TRACKING OF MICROTUBULES IN ANISOTROPIC VOLUMES OF NEURAL TISSUE

Julia M. Buhmann[⋆] *Stephan Gerhard*[⋆†] *Matthew Cook*[⋆] *Jan Funke*[⋆‡]

[⋆] Institute of Neuroinformatics
University of Zurich/
ETH Zurich

[†] Janelia Research Campus
VA, Ashburn

[‡] Institut de Robòtica i
Informàtica Industrial
UPC Barcelona

ABSTRACT

For both the automatic and manual reconstruction of neural circuits from electron microscopy (EM) images, the detection and identification of intracellular structures provide useful cues. This is particularly true for microtubules which are indicative of the scaffold of neuronal morphology. However, to our knowledge, the automated reconstruction of microtubules from EM images of neural tissue has received no attention so far.

In this paper, we present an automatic method for the tracking of microtubules in 3D EM volumes of neural tissue. We formulate an energy-based model on short candidate segments of microtubules found by a local classifier. We enumerate and score possible links between candidates, in order to find a cost-minimal subset of candidates and links by solving an integer linear program. The model provides a way to incorporate biological priors including both hard constraints (e.g. microtubules are topologically chains of links) and soft constraints (e.g. high curvature is unlikely). We test our method on a challenging EM dataset of *Drosophila* neural tissue and show that our model reliably tracks microtubules spanning many image sections.

1. INTRODUCTION

A current bottleneck of Connectomics, i.e. the reconstruction and study of the connectivity of neurons, is the extraction of neural circuits from large volumes of raw image data. Only electron microscopy (EM) methods provide sufficient resolution to resolve entire neural circuits at synaptic level. The size of resulting datasets can easily reach several hundred terabytes, even for small brains such that of *Drosophila* [1]. The manual labeling of individual neurons and their synaptic connections is tedious and time consuming, and can practically only be done for selected neurons in small circuits.

Currently, the accuracy of state of the art image processing methods for the automation of this process is insufficient to replace manual reconstruction efforts. Common approaches of neuron segmentation include neuron membrane detection and the clustering of pixels and superpixels into larger neuron segments [2]. A promising way to improve the accuracy is to design models that account for more cell components than just neuron membranes [3]. Consequently, some effort has been devoted to the automated segmentation of additional cell components found in EM recordings, such as synapses [4, 5], mitochondria [6], and myelin [7].

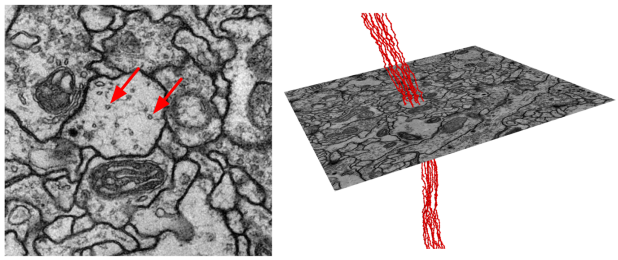

(a) (b)

Fig. 1: (a) Appearance of perpendicularly cut microtubules in an EM 2D section (red arrows). (b) 3D rendering of ground truth tracings comprising 12 microtubules. As can be seen here, microtubules in neurons run in parallel along the neuron's trajectory.

In this paper, we propose a method for the automated reconstruction of microtubules in neurons. The reconstruction of microtubules, which form part of the neuron's cytoskeleton, can aid analysis of neural circuits in two ways: First, microtubules are a valuable source of constraints for the automated reconstruction of neurons [8, 9]. Second, microtubules follow the coarse morphology of a neuron, its so-called *backbone*. The backbone allows to quickly identify neuron types or even individual neurons [8], without having to segment the whole neural arbor. Additionally, reconstructing the backbone morphology is useful to find starting points for the manual reconstruction of selected circuits, a task that is currently done by hand [10].

So far, to our knowledge, there is no work considering the automated tracking of microtubules in EM volumes of neural tissue. Methods have been developed to trace [11] and stitch [12] microtubules in serial EM tomograms. However, [12] operates on human proofread tracings and only allows the pairwise matching between two consecutive sections. Furthermore, the used EM tomograms differ significantly from EM images of neural tissue.

For the tracking of microtubules in neural tissue, we de-

978-1-4799-2349-6/16/$31.00 ©2016 IEEE

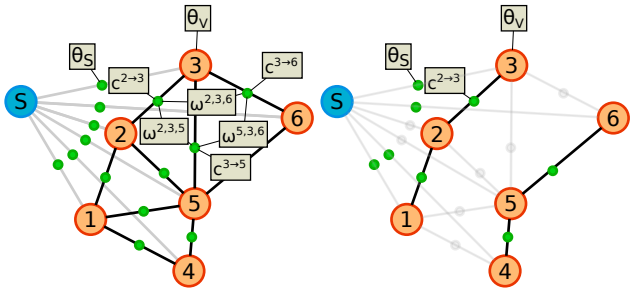

(a) example graph (b) consistent solution

Fig. 2: Illustration of the graphical model used for inference. Candidate segments of microtubules are shown as orange nodes, possible links between them as thick lines. In (a), all cost contributions involving candidate 3 are shown (unary potentials for selecting a candidate or a link and pairwise potentials for the joint selection of two links). One consistent solution is shown in (b).

signed an energy-based model on candidate detections, a method that has already been successfully applied to biomedical image processing problems, such as the reconstruction of neurons from EM sections [9, 13] or the tracking of cells in time-lapse videos [14]. Possible links between candidates, which in our case are small microtubule segments, are enumerated and scored. The scoring of links is a suitable mechanism to include biological priors: In our model, strong curvature is penalized because microtubules are physically restricted of how strong they can bend. With hard constraints we ensure that selected links form a chain. The resulting optimization problem can be solved to global optimality using an integer linear programming (ILP) solver.

2. METHOD

In this section, we describe our candidate based assignment model for the tracking of 3D trajectories of microtubules. First, we extract possible candidates for microtubule segments based on a pixel classification. Between these candidates, we enumerate possible links within a threshold distance. We introduce binary indicator variables for the selection of each candidate, each possible link, and each pair of links adjacent to the same candidate. Based on priors and orientation features, we compute a cost for each indicator. We introduce hard constraints to ensure that each candidate is linked to at most two other candidates. Finally, we find a globally cost-minimal and consistent assignment of the indicators by solving an ILP.

Extraction of Microtubule Candidates. We use the toolkit ilastik [15] to train a random-forest pixel classifier to predict microtubules (interactive training time: 2 hours). The classification output consists of a pixelwise probability map, high intensity indicating a high probability for the pixel to belong to a microtubule. In order to extract individual objects, we perform a connected-component analysis on the thresholded probability map. The centroids of the extracted objects are taken as the microtubule candidates.

Indicator Variables We enumerate possible links between microtubule candidates within a threshold distance and assign each a binary decision variable. We can now picture the task of tracking microtubules as a task of identifying and selecting the right links.

More formally, we define an undirected graph $G = (V \cup \{S\}, E)$, in which the vertices V correspond to microtubule candidates, and the edges $E \subset V \times V$ to possible links. For each candidate $i \in V$, we introduce a binary variable $v^i \in \{0, 1\}$, indicating the selection ($v^i = 1$) or rejection ($v^i = 0$) of this candidate. Analogously, we introduce binary variables $e^{i \to j} \in \{0, 1\}$ to indicate the selection of link $(i, j) \in E$. To explicitly represent the beginning or end of a microtubule trajectory, we add a special sink node S, and introduce additional link indicators $e^{i \to S}$ for each candidate $i \in V$. An example graph is shown in Fig. 2.

Local evidence is used for estimating the cost for each link to be selected. The cost should ideally reflect the validity of a link, e.g, a correct link should have a low cost assigned. A cost vector $\boldsymbol{c} \in \mathbb{R}$, associated with the link indicators $\boldsymbol{e}$, is constructed and costs are calculated as follows:

$$c^{i \to j} = \theta_D d_{ij} + \theta_A(\phi_{(i,j),i} + \phi_{(i,j),j}) \quad (1)$$

$$c^{i \to S} = \theta_S, \quad (2)$$

where θ_D and θ_A are parameters of the model, d_{ij} denotes the Euclidean distance between candidates i and j, and $\phi_{(i,j),i}$ is the spanning angle between the direction of link (i, j) (measured between the centroids of the candidates) and a locally estimated direction vector for the candidate i. In other words, long distance and high deviation from the estimated direction vector is penalized in the model. θ_S is the cost for beginning or end of a trajectory and the same for all candidates. Analogously, the selection of a candidate, encoded by the binary vector $\boldsymbol{v}$ is associated with a cost vector p, where we set the cost for selecting a candidate i uniformly to $p_i = \theta_V$.

We introduce additional costs for the selection of pairs of incident links. Our motivation is here that microtubules are rigid structures: link combinations whose spanning angle is straight are thus much more likely to represent a true microtubule trajectory. Let (i, j) and (j, k) be two incident links on i and $\phi_{i,j,k}$ the spanning angle between the direction vectors of (i, j) and (j, k). The cost for selecting two incident links is then given as

$$\omega^{i,j,k} = (\theta_C \phi_{i,j,k})^2. \quad (3)$$

To account for these costs, we construct binary indicators $a^{i,j,k}$ for each pair of links (i, j) and (j, k). The following constraints ensure that indicator $a^{i,j,k} = 1$, iff $e^{i,j} = 1$ and $e^{j,k} = 1$:

$$2a^{i,j,k} - e^{i,j} - e^{j,k} \leq 0 \text{ and } e^{i,j} + e^{j,k} - a^{i,j,k} \leq 1 \quad (4)$$

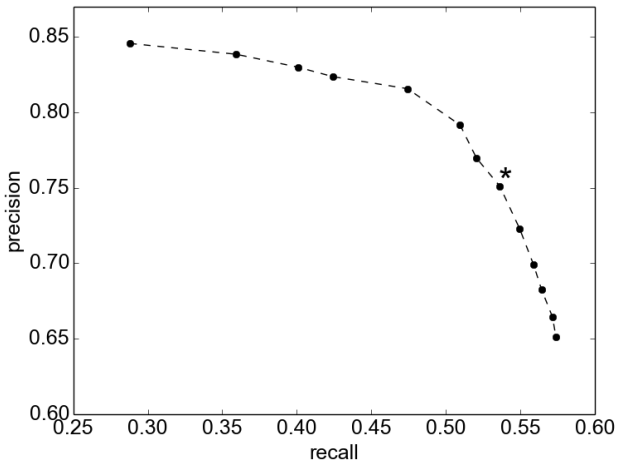

Fig. 3: Performance measure for different values of θ_S. One data point corresponds to one value of θ_S. The point with the best f-measure is marked with an asterisk (f-measure = 0.63, precision = 0.75, recall = 0.54).

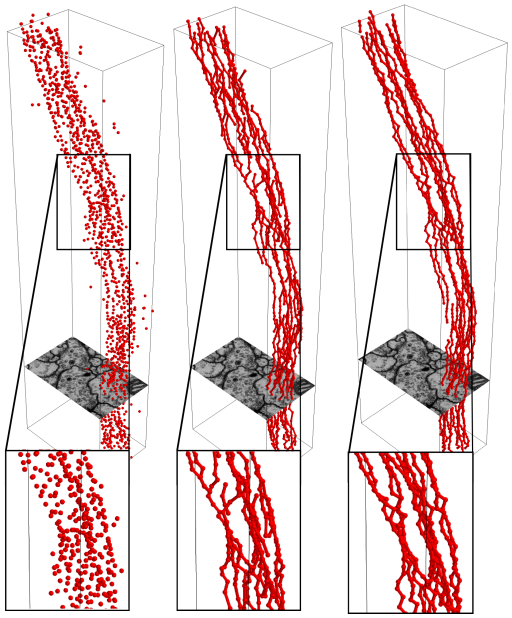

(a) candidates (b) results (c) ground truth

Fig. 4: 3D rendering of (a) initial microtubule candidates, (b) automatically reconstructed trajectories and (c) the corresponding ground truth used for testing. In (a) and (b), only microtubules are displayed that are located within the ground truth neuron.

Consistency Constraints Unlike a neuron, a microtubule does not branch. This means that if a candidate is selected, exactly two of its links should also be selected (one of them might be the link to sink S). This is guaranteed if the following equality holds:

$$2v^i - \sum_{e \in \boldsymbol{e}^{\rightarrow i}} e = 0 \qquad \forall i \in V, \tag{5}$$

where $\boldsymbol{e}^{\rightarrow i}$ denotes the set of variables representing all incident links on candidate i. We further have to guarantee, that, if a link is selected, the two incident candidates on this link also need to be selected:

$$2e^{i \rightarrow j} - v^i - v^j \leq 0 \qquad \forall (i,j) \in E. \tag{6}$$

A solution for the optimization problem can be found solving the following ILP:

$$\underset{\substack{e,a,v \\ \text{subject to 4, 5 and 6}}}{\arg\min} \quad E(e,a,v) = c^T e + \omega^T a + p^T v \tag{7}$$

Our implementation is publicly available at `https://github.com/juliabuhmann/microtubules`.

3. RESULTS

We evaluated our method on a serial-section transmission electron microscopy (ssTEM) volume of *Drosophila* neural tissue [16] ($5 \times 5 \times 50$ nm resolution). Due to the anisotropy of this volume, we consider each 2D section of the stack individually for candidate extraction. Furthermore, we focus on candidates for microtubules that roughly run perpendicular to the cutting plane, although our model allows candidates in general positions.

We measure performance in terms of correctly found links compared to human generated ground-truth, acquired by linking microtubule points across sections. For that, we first match each selected candidate to the closest ground-truth point within 25 nm using Hungarian Matching. We consider a selected link and a ground-truth link as a match, if their points were matched to each other. Unmatched selected links are considered as False Positives (FP), unmatched ground-truth links as False Negatives (FN). Candidates that are outside the volumetric ground truth of test neurons are excluded from analysis.

We obtain the parameters θ_S, θ_V, θ_A, θ_D, and θ_C of our model by performing a grid search on a validation neuron (1382 links, spanning 100 sections) to maximize the f-measure. The f-measure combines precision and recall into a single value: $F = 2\frac{\text{precision} * \text{recall}}{\text{precision} + \text{recall}}$. On three test neurons (in total 3006 links, spanning 100 sections each), we achieve a final f-measure of 0.63. Fig. 3 shows results in terms of precision vs. recall for various values of θ_S, the prior for the number of microtubule trajectories. Depending on the application, it may be preferable to focus on either a high precision or a high recall. A volume of 250x250x250 voxels (35600 links) has an ILP running time of 7.5 min using the Gurobi solver.

4. DISCUSSION

To our knowledge, we presented the first method for the reconstruction of microtubules in EM volumes of neural tissue. We showed that our model reliably tracks microtubules in a difficult to analyze, highly anisotropic dataset. Our model conveniently incorporates biological priors in form of hard constraints (microtubules form chains) and soft constraints (microtubules have small curvature).

While our model scores higher maximal precision (0.84) the maximal reached recall is comparatively low (0.57). Although the model can cope with False Negatives to a certain degree, performance of the model is ultimately limited by the quality of the initial candidates. The here proposed candidate extraction method is based on a random-forest classifier. However, convolutional neural networks currently outperform other classification algorithms in image recognition tasks. Using a convolutional neural network for candidate extraction is thus likely to improve overall performance results.

Even with a low recall, given the high density of microtubules in neural tissue, the selected trajectories are already potentially sufficient to provide helpful cues for neural reconstruction. As can be seen in Fig. 4b, our automatically tracked microtubules provide continuous evidence for the underlying neuron.

Acknowledgements

We thank the Albert Cardona Lab at the Howard Hughes Medical Institute Janelia Research Campus for providing the EM data.

5. REFERENCES

[1] S. M. Plaza, L. K. Scheffer, and D. B. Chklovskii, "Toward large-scale connectome reconstructions," *Current opinion in neurobiology*, 2014.

[2] J. Nunez-Iglesias, R. Kennedy, T. Parag, J. Shi, D. B. Chklovskii, and X.-N. Zuo, "Machine learning of hierarchical clustering to segment 2d and 3d images," *PloS one*, 2013.

[3] N. Krasowski, T. Beier, G. Knott, U. Koethe, F. Hamprecht, and A. Kreshuk, "Improving 3d em data segmentation by joint optimization over boundary evidence and biological priors," *Proceedings of ISBI*, 2015.

[4] C. Becker, K. Ali, G. Knott, and P. Fua, "Learning context cues for synapse segmentation," *IEEE Transactions on Medical Imaging*, 2013.

[5] A. Kreshuk, C. N. Straehle, C. Sommer, U. Koethe, G. Knott, and F. A. Hamprecht, "Automated segmentation of synapses in 3d em data," *Proceedings of ISBI*, 2011.

[6] A. Lucchi, K. Smith, R. Achanta, G. Knott, and P. Fua, "Supervoxel-based segmentation of mitochondria in em image stacks with learned shape features," *IEEE Transactions on Medical Imaging*, 2012.

[7] A. Kreshuk, R. Walecki, U. Koethe, M. Gierthmuehlen, D. Plachta, C. Genoud, K. Haastert-Talini, and F. Hamprecht, "Automated tracing of myelinated axons and detection of the nodes of ranvier in serial images of peripheral nerves," *Journal of microscopy*, 2015.

[8] C. M. Schneider-Mizell, S. Gerhard, M. Longair, T. Kazimiers, F. Li, M. F. Zwart, A. Champion, F. Midgley, R. Fetter, S. Saalfeld, *et al.*, "Quantitative neuroanatomy for connectomics in drosophila," *bioRxiv*, 2015.

[9] J. Funke, B. Andres, F. Hamprecht, A. Cardona, M. Cook, *et al.*, "Efficient automatic 3d-reconstruction of branching neurons from em data," *Proceedings of CVPR*, 2012.

[10] T. Ohyama, C. M. Schneider-Mizell, R. D. Fetter, J. V. Aleman, R. Franconville, M. Rivera-Alba, B. D. Mensh, K. M. Branson, J. H. Simpson, J. W. Truman, *et al.*, "A multilevel multimodal circuit enhances action selection in drosophila," *Nature*, 2015.

[11] B. Weber, G. Greenan, S. Prohaska, D. Baum, H.-C. Hege, T. Müller-Reichert, A. A. Hyman, and J.-M. Verbavatz, "Automated tracing of microtubules in electron tomograms of plastic embedded samples of caenorhabditis elegans embryos," *Journal of Structural Biology*, 2012.

[12] B. Weber, E. M. Tranfield, J. L. Höög, D. Baum, C. Antony, T. Hyman, J.-M. Verbavatz, and S. Prohaska, "Automated stitching of microtubule centerlines across serial electron tomograms," *PloS ONE*, 2014.

[13] V. Kaynig, T. J. Fuchs, and J. M. Buhmann, "Geometrical consistent 3d tracing of neuronal processes in sstem data," *Proceedings of MICCAI*, 2010.

[14] M. Schiegg, P. Hanslovsky, B. X. Kausler, L. Hufnagel, F. Hamprecht, *et al.*, "Conservation tracking," *Proceedings of ICCV*, 2013.

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

[16] S. Gerhard, C. Schneider-Mizell, I. Andrade, R. Fetter, and A. Cardona, "Comparative connectomics of drosophila nociceptive circuits." In preparation.