IMPROVING 3D EM DATA SEGMENTATION BY JOINT OPTIMIZATION OVER BOUNDARY EVIDENCE AND BIOLOGICAL PRIORS

N. Krasowski[⋆] *T. Beier*[⋆] *G. W. Knott*[†] *U. Koethe*[⋆] *F. A. Hamprecht*[⋆] *A. Kreshuk*[⋆]

[⋆] HCI/IWR, University of Heidelberg, Germany
[†] Ecole Polytechnique Fédérale de Lausanne, Switzerland

ABSTRACT

We present a new automated neuron segmentation algorithm for isotropic 3D electron microscopy data. We cast the problem into the *asymmetric multiway cut* framework. The latter combines boundary-based segmentation (clustering) with region-based segmentation (semantic labeling) in a single problem and objective function. This joint formulation allows us to augment local boundary evidence with higher-level biological priors, such as membership to an axonic or dendritic neurite. Joint optimization enforces consistency between evidence and priors, leading to correct resolution of many difficult boundary configurations. We show experimentally on a FIB/SEM dataset of mouse cortex that the new approach outperforms existing hierarchical segmentation and multicut algorithms which only use boundary evidence.

Index Terms— Segmentation, Electron Microscopy, graphical model

1. INTRODUCTION

Connectomics is a domain of neuroscience, which seeks to establish relation between structure (physical connectivity) and function of neural circuits. While large-scale EM imaging, necessary for connectomics experiments, is now close to being a solved problem, image *analysis* and actual reconstruction of neural circuits from image volumes still requires an enormous amount of manual processing [1]. Our contribution aims to reduce the manual processing load by proposing an algorithm for automated segmentation of neurons in image volumes of high isotropic resolution, such as the ones produced by a FIB/SEM microscope [2]. An example of such a dataset is shown in Fig. 1 (left).

Our main contribution is an approach to introduce higher-level biological priors into a low-level segmentation procedure. Like most current automated segmentation methods, we start from an over-segmentation of the volume into 3-dimensional superpixels. But, when deciding if two superpixels need to be merged, our algorithm is not limited to local superpixel affinity. In particular, we exploit the fact, that in mammalian cortex neurites can be part of an axon or of a dendrite, but not both (current high-resolution mammalian image volumes do not encompass the neuron soma). While axons and dendrites are generally not distinguishable on the superpixel level, some superpixels contain sparse indicators for one or the other, such as neurotransmitter vesicle clouds. This sparse neuron type information is used to implicitly introduce long range repulsion between axonic and dendritic superpixels in our model. The MAP state of the graphical model can be found by solving the asymmetric multiway cut problem [3]. Inference based on cutting planes allows us to solve this problem to global optimality.

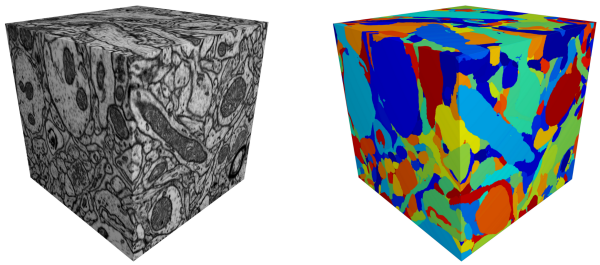

Fig. 1. The raw 700^3 px (approx. $50\mu m^3$) test volume (left) and the automated segmentation results (right). In total, the block contains 374 neurites, connected by 71 synapses.

2. RELATED WORK

Automated neuron segmentation algorithms can roughly be divided into two groups: those that perform the segmentation in 2D with subsequent linking along the z-axis, and those that work in 3D directly. The former are usually applied to data with anisotropic resolution [4, 5, 6, 7], while the latter, including the algorithm proposed in this contribution, are best suited for isotropic data [8, 9, 10, 11]. All modern automated segmentation algorithms are based on supervised learning, which is either applied to find the neuron membranes directly [11] or to learn the affinity of pixels or superpixels and then cluster them into separate objects [8, 9, 10]. For all these methods, learning and/or clustering is performed based on local information only and no high-level biological priors are used to

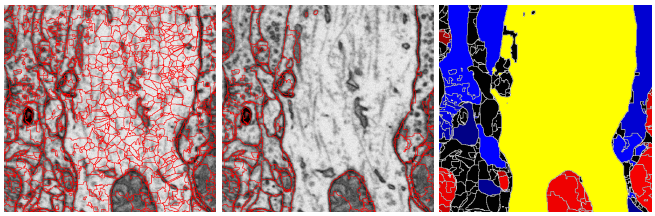

Fig. 2. Left: small superpixels, obtained by watershed on membrane probability map. Center: large superpixels, obtained by hierarchical clustering. Right: probabilities for "dendrite" (yellow), "mitochondria" (red) and "vesicle cloud" (blue) classes, assigned to superpixels.

improve the segmentation.

Our work builds on that of [9], but uses the asymmetric multiway cut [3] instead of the multicut partitioning. The new formulation allows us to introduce sparse implicit long-range repulsive potentials between some of the model variables and thus exploit prior information on the pre- or post-synaptic role of the neurons.

3. METHODS

3.1. Superpixel generation

To generate 3D superpixels, we start by creating a pixelwise membrane probability map for the volume [12]. The probability map is then denoised by the non-local means algorithm [13] and its minima are taken as seeds for watershed segmentation. Seeded watershed partitions the volume into very small superpixels, which we then agglomerate by applying the GALA hierarchical clustering algorithm [10]. Here we can stop at a conservative threshold to ensure oversegmentation. The edge weights for GALA are found by a Random Forest trained on the features of the superpixels adjacent to the edge. Fig. 2 illustrates the algorithm. Compared to initial watershed superpixels (Fig. 2 (left)), GALA superpixels (Fig. 2 (center)) greedily incorporate many "easy" merge decisions and significantly reduce the problem space for the next, globally optimal, step of the pipeline.

3.2. Biological Priors

In general, axonic and dendritic cytoplasm is not locally distinguishable even for domain experts. However, for some of the superpixels an educated guess can be made regarding their axonic or dendritic nature: i) at the synaptic contact locations, superpixels on the presynaptic side correspond to axons, on the post-synaptic side to dendrites; ii) large vesicle clouds are primarily found in axonic superpixels; iii) large "empty" areas are primarily found in the dendritic shafts (axons can also be fairly thick at presynaptic varicosities, but then, unlike empty dendrites, they are filled with vesicles). Based on this evidence, we propose the following procedure for assigning neuron type to superpixels:

1. Detect synapses in the volume, along with their directions [14, 15]. Assign superpixels, touching the synapse, to either dendritic or axonic class, depending on which side of the synapse they are found.
2. Perform pixelwise classification of the volume with ilastik [12], using classes "vesicle cloud", "dendritic cytoplasm" and background.
3. For smaller superpixels (as found by watershed before agglomerative clustering is applied), assign the maximum pixelwise vesicle/dendrite probability to the superpixel. Then, for larger superpixels, assign the average of smaller superpixels.

Empirically, we observed that introducing a separate class for mitochondria superpixels improves segmentation accuracy. Therefore, we also perform pixel classification into "mitochondria" and "background" classes and assign to superpixels the mean mitochondria probability of their pixels. Fig. 2 (right) shows the priors in a small region of raw data, and projected on superpixels.

3.3. Graphical model

To solve the neuron segmentation problem, we construct a probabilistic graphical model, where superpixels from Sec. 3.1 are represented as random variables. Pairwise factors are introduced between the neighboring superpixels, representing the likelihood that the superpixels should be merged. Unary factors for each superpixel reflect the cost of assignment to one of the three classes: "axon", "dendrite" or "mitochondrion". On this superpixel graph we then simultaneously solve the partitioning problem and the semantic segmentation problem with C classes to global optimality. Our motivation for the inclusion of the semantic segmentation problem is that in difficult cases, where the edge evidence by itself is inconclusive, semantic labels bring additional evidence that certain superpixels should belong to different elements of the final partitioning. The work of [3], where the asymmetric multiway cut problem was originally introduced, shows that solving the semantic segmentation problem first and following it with the partitioning problem is not equivalent to the joint solution.

More formally, define a graph $G = G(\mathcal{V}, \mathcal{E})$, formed by the region adjacency graph of the N superpixels. The label l_i of the superpixel i can then be presented as:

$$l_i = \left(l_i^p, l_i^s\right)^T,\ l_i^p \in [1, \cdots, N],\ l_i^s \in [1, \cdots, C], \qquad (1)$$

where l_i^p corresponds to the partitioning problem and l_i^s corresponds to the semantic segmentation problem. Since the number of partitions is not known beforehand we have to reserve N different labels to account for the possibility that no superpixels are merged.

The MAP state of the graphical model can be found by minimizing the following energy function:

$$\underset{l}{\operatorname{argmin}} \left(w \sum_{i \in \mathcal{V}} E_i(l_i^s) + \sum_{(i,j) \in \mathcal{E}} E_{ij}(l_i, l_j) \right). \quad (2)$$

The unary terms depend only on the semantic labeling and the respective prior probabilities p_i^c for the C classes, computed as described in Sec. 3.2:

$$E_i(l_i^s = c) = \begin{cases} -\log\left(\frac{p_i^c}{1-p_i^c}\right) & \text{if } c = \text{“mito”} \\ -\log\left(\frac{\max(p_i^c, 0.5)}{\min\left(1-p_i^c, 0.5\right)}\right) & \text{else} \end{cases} \quad (3)$$

For classes other than mitochondria our priors are sparse, so we only consider positive evidence for the “axon” and “dendrite” classes. In other words, seeing no evidence for these classes brings no penalty for assigning a label to one of them. The pairwise potentials are constructed in the following way:

$$E_{ij}\left(l_i, l_j\right) = \begin{cases} 0 & \text{if } l_i^p = l_j^p, l_i^s = l_j^s \\ -\log\left(\frac{p_{ij}}{1-p_{ij}}\right) & \text{if } l_i^p \neq l_j^p \\ \infty & \text{if } l_i^p = l_j^p, l_i^s \neq l_j^s \end{cases} \quad (4)$$

The Random Forest classifier is used to compute the pseudo-probabilities p_{ij} that superpixels i and j are separated in the final partitioning. Features for this prediction are computed from the superpixels i and j similar to the method, described in [9]. If the semantic segmentation part of the problem is ignored, it reduces to the multicut partitioning or correlation clustering of [9], optimization of which is already an NP-hard problem. Introduction of the unary terms brings the problem into the multiway cut framework [16], while allowing for edges between objects of the same class makes it possible to formulate it as an asymmetric modified multiway cut [3]. In order to actually solve this problem, we transform it into an equivalent graphical model where semantic labels are represented by auxilliary terminal nodes. All random variables and constraints are associated with the edges of the augmented graph, and the resulting optimization problem is solved using the cutting-plane algorithm from [3].

3.4. Post-processing

Since our aim is to aid in the reconstruction of connectomes, where mitochondria do not play any special role, our pipeline is augmented by a post-processing step, which merges objects of mitochondria class objects of axon or dendrite class, which surround them. We also merge all small objects that are completely contained in a larger object into this larger object.

4. EXPERIMENTS

The approach was tested on a FIB/SEM dataset from adult mouse somatosensory cortex with almost isotropic resolution of $5 \times 5 \times 6$ nm. The training of all classifiers was performed on a $900 \times 900 \times 200$ px sub-volume, while the tests were run on a different 700^3 px subvolume of the same tissue block. The ground truth segmentation was performed by an expert neuroscientist, using the interactive segmentation algorithm of [17].

Three algorithms have been compared: hierarchical clustering (HC) as in [10], multicut (MC) as in [9] and multiway cut-based (AMWC), proposed in this contribution. Edge probabilities for hierarchical clustering and multicut were obtained by training a Random Forest classifier on neuron membranes only (mitochondria membranes labeled “off”), while for the multiway cut, where mitochondria constitute a separate class, mitochondria membranes were additionally labeled as “on”. To obtain large superpixels for the multicut and multiway cut algorithms, we performed several levels of hierarchical clustering and selected the thresholds so that training and test datasets have approximately the same number of superpixels per unit volume without visible merge errors. For the hierarchical clustering experiment, we continued the clustering procedure until edge probability reached 0.5. Our own re-implementation of the algorithm in [10] was used for this experiment. The best w parameter of the multiway cut-based algorithm was found by grid search and set to 300. Merging of small, fully contained objects as described in Sec. 3.4 was performed for both MC and AMWC.

Table 1 shows the Rand Index and Variation of Information estimates for the three algorithms, as well as for the direct projection of the pixelwise groundtruth on the large superpixels, used for the multicut and multiway cut algorithms.

	HC	MC	AMWC	GT
RI	0.987	0.990	**0.992**	0.997
VI	0.894	0.798	**0.789**	0.406

Table 1. Segmentation quality of the hierarchical clustering, the multicut and the asymetic multiway cut against the groundtruth segmentation, as measured by the Rand Index (RI, higher is better) and Variation of Information (VI, lower is better). For comparison, the last column (GT) shows the pixelwise groundtruth, projected on the superpixels, used as input for MC and AMWC methods.

5. DISCUSSION

We have presented an approach to incorporate sparse high-level biological priors into the multicut-based automated segmentation algorithm of [9]. The global pre- or post-synaptic prior for the whole segmented neurite can be automatically induced from local evidence for characteristic cell organelles or cytoplasm appearance. Fig. 3 shows two typical examples of the improvement, brought by the new approach. At certain locations the edge probability can not be estimated correctly due to, for instance, a mitochondrion located close to the cell

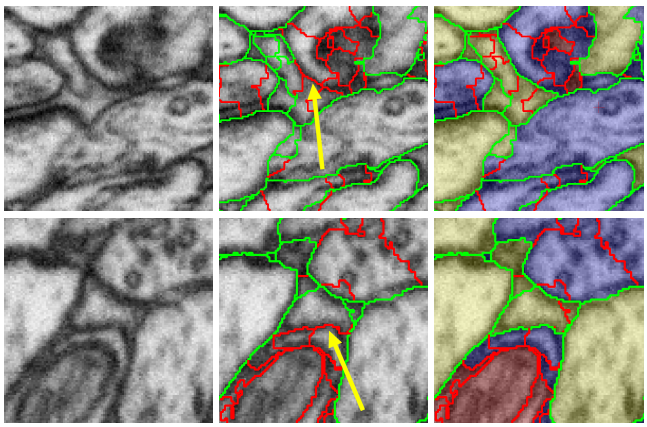

Fig. 3. Left: a difficult location in the raw data. Center: multicut fails to corectly close the edge, pointed out by the arrow. Right: the asymmetric multiway cut overcomes this problem, since the two neurons, separated by the edge, belong to different classes (shown by different object background colors).

membrane or other variations in the local membrane appearance. In this case, the multicut-based algorithm prefers to remove the edge for which it does not have sufficient evidence. The AMWC algorithm, however, has additional information, that the respective superpixels likely belong to different neuron types. This information makes removal of the edge much more costly (a penalty has to be paid for switching the superpixel class labels) and thus pushes the algorithm to the correct solution. So far, we limited the local evidence we consider to synapse locations, vesicle clouds and large areas of dendritic cytoplasm. In future, we plan to incorporate more local prior information, and, to make the approach scale better to very large datasets, to work on approximate solvers for the multiway cut problem.

6. REFERENCES

[1] M. Helmstaedter, "Cellular-resolution connectomics: challenges of dense neural circuit reconstruction," *Nature Methods*, 2013.

[2] G. Knott, H. Marchman, D. Wall, and B. Lich, "Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling.," *The Journal of Neuroscience*, 2008.

[3] T. Kroeger, J. Kappes, T. Beier, U. Koethe, and F.A. Hamprecht, "Joint supervised-unsupervised segmentation," *Proceedings of GCPR*, 2014.

[4] A. Vazquez-Reina, M. Gelbart, D. Huang, J. Lichtman, E. Miller, and H. Pfister, "Segmentation fusion for connectomics," *Proceedings of ICCV*, 2011.

[5] M. Seyedhosseini, M. Sajjadi, and T. Tasdizen, "Image Segmentation with Cascaded Hierarchical Models and Logistic Disjunctive Normal Networks," *Proceedings of ICCV*, 2013.

[6] J. Funke, J. Martel, S. Gerhard, B. Andres, et al., "Candidate Sampling for Neuron Reconstruction from Anisotropic Electron Microscopy Volumes," in *Proceedings of MICCAI*, 2014.

[7] Ting Liu, Cory Jones, Mojtaba Seyedhosseini, and Tolga Tasdizen, "A modular hierarchical approach to 3d electron microscopy image segmentation," *Journal of Neuroscience Methods*, vol. 226, no. 0, pp. 88 – 102, 2014.

[8] V. Jain, S.C. Turaga, K.L. Briggman, M.N. Helmstaedter, et al., "Learning to Agglomerate Superpixel Hierarchies," in *Proceedings of NIPS*, 2011.

[9] B. Andres, T. Kroeger, K.L. Briggman, W. Denk, et al., "Globally Optimal Closed-surface Segmentation for Connectomics," *Proceedings of ECCV*, 2012.

[10] J. Nunez-Iglesias, R. Kennedy, T. Parag, J. Shi, and D.B. Chklovskii, "Machine Learning of Hierarchical Clustering to Segment 2D and 3D Images.," *PloS one*, 2013.

[11] G.B. Huang and V. Jain, "Deep and Wide Multiscale Recursive Networks for Robust Image Labeling," *Proceedings of ICLR*, 2014.

[12] C. Sommer, C.N. Straehle, U. Koethe, and F.A. Hamprecht, "Ilastik: Interactive learning and segmentation toolkit," *Proceedings of ISBI*, 2011.

[13] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," *Proceedings of CVPR*, 2005.

[14] A. Kreshuk, C.N. Straehle, C. Sommer, U. Koethe, et al., "Automated detection and segmentation of synaptic contacts in nearly isotropic serial electron microscopy images.," *PloS ONE*, 2011.

[15] C. Becker, K. Ali, G. Knott, and P. Fua, "Learning Context Cues for Synapse Segmentation," *IEEE Transactions on Medical Imaging*, 2013.

[16] J. Kappes, M. Speth, G. Reinelt, and C. Schnoerr, "Higher-order Segmentation via Multicuts," *ArXiv e-prints*, , no. http://arxiv.org/abs/1305.6387, 2013.

[17] C.N. Straehle, U. Koethe, G.W. Knott, and F.A. Hamprecht, "Carving: scalable interactive segmentation of neural volume electron microscopy images.," *Proceedings of MICCAI*, 2011.