Geometrical Consistent 3D Tracing of Neuronal Processes in ssTEM Data

Verena Kaynig[1,2], Thomas J. Fuchs[1], and Joachim M. Buhmann[1]

[1] Department of Computer Science, ETH Zurich, Switzerland
verena.kaynig@inf.ethz.ch
[2] Electron Microscopy ETH Zurich (EMEZ), Switzerland

Abstract. In neuroanatomy, automatic geometry extraction of neurons from electron microscopy images is becoming one of the main limiting factors in getting new insights into the functional structure of the brain. We propose a novel framework for tracing neuronal processes over serial sections for 3d reconstructions. The automatic processing pipeline combines the probabilistic output of a random forest classifier with geometrical consistency constraints which take the geometry of whole sections into account. Our experiments demonstrate significant improvement over grouping by Euclidean distance, reducing the split and merge error per object by a factor of two.

1 Introduction

Neuroanatomists build 3d reconstructions of neuronal structures and their synaptic connections in order to gain insight into the functional structure of the brain. As the identification of post synaptic densities is crucial for this task, serial section electron microscopy is the only imaging technique so far, which can provide sufficient resolution. New advances in sample preparation and the imaging process make the acquisition of large data volumes possible [1,2], but the image processing work flow needed to evaluate these data sets still relies heavily on manual labor [3]. This manual intervention renders the process not only error prone and very tedious for the neuroanatomist, but nowadays also becomes a serious bottleneck for the evaluation work flow. In order to build 3d reconstructions of neuronal tissue based on transmission electron microscope (TEM) images, the sample first is embedded into resin, then cut into ultra thin sections of about 50 nm thickness and finally each section is recorded with the microscope. The following image processing work flow consists of aligning the image stack, segmenting structures of interest and building 3d reconstructions out of these segmentations. While the image resolution is dependent on the microscope and can easily achieve 4 nm per pixel, the z-resolution is limited by the section thickness around 40 nm. As a consequence segmentation is usually performed in two dimensions, using the fine resolution to identify membranes of neuronal processes like dendrites and axons [4,5,6,7]. The regions surrounded by the detected membranes then need to be grouped over consecutive sections to extract the geometry of neuronal processes (see Figure 1). Previous work

T. Jiang et al. (Eds.): MICCAI 2010, Part II, LNCS 6362, pp. 209–216, 2010.

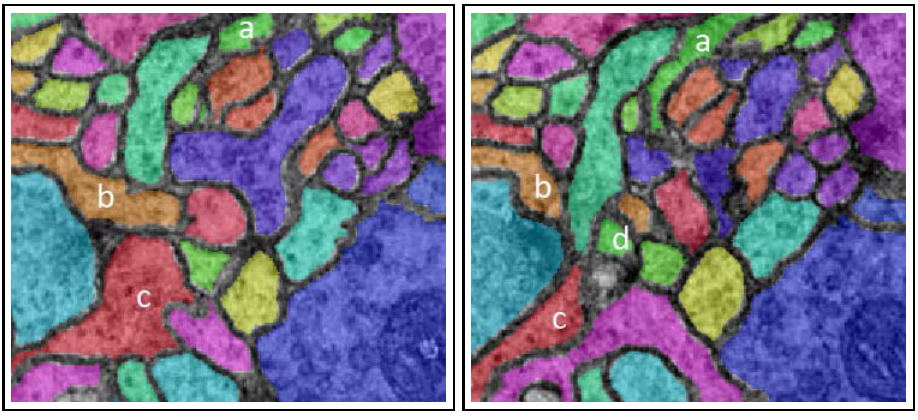

Fig. 1. Example groupings of regions from two adjacent section images (correspondence is indicated by color). The grouping problem is especially hard for thin processes, which have greater flexibility than large structures. In addition structures running longitudinal to the cutting plane, express significant changes in appearance between sections (a-c). Example d has no correspondence in the left section.

has addressed this problem by tracking single processes through the image stack [8,9]. We extend the previous approaches with respect to three important points: (i) instead of tracking single processes the labeling of the whole data volume is optimized, allowing for neuronal processes to start or end inside the volume, (ii) similarity of regions is learned from annotated data, (iii) geometrical consistency between whole sections is taken into account.

2 Method

We regard the problem of three dimensional geometry extraction as partitioning an edge weighted graph into connected components representing an image volume belonging to the same neuronal process. The regions are represented by the vertices V of the graph and the set of edges E connects each region to all regions of the two adjacent sections. Each edge is assigned a weight w_{ij} according to the similarity between regions i and j.

We propose the following processing pipeline to build the edge weight matrix W and to find connected components representing neuronal processes (see Figure 2). First, a set of weight matrices based on features like region overlap or similarity of texture is created. A detailed description of the features is given in Section 2.1. A random forest classifier [10] is trained on manual annotations to predict the similarity between two regions. The weight matrix predicted by the random forest classifier only captures the similarity of pairwise regions. Therefore, a further step refines the weight matrix using geometrical consistent constraints that take the geometry of all neuronal processes included in the section into account. Finally, agglomerative clustering is employed to partition the graph into connected components representing neuronal processes. The hierarchical clustering

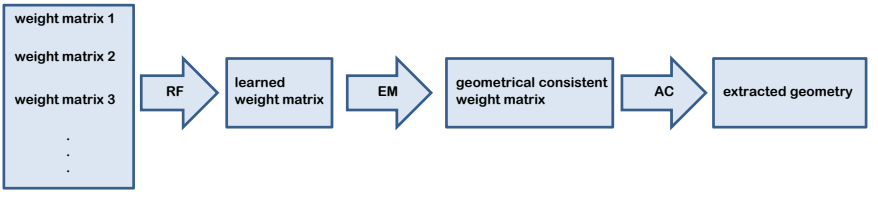

Fig. 2. Processing pipeline for the extraction of 3d geometry of neuronal processes as proposed in this paper. First, a similarity matrix of pairwise regions is learned by a random forest classifier. The learned weight matrix is then combined with geometrical constraints, taking the geometry of all neuronal processes from the whole section into account. Optimization is performed by expectation maximization. Finally agglomerative clustering is used to extract continuous neuronal processes.

scheme starts from individual objects and then progressively merges the regions which are most similar to each other. This system mirrors the approach of the neuroanatomist, who first establishes correspondences between regions that are easy to detect and then refines the partitioning.

2.1 Similarity Features between Regions

The following paragraphs describe the features that are used to train the random forest classifier from manually annotated data. For each feature we build a weight matrix W, each entry representing the edge weight of the corresponding edge in the graph.

Euclidean distance of region center: Each region i is represented by its center of mass $c_i \in \mathbb{R}^3$. The distance of two regions is then given by the Euclidean distance between the two centers:

$$W_{distance}(i,j) = \sqrt{(c_i - c_j) \cdot (c_i - c_j)^T} \tag{1}$$

Overlap of region areas: For each region i, the set P_i contains the position of all pixels belonging to the region ($P_i \in \mathbb{R}^3$). The overlap of two regions is measured by projecting both regions orthogonally to the same plane and building the intersection of both projections:

$$W_{overlap}(i,j) = \#(P_i \cdot A \cap P_j \cdot A), \text{ with } A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{2}$$

Difference in region size: Neuronal processes have only smooth variations in diameter. Therefore the size of corresponding regions should be similar to each other.

$$W_{size}(i,j) = \frac{(\#P_i - \#P_j)^2}{\#P_i + \#P_j} \tag{3}$$

Here $\#P_i$ describes the size of region i in number of pixels. The size difference between two regions is measured by the fraction of the difference in pixels in relation to the total size of both regions. This normalization accounts for the comparability of processes with large or small diameter.

Texture similarity: For the neuroanatomist, texture is an important clue for the extraction of neuronal processes. Intracellular structures like vesicles or microtubules provide information about the type of neuronal process, e.g. bouton or axon, and about the consistent grouping of regions. Following the approach described in [8], we measure the similarity in texture by the cross correlation coefficient of two regions

$$W_{xcorr}(i,j) = X_{\max}(r_i, r_j). \tag{4}$$

Where r_i represents the gray values of region i and $X_{\max}$ denotes the maximal cross correlation between the two regions.

Smooth continuation: This feature weights the connection between two regions i and j according to the smoothest continuation to the next sections. The smoothness of a possible continuation is given by the angle θ_{hij} between the three region centers c_h, c_i and c_j (see Figure 3).

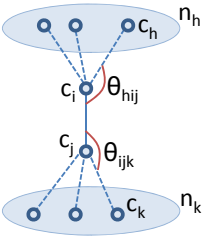

Fig. 3. Illustration of the smooth continuation feature. The smoothness of a possible continuation is given by the function $\theta(c_h, c_i, c_j)$ which measures the angle between the three region centers c_h, c_i and c_j.

$$W_{smooth}(i,j) = \frac{1}{2} \cdot (\min_{h \in n_h} \theta(c_h, c_i, c_j) + \min_{k \in n_k} \theta(c_i, c_j, c_k)), \tag{5}$$

with $\theta(c_h, c_i, c_j) = abs(\pi - \angle(c_h, c_i, c_j))$. The set n_h contains all regions from the section above region i and the set n_k contains all regions from the section below region j. Reflection is employed as border treatment to compute the smooth continuation feature for the first and last section of the stack.

2.2 Geometrical Consistency across Sections

Region correspondences should be assigned in consistency with the overall geometry changes from one section to the next. We address this problem by establishing geometrical consistency of the correspondences between sections. The

approach allows for a non linear but smooth transformation between sections to match correspondent points. Correspondences are not fixed beforehand, but obtained during the optimization [11] For the non-linear transformation we use an explicit polynomial kernel expansion off the points c_i: $\phi(c_i) = [1, c_{i1}, c_{i2}, c_{i1}^2, c_{i1}c_{i2}, c_{i2}^2, \ldots, c_{i2}^d]^T$.

The transformation matrix β projects these points back into the image plane, leading to a nonlinear transformation. Correspondences are assigned by a binary matrix M whose entry m_{ij} is one, if point c_i in one section corresponds to point c_j in the adjacent section and zero otherwise. The energy function to be optimized depends on the similarity of the correspondent regions as classified by the random forest, as well as on the quality of the geometric fit:

$$E(\beta, M) = \sum_{i=1}^{n_i} \sum_{j=1}^{n_j} -m_{ij} ||\phi(c_i)\beta - c_j||^2 + m_{ij} \cdot \ln(W(i,j)) \tag{6}$$

Here the index i runs over the number of regions n_i from one section and j over the number of regions n_j from the adjacent section. The variable m_{ij} contains the associated value of the assignment matrix M and $W(i,j)$ corresponds to the edge weight given by the random forest classifier. Maximizing this energy function can be interpreted as maximizing the data likelihood $p(C_i, C_j|\beta, M)$ where C_i and C_j are matrices containing all points from two adjacent sections. We use expectation maximization to optimize the joint log-posterior, treating the correspondences as unobservable. The algorithm iterates between estimating the expectation of the latent variables m_{ij} while keeping β fix and maximizing the joint log-posterior while keeping the expectation values of M constant.

E-step: In each iteration the variables m_{ij} are replaced by their conditional expectation given β. The expectation values are calculated using the currently optimized β. Under the condition that M is a valid assignment matrix ($\sum_j^{n_2} m_{ij} = 1$, for all $i = 1, \ldots n_2$), we derive the following result:

$$\gamma_{ij} = \mathbf{E}[m_{ij}|C_i, C_j, \beta] = \frac{p(C_i, C_j|\beta, m_{ij} = 1)}{\sum_{l=1}^{n_2} p(C_i, C_j|\beta, m_{il} = 1)} \tag{7}$$

M-step: The expectation of the joint log posterior has the same form as the joint log posterior itself, but with m_{ij} replaced by γ_{ij}. Under the assumption that β is smooth, i.e the components of β are assumed to be normally distributed, maximizing for β yields a weighted ridge regression problem with weights γ_{ij}:

$$\beta \leftarrow (\phi(\widetilde{\mathbf{C_i}})^T \mathbf{\Gamma} \phi(\widetilde{\mathbf{C_i}}) + 2\lambda \mathbf{I})^{-1} \phi(\widetilde{\mathbf{C_i}})^T \mathbf{\Gamma} \mathbf{C_j} \tag{8}$$

where Γ is a $(n_i \cdot n_j) \times (n_i \cdot n_j)$-dimensional diagonal matrix of the weights γ_{ij}. The $(n_i \cdot n_j) \times 2$ matrix $\widetilde{\mathbf{C_i}}$ contains n_j copies of each center point c_i from the first section and the $(n_i \cdot n_j) \times 2$ matrix $\mathbf{C_j}$ contains n_j possible correspondence points from the adjacent section for each point c_i. The parameter λ is the regularization parameter defined by the prior distribution $p(\beta)$. In our experiments λ is set to 0.001.

3 Evaluation

The proposed method is evaluated on ssTEM images, resembling average image quality from neuroanatomy projects. The data set depicts part of the dorsolateral fasciclin-II tract of the ventral nerve cord of the first instar larva of drosophila, at abdominal segment 5. It consists of 30 images with 512x512 pixels. The resolution is 3.7 nm per pixel in the image plane and section thickness is 50nm. The whole data set was annotated exclusively by a neuroanatomist, providing the ground truth for the evaluation. The random forest classifier was trained on this data set using ten fold cross validation to obtain the test error. The remaining pipeline is free of tuning parameters and therefore just applied to the test results of the classifier.

As demonstrated by the plots in Figure 4, each step of our processing pipeline yields significant improvement for the geometry extraction in terms of split and merge error per object. A perfect solution would assign exactly one label per ground truth cluster. For each additional label a split error is counted. A merge error occurs when two ground truth clusters are assigned the same label. If two ground truth clusters are merged more than once, we follow the definition of [12] and count this as one error as the same two objects are involved.

The agglomerative clustering is restricted to establish a maximum of two correspondences for each region, one to the upper and one to the lower section.

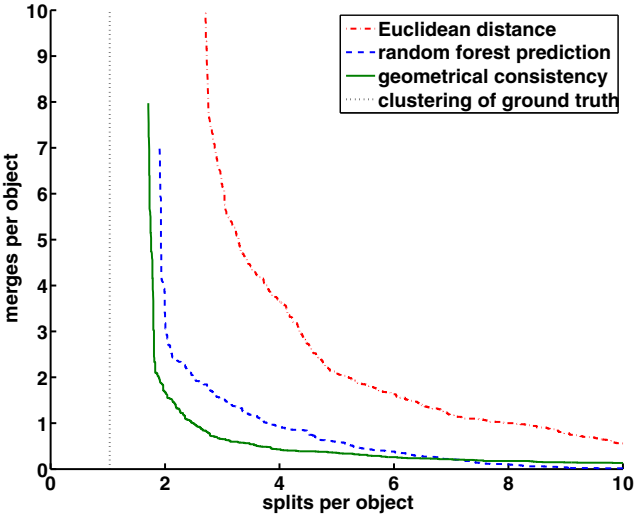

Fig. 4. Evaluation of clustering results according to split/merge error per neuronal process. Depicted are the results for different weight matrices: (i) Euclidean distance of region centers only, (ii) weights learned by the random forest classifier, and geometrical consistent weights. The dotted line corresponds to the best result obtainable without considering branching processes.

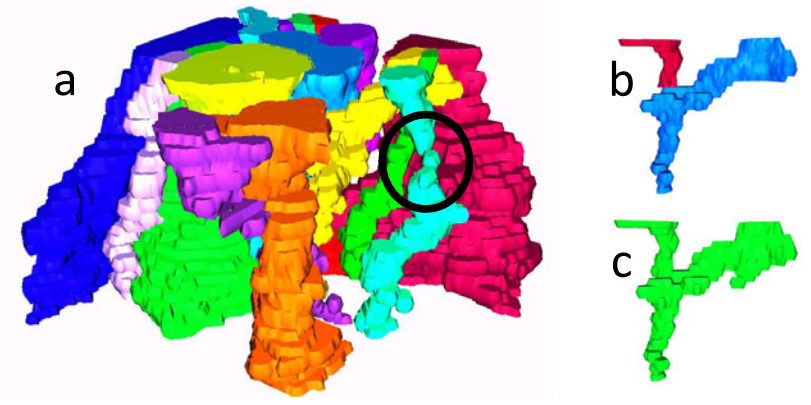

Fig. 5. 3d reconstruction of neuronal processes that were correctly tracked over all 30 sections. A black circle marks an example where regions were correctly grouped despite not having any overlap in adjacent sections. The neuronal process shown in Figure b (ground truth) and c (clustering result) shows an example for a split. The large part including regions moving longitudinal to the cutting direction was correctly grouped and the remaining part was also identified as one object.

Thus, our model allows for starting and ending of new neuronal processes inside the volume, but does not account for branching of processes. The dotted line in Figure 4 marks the best clustering performance achievable by this model.

Examples of extracted geometries are given in Figure 5. The examples demonstrate, that the proposed method is capable of extracting correct geometries also in difficult cases of neuronal processes running longitudinal to the cutting plane and in cases of discontinuities in the geometry due to alignment errors.

4 Conclusion

In this paper we introduced a novel framework for global tracing of neuronal processes in stacks of serial section transmission electron microscopy images. The setting is formulated as a partitioning problem on edge weighted region-graphs. The main contributions of this work are threefold: (i) On the modeling side we propose the use of a random forest classifier to learn a predictor for neighborhood relations of regions within the 3d volume. (ii) Predicted region correspondences are refined taking the geometrical consistency of whole sections into account. (iii) The unsupervised clustering approach results in a non parametric robust procedure for partitioning the graph. In depth evaluation of all single steps of the pipeline and cross validation of the similarity classification demonstrate significant improvement in terms of split and merge error per object. We are convinced that the proposed algorithm is a valuable contribution to the field of neuroscience due to its robustness and general applicability for neuronal process tracing in 3d volumes.

Acknowledgement

We like to thank Nuno Maçarico da Costa, Kevan Martin, and Albert Cardona, Institute of Neuroinformatics UNI-ETH Zurich, for valuable discussions on neuroanatomy and for providing the TEM images.

References

1. Knott, G., Marchman, H., Wall, D., Lich, B.: Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. J. Neurosci. (2008)
2. Denk, W., Horstmann, H.: Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. Curr. Opin. Neurobiol. (2006)
3. Hoffpauir, B., Pope, B., Spirou, G.: Serial sectioning and electron microscopy of large tissue volumes for 3D analysis and reconstruction: a case study of the calyx of held. Nat. Protoc. 2(1), 9–22 (2007)
4. Kaynig, V., Fuchs, T., Buhmann, J.M.: Neuron geometry extraction by perceptual grouping in sstem images. In: CVPR 2010 (2010)
5. Mishchenko, Y.: Automation of 3d reconstruction of neural tissue from large volume of conventional serial section transmission electron micrographs. J. Neurosci. Methods 176, 276–289 (2009)
6. Reina, A.V., Miller, E., Pfister, H.: Multiphase geometric couplings for the segmentation of neural processes. In: CVPR (June 2009)
7. Kannan, U., Paiva, A., Jurrus, E., Tasdizen, T.: Automatic markup of neural cell membranes using boosted decision stumps. In: ISBI (July 2009)
8. Jurrus, E., Whitaker, R., Jones, B.W., Marc, R., Tasdizen, T.: An optimal-path approach for neural circuit reconstruction. In: ISBI, pp. 1609–1612 (2008)
9. Jurrus, E., Tasdizen, T., Koshevoy, P., Fletcher, P.T., Hardy, M., Chien, C., Denk, W., Whitaker, R.: Axon tracking in serial block-face scanning electron microscopy. Med. Image Anal. 13(1), 180–188 (2009)
10. Breiman, L.: Random forests. Mach. Learn. 45(1), 5–32 (2001)
11. Chui, H., Rangarajan, A.: A new algorithm for non-rigid point matching. In: CVPR, vol. 2, pp. 44–51 (2000)
12. Turaga, S.C., Murray, J.F., Jain, V., Roth, F., Helmstaedter, M., Briggman, K., Denk, W., Seung, H.S.: Convolutional networks can learn to generate affinity graphs for image segmentation. Neural Comput. 22(2), 511–538 (2010)