journal.pone.0125825.g008

rectangular enclosed region of Fig 8 (left). If several boundaries within a chain of supervoxel faces receive very high predictor confidences, then, by construction, the Global method tends to retain the another boundary e within the same chain with a low $h_c(e)$. Such tendency may be the reason behind the high concentration of false splits with low h_c within the rectangular region in Fig 8 (left).

3.3.3 Delayed vs standard agglomeration. In order to illustrate the improved accuracy attained by the delayed agglomeration over the standard one, we collected all faces that were incorrectly dissolved by standard agglomeration algorithm (LASH) and examine their confidences under a delayed scheme (LASH-D) operating at $\delta_c = 0.14$. The confidences (clipped to 0.25) of these 534 edges generated by standard and delayed agglomeration are plotted in Fig 8 (right) in blue square and red + respectively. The proposed delayed agglomeration accurately increased the confidences h_c of many of these faces, among which, 41 exceeded the threshold of 0.14 (green line) and avoided a false merge. In addition to these common supervoxel edges, the standard and delayed algorithms independently generated 163 and 4 more incorrect merges respectively.

3.4 Segmentation Performances-ssTEM data

This section reports the 2D segmentation results that our method and others produced on a different data modality, namely ssTEM images. These images were part of those generated for the work of [2] and were collected from the authors. Fifteen 500×500 images were used for training both the pixel and superpixel boundaries. We follow the techniques and 2D versions of features described in Section 3.1 for ssTEM data. The same pixel prediction and watershed regions are provided as input to all competing methods. The segmentation is performed on each image without connecting them across planes. Fig 9 plots the average of split-VI and split-RE errors over 15 images of size 1000×1000 of the proposed CADA-L and GALA [15]

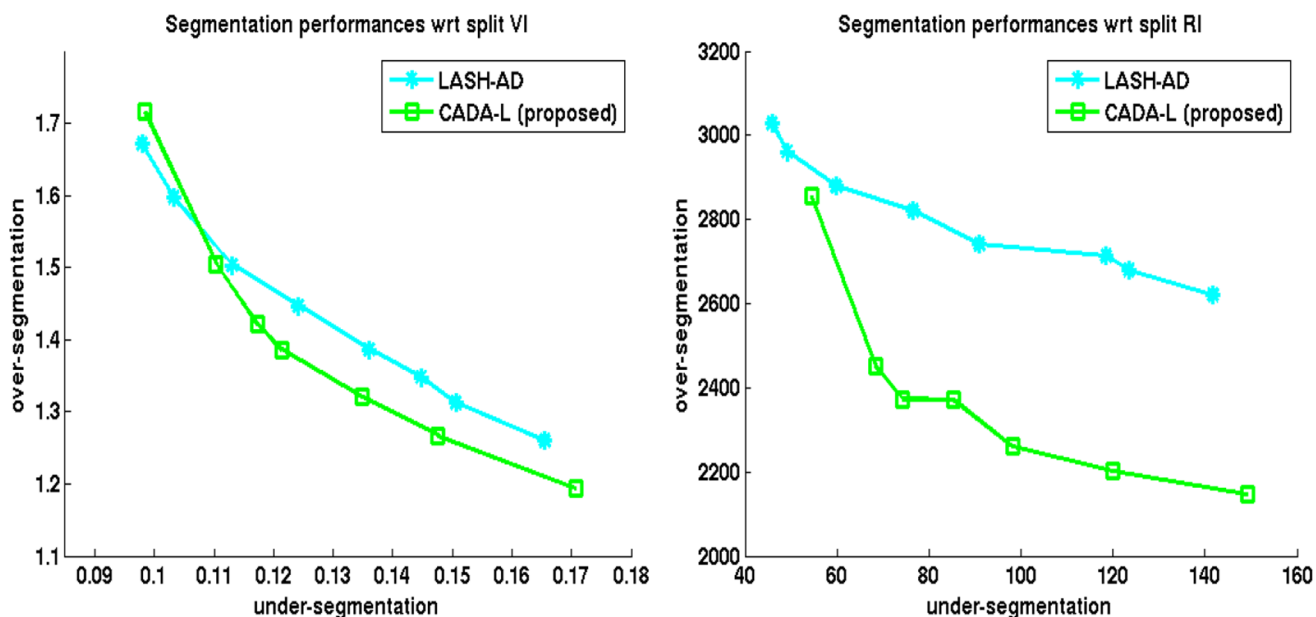


Fig 9. Segmentation errors on TEM data. Left column shows split-VI error: V_{UE} in x-axis, V_{OE} in y-axis; right column shows split-RE: RE_{UE} in x-axis, RE_{OE} in y-axis. The curves are averages of errors on 15 1000×1000 images. The results of Global method [6] were too poor to plot.

doi:10.1371/journal.pone.0125825.g009

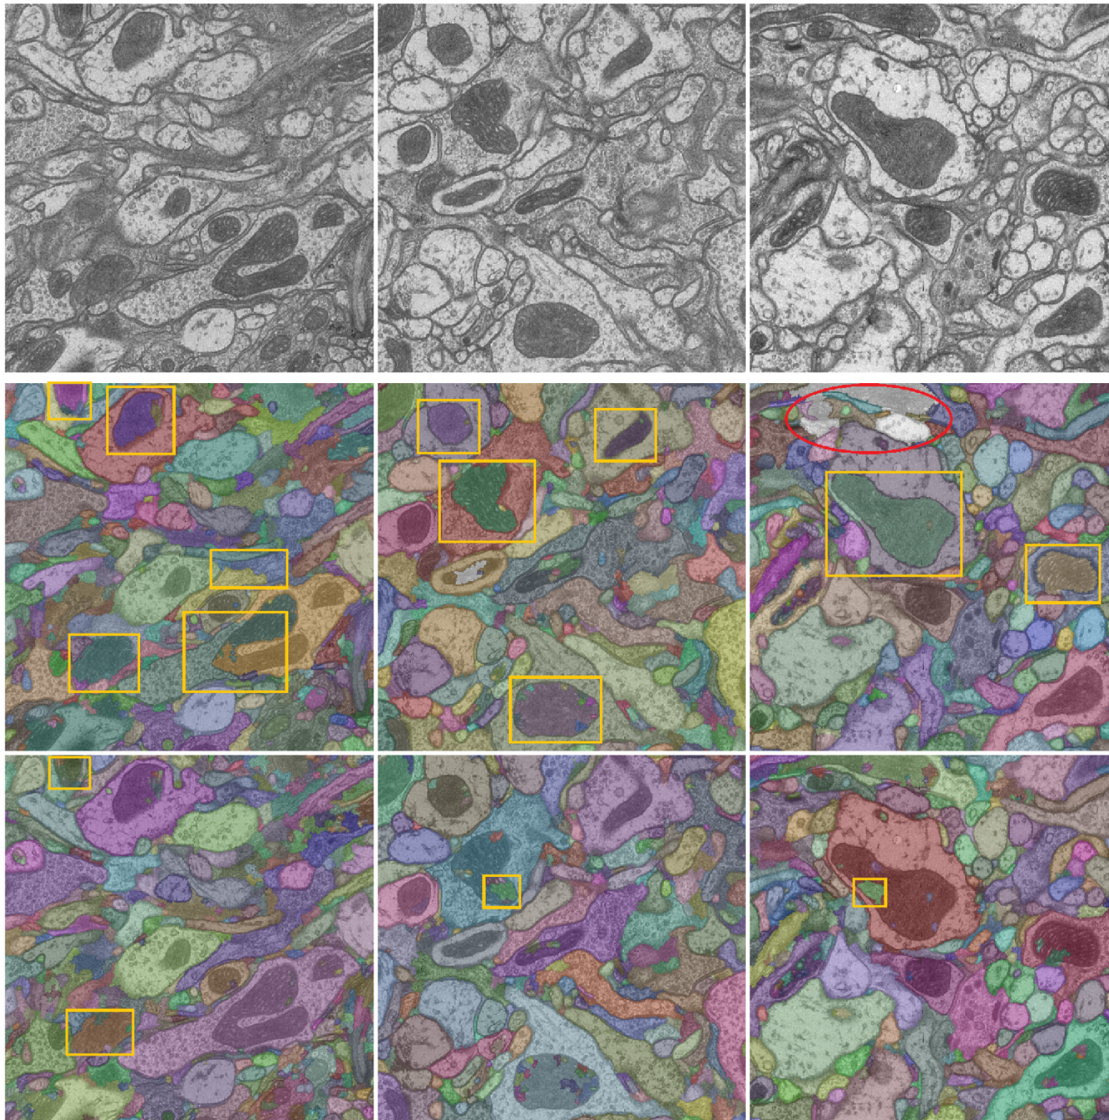


Fig 10. Qualitative comparison of GALA and proposed method segmentation on TEM data. The segmentation outputs are overlaid with random colors on the grayscale images. Top row: input, middle GALA and bottom: proposed CADA-L. Significant over-segmentation errors and under-segmentation errors are marked in yellow rectangles and red ellipses respectively.

doi:10.1371/journal.pone.0125825.g010

methods. Our method CADA-L produces less over-segmentation than GALA [15] in almost all threshold values. The result of the Global method [7] were too poor to show on this plot—lowest over-segmentation error at 4.11 with 0.13 under-segmentation average.

In Fig 10, we show input images and the segmentation results (overlaid on the image) of GALA and our methods at the same under-segmentation error. Examining the qualitative output in Fig 10, GALA seems to struggle to absorb the mitochondria regions despite multiple learning iterations and even merges two cells in one occasion. While a more accurate mitochondria predictor could potentially reduce the segmentation errors of the proposed method, context-invariant algorithms such as GALA would be less effective around mitochondria

regions. Compared to GALA, CADA-L used less than half of the training examples (42.46%) collected without training iterations (i.e., significantly more efficient in training).

4 Discussion

We argue that, due to considerable ambiguity in appearances, it is only rational for an EM segmentation algorithm to be context-aware in each of its stages, i.e., in both pixel and superpixel levels (and in alignment for anisotropic data). The results reported in this paper support our claim that a context-aware clustering of sub-classes such as cytoplasm and mitochondria can improve segmentation accuracy significantly given fairly accurate sub-class detection. Our examination of both isotropic and anisotropic data suggests cell structures cannot be meaningfully identified without mitochondria regions and it is non-trivial to combine detection with a segmentation that ignores it (e.g., [7]) in order to produce the final segmentation. Our analysis also illustrates how a delayed agglomerative procedure benefits from the intermediate boundary probabilities and improves the efficiency of the segmentation process significantly.

In addition to reducing the over- and under-segmentation errors, one of the variants of our classifier, namely CADA-L, can be trained considerably faster than those in other methods because CADA-L demands substantially fewer training examples and no training iterations. A context-oblivious strategy gain significantly (compare LASH-D with GALA in Fig 3) by accumulating training set over multiple iterations. However, in context-aware approach, one does not benefit much by accumulating the training set (CADA-F in Fig 3) over a classifier trained from a single iteration (CADA-L). One possible explanation is that previous context oblivious strategies require the extra iterations to mitigate the impact of the noise introduced by mitochondria superpixels. This explanation implies that detecting the sub-classes, and considering them separately as necessary, is perhaps the key to train a boundary classifier accurately and efficiently.

We further investigated this conjecture and developed a semi-supervised active learning algorithm to train the supevoxel boundary classifier with as few as < 20% of the total examples [20]. The requirement of exhaustive labels is a critical bottleneck for automatic EM segmentation, especially for reconstructing larger brain regions, or whole animal brain, where one may anticipate the necessity to train several different classifiers [21]. The interactive training of both pixels (using Ilastik [24] for example) and superpixel boundaries (using [20]) holds the promise of removing the need for such complete groundtruth and paves the way for scaling up the EM reconstruction algorithms.

We have applied our context aware algorithm to segment 216 FIBSEM volumes of 520^3 voxels each, with a 10nm isotropic resolution, from the Medulla region of fly retina. To our knowledge, this is an attempt to reconstruct one of the largest volumes for such animal. Compared to the result of [15] on two of the 520^3 blocks, our segmentation resulted in an estimated 30% reduction in subsequent manual correction time. In addition, our segmentation was sufficiently accurate for regions that pertain to Post-synaptic densities (PSD), i.e., the synaptic partners of a cell. During the manual annotation of these PSDs, the output of our segmentation method assisted the experts to improve their performances [26].

Acknowledgments

The authors wish to thank Stuart Berg and Bill Katz for their contributions in software development; Pat Rivlin and Shinye Takemura for providing the annotated dataset of [2] and for the discussion on the neurobiological properties of our dataset; Matt Saunders and Lei-Ann Chnag for their effort in generating pixel classification ilps and groundtruth annotation; Don Olbris

for his support with the visualization software Raveler and Margaret Jefferies for her assistance in writing.

Author Contributions

Conceived and designed the experiments: TP AC. Performed the experiments: TP AC. Analyzed the data: TP. Contributed reagents/materials/analysis tools: TP AC SP. Wrote the paper: TP SP LS. Assisted with writing and analysis: LS.

References

- Helmstaedter M., Briggman K., Turaga S., Jain V., Seung S., Denk W.: Connectomics reconstruction of the inner plexiform layer in the mouse retina. *Nature* 500 (7461) (2013) 168–174 doi: [10.1038/nature12346](https://doi.org/10.1038/nature12346) PMID: [23925239](https://pubmed.ncbi.nlm.nih.gov/23925239/)
- Takemura S. Y., Bharioke A., Lu Z., Nern A., Vitaladevuni S., Rivlin P. K., et al.: A visual motion detection circuit suggested by *Drosophila* connectomics. *Nature* 500 (7461) (2013) 175–181 doi: [10.1038/nature12450](https://doi.org/10.1038/nature12450) PMID: [23925240](https://pubmed.ncbi.nlm.nih.gov/23925240/)
- Arbelaez P., Maire M., Fowlkes C., Malik J.: Contour detection and hierarchical image segmentation. *PAMI, IEEE Transactions on* 33(5) (2011) 898–916 doi: [10.1109/TPAMI.2010.161](https://doi.org/10.1109/TPAMI.2010.161)
- Krahenbuhl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. In: NIPS. (2011)
- Ladicky, L., Russell, C., Kohli, P., Torr, P.H.S.: Associative hierarchical crfs for object class image segmentation. In: ICCV. (2009)
- Andres B., Köthe U., Helmstaedter M., Denk W., Hamprecht F.: Segmentation of SBFSEM Volume Data of Neural Tissue by Hierarchical Classification. *Pattern Recognition* 5096(15) (2008) 142–152 doi: [10.1007/978-3-540-69321-5_15](https://doi.org/10.1007/978-3-540-69321-5_15)
- Andres, B., Kroeger, T., Briggman, K., Denk, W., Korogod, N., Knott, G., Koethe, U., Hamprecht, F.: Globally optimal closed-surface segmentation for connectomics. In: ECCV. (2012)
- Chklovskii D.B., Vitaladevuni S., Scheffer L.K.: Semi-automated reconstruction of neural circuits using electron microscopy. *Current Opinion in Neurobiology* 20(5) (2010) 667–675 doi: [10.1016/j.conb.2010.08.002](https://doi.org/10.1016/j.conb.2010.08.002) PMID: [20833533](https://pubmed.ncbi.nlm.nih.gov/20833533/)
- Funke, J., Andres, B., Hamprecht, F., Cardona, A., Cook, M.: Efficient automatic 3d-reconstruction of branching neurons from em data. In: CVPR. (2012)
- Jain V., Turaga S.C., Briggman K., Helmstaedter M.N., Denk W., Seung H.S.: Learning to agglomerate superpixel hierarchies. In: NIPS 24. (2011) 648–656
- Vazquez-Reina, A., Gelbart, M., Huang, D., Lichtman, J., Miller, E., Pfister, H.: Segmentation fusion for connectomics. In: ICCV. (2011)
- Beucher, S., Meyer, F.: The Morphological Approach to Segmentation: The Watershed Transformation. *Mathematical Morphology in Image Processing* (1993) 433–481
- Liu, T., Jurrus, E., Seyedhosseini, M., Ellisman, M., Tasdizen, T.: Watershed merge tree classification for electron microscopy image segmentation. In: ICPR. (2012)
- Uzunbas, M., Chen, C., Metaxas, D.: Optree: A Learning-Based Adaptive Watershed Algorithm for Neuron Segmentation. In: MICCAI. (2014)
- Nunez-Iglesias J., Kennedy R., Parag T., Shi J., Chklovskii D.B.: Machine learning of hierarchical clustering to segment 2d and 3d images. *PLoS ONE* 8(8) (2013). doi: [10.1371/journal.pone.0071715](https://doi.org/10.1371/journal.pone.0071715) PMID: [23977123](https://pubmed.ncbi.nlm.nih.gov/23977123/)
- Ciresan, D.C., Giusti, A., Gambardella, L.M., Schmidhuber, J.: Deep neural networks segment neuronal membranes in electron microscopy images. In: NIPS. (2012)
- Jain, V., Bollmann, B., Richardson, M., Berger, D., Helmstaedter, M., Briggman, K., et al: Boundary learning by optimization with topological constraints. In: CVPR. (2010)
- Jurrus E., Paiva A., Watanabe S., Anderson J., Jones B., Whitaker R., et al: Detection of neuron membranes in electron microscopy images using a serial neural network architecture. *Medical Image Analysis* 14(6) (2010) 770–783 doi: [10.1016/j.media.2010.06.002](https://doi.org/10.1016/j.media.2010.06.002) PMID: [20598935](https://pubmed.ncbi.nlm.nih.gov/20598935/)
- Liu T., Jones C., Seyedhosseini M., Tasdizen T.: A modular hierarchical approach to 3D electron microscopy image segmentation. *Journal of Neuroscience Methods* 226(15) (2014) 88–102. doi: [10.1016/j.jneumeth.2014.01.022](https://doi.org/10.1016/j.jneumeth.2014.01.022) PMID: [24491638](https://pubmed.ncbi.nlm.nih.gov/24491638/)

20. Parag, T., Plaza, S., Scheer, L.: Small sample learning of superpixel classifiers for EM segmentation. In: MICCAI. (2014)
21. Helmstaedter M.: Cellular-resolution connectomics: challenges of dense neural circuit reconstruction. *Nat Methods* 10(6) (2013) 501–7 doi: [10.1038/nmeth.2476](https://doi.org/10.1038/nmeth.2476) PMID: [23722209](https://pubmed.ncbi.nlm.nih.gov/23722209/)
22. Lucchi A., Smith K., Achanta R., Knott G., Fua P.: Supervoxel-based segmentation of mitochondria in em image stacks with learned shape features. *Medical Imaging, IEEE Transactions on* 31(2) (2012) 474–486 doi: [10.1109/TMI.2011.2171705](https://doi.org/10.1109/TMI.2011.2171705)
23. Breiman L.: Random forests. *Machine Learning* 45(1) (October 2001) 5–32 doi: [10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324)
24. Sommer, C., Straehle, C., Koethe, U., Hamprecht, F.A.: “ilastik: Interactive learning and segmentation toolkit”. In: ISBI. (2011)
25. Meila, M.: Comparing clusterings by the variation of information. In: COLT’03. (2003) 173–187
26. Plaza, S., Parag, T., Huang, G., Olbris, D., Saunders, M., Rivlin, P.: Annotating Synapses in Large EM Datasets. arXiv:1409.1801. (2014)