

Neuron Geometry Extraction by Perceptual Grouping in ssTEM Images

Verena Kaynig^{1,2}, Thomas Fuchs¹, Joachim M. Buhmann¹

{verena.kaynig, thomas.fuchs, jbuhmann}@inf.ethz.ch

¹Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland

²Electron Microscopy
ETH Zurich
8093 Zurich, Switzerland

Abstract

In the field of neuroanatomy, automatic segmentation of 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 the segmentation of thin elongated structures like membranes in a neuroanatomy setting. The probability output of a random forest classifier is used in a regular cost function, which enforces gap completion via perceptual grouping constraints. The global solution is efficiently found by graph cut optimization. We demonstrate substantial qualitative and quantitative improvement over state-of-the-art segmentations on two considerably different stacks of ssTEM images as well as in segmentations of streets in satellite imagery. We demonstrate that the superior performance of our method yields fully automatic 3D reconstructions of dendrites from ssTEM data.

1. Introduction

Neuroanatomists build 3D reconstructions of neuronal structures and their synaptic connections in order to gain insight in the functional structure of the brain. The identification of post synaptic densities is crucial for this task, and currently electron microscopy is the only imaging technique which can provide sufficient resolution. Recent advances in sample preparation and the imaging process make the acquisition of large data volumes possible [15, 10, 8]. In contrast to these technological advances for acquisition, the subsequent imaging work flow required for analyzing these data still relies heavily on manual labor [11]. Not only does this very tedious manual intervention by the neuroanatomist make the process susceptible to errors, it is also the major bottleneck for automatic evaluation and reconstruction.

To build 3D reconstructions of neuronal tissue based on transmission electron microscope (TEM) images, the sample is first embedded into resin, which is subsequently cut

into ultra thin sections of about 50 nm thickness. Each section is then recorded with the TEM. The image processing work flow that follows consists of (i) registering the image stack, (ii) segmenting structures of interest, and (iii) building 3D reconstructions out of these segmentations. As dendrites and axons are surrounded by membranes, a perfect membrane segmentation allows for a full reconstruction of the data volume.

Significant progress has been made in recent years on the front for the registration and warping of serial section TEM (ssTEM) images from single sections into a single image volume [4, 14, 18]. However, the automatic segmentation of ssTEM data is still an unsolved problem. The images typically show highly textured dense biological structures, which renders the detection of membranes difficult. Furthermore, variations such as different animal species, sample preparation, staining protocols e.t.c., can lead to very different image characteristics (see Figure 4) which poses an additional challenge for the automatic segmentation.

In order to cope with the data annotation problem, semi-automatic tools have been developed to speed up manual annotation [19, 20, 24, 21]. Recent works in automatic reconstruction use intracellular staining procedures to simplify the segmentation task [12, 3, 22]. This approach sacrifices important anatomical details like dendrite and bouton textures which are necessary to identify synapses and to build neural circuit reconstructions. In [13], the authors developed a method for neuronal circuit 3D reconstructions from ssTEM images. The authors tackle the membrane segmentation problem by thresholding on linear diffusion filtering, but this approach is only applicable to unusually high data quality.

For thin and elongated structures like membranes, graph cut is well known to have problems with “shrinking bias”. Current state of the art segmentation methods overcome this problem by combining overall smoothness with gradient flux, to enhance the segmentation result [23, 5]. In [24] gradient flux is additionally used to segment the in-

terior region of a dendrite. But, in images with textured background, like electron microscopy images, gradient flux leads to false positive detections, due to the high gradient in the background.

In this paper, we propose a novel energy term to overcome the shortcomings of gradient flux for the automatic segmentation of membranes. We improve the segmentation of thin elongated structures by enhancing gap completion. The energy term is regular and thus can be efficiently globally optimized using max-flow/min-cut computation. The novel energy framework combines a discriminative model for membrane appearance learned by a random forest classifier with perceptual grouping constraints for contour completion in a single energy minimization framework. The gap completion term follows the principle of good continuation, which states that elongated structures, which form a continued visual line should be grouped together. Thus, the proposed energy term focuses on the main characteristics of membranes as thin elongated structures, which are biologically given and therefore not influenced by different sample preparations. We also take information of adjacent sections into account to support the segmentation of membranes which are not prominent in one image, but better detectable in corresponding regions of nearby sections. This corresponds to the principle of non accidentalness, which states that elements should be grouped, if their configuration is unlikely to occur by chance.

The framework is evaluated on two different data sets of conventional ssTEM images from neuroanatomy. The image stacks differ not only in the type of animal brain shown (mammal and insect), but also in the staining protocols used, leading to very different image characteristics. On both data sets, the proposed cost function with perceptual grouping constraints outperforms the state-of-the-art segmentation using gradient flux. These results point out the robustness of the proposed perceptual grouping constraints to different staining protocols and animal types. The high quality of the membrane segmentations allows for fully automatic 3D reconstructions of neuronal structures. To demonstrate the wide applicability of the proposed framework we also provide segmentation results for streets from satellite imagery.

2. Perceptual grouping constraints via graph cut

In the graph cut framework each pixel p is mapped to corresponding labels $y_p \in \{0, 1\}$ such that the entire labeling y for all pixels minimizes a given energy function $E(y)$. Typically the energy function $E(y)$ consists of a summation over the data term $E_d(y_p)$ and a smoothness

term $E_s(y_p, y_q)$ over neighboring pixels:

$$E(y) = \sum_{p \in P} E_d(y_p) + \lambda \sum_{p \in P, q \in N_2(p)} E_s(y_p, y_q), \quad (1)$$

where P denotes the set of all pixels and $N_2(p)$ the set of all pixels adjacent to a pixel p in the 2D image plane.

As long as E_s is regular, i.e. $E_s(0, 0) + E_s(1, 1) \leq E_s(1, 0) + E_s(0, 1)$, the global minimum of $E(y)$ can be efficiently found by max-flow/min-cut computation [17, 1]. For this purpose, a graph $G = (\mathcal{V}, \mathcal{E})$ is defined. The set of graph nodes \mathcal{V} consists of all pixels $p \in P$ and two additional terminal nodes s and t which represent foreground and background in the segmentation. The set of directed edges \mathcal{E} connects all pixels p to their neighbors $q \in N_2(p)$. Weights are assigned to these edges as specified by the smoothness term $E_s(y_p, y_q)$. In addition the set of edges \mathcal{E} connects each pixel to the two terminal nodes s and t with weights specified by $E_d(y_p)$. Minimizing $E(y)$ corresponds to finding the optimal cut $\mathcal{C} \subset \mathcal{E}$ such that no path exists between the terminal nodes s and t in the graph $G_{cut} = (\mathcal{V}, \mathcal{E} - \mathcal{C})$. The cut is optimal in the sense that the sum of all edge weights of all edges included in the cut is minimal.

Often graph cut approaches use a definition of E_s which penalizes for discontinuities in the segmentation for neighbored pixels of similar intensities [5]:

$$E_s(y_p, y_q) = \exp\left(-\frac{(x_p - x_q)^2}{2\sigma_s^2}\right) \cdot \frac{\delta(y_p, y_q)}{dist(p, q)}, \quad (2)$$

where x_p is the gray value of the image at pixel p and $dist(p, q)$ takes the distance between neighbored pixels into account. The Kronecker delta function $\delta(y_p, y_q)$ equals 0 if $y_p = y_q$ and 1 otherwise. This ensures that the energy term is regular.

For the segmentation of thin and elongated structures, like blood vessels, it is common to use an additional term $E_{gf}(y_p)$ that incorporates the flux of the gradient vector field into the segmentation. It has been shown that this can overcome the problem of “smoothing away” thin structures [23]. Flux is defined according to

$$F(p) = \sum_{q \in N_2(p)} \langle u_{pq}, v_q \rangle, \quad (3)$$

where u_{pq} is a unit vector oriented from pixel p to the neighboring pixel $q \in N_2(p)$ and vector v_p corresponds to the gradient vector at pixel p . This term can be seen as the flow of the gradient vector field through the contour of the segmented region. The corresponding unary potential $E_{gf}(y_p)$ is defined as:

$$E_{gf}(y_p) = \begin{cases} \max(0, F(p)) & \text{for } y_p = 1 \\ -\min(0, F(p)) & \text{for } y_p = 0 \end{cases} \quad (4)$$

A detailed description on how to define edge weights for flux in graph cut is given in [16].

In a simple setting, the term $E_d(y_p)$ of Equation (1) can be defined as relying directly on the pixel intensities in the original gray value image. But, structures in electron microscopy images are often only recognizable by their texture in the local context. Therefore, we use the probabilistic output of a random forest classifier [6] trained on annotated data for membrane detection, similar to the approach in [25, 9]. We use Haar-like features as well as histograms over context regions to capture a discriminative representation of the central pixel with little computational cost. To account for the random forest classifier, we rename the data term to $E_{rf}(y_p)$ throughout the paper.

Taking the details explained above into account, our implementation of the state of the art segmentation method looks as follows:

$$E(y) = \sum_{p \in P} E_{rf}(y_p) + \lambda_s \sum_{\substack{p \in P \\ q \in N_2(p)}} E_s(y_p, y_q) + \lambda_{gf} \sum_{p \in P} E_{gf}(y_p). \quad (5)$$

Using gradient flux to enhance the segmentation of thin objects also has a drawback. In textured images the image gradient is not only very high at the desired segmentation borders, but also at other image regions with high contrast. Therefore the gradient flux can cause a large amount of false positives in the resulting segmentation. In addition we want to use the output of a trained membrane detector as data term for the segmentation. Experiments showed that gradient flux and smoothness alone is not sufficient to compensate for weakly detected membranes, as is illustrated in the following toy data setting. We generate an image, of a perfect membrane represented as straight black line on a white background. A weak classifier response is simulated by fading out a section of the line (Figure 1). Although the gradient flux and smoothness terms were calculated on the perfect, non-faded line, they cannot compensate for the weak E_{rf} input. The gradient enhances segmentation of the rim of the lines, but any attempt to make the segmented regions solid by using the smoothness term E_s leads to gaps in the membranes segmented. This problem is more aggravated on real data, since weak classifier responses often occur in the case of membranes which appear fuzzy in the image due to non orthogonal cutting or staining conditions. In these cases the gradient along the membrane is small and thus further limits the use of the gradient vector flux in the segmentation. To overcome this problem we introduce a novel energy term, that focuses on the principle of good continuation to close gaps along membranes.

To overcome the shortcomings of gradient flux, we in-




Figure 1. Toy example for membrane segmentation. The good continuation energy term is able to produce a solid segmentation where gradient flux fails. From top to bottom: (1) original perfect line, (2) line with a faded out segment as input for the data term E_{rf} , (3) with gradient flux, segmentation of borders is improved, (4) attempt to close segmented structures by additional use of the smoothness term E_s , (5) solid segmentation using only E_{rf} , and the directed term E_{gc} .

troduce a directional energy term that is based on the perceptual concept of good continuation. Intuitively, lines as well as membranes are directed structures. By the principle of good continuation well classified parts of directed structures should enforce smoothness in labels along their orientation. This is formulated by $E_{gc}(y_p, y_q)$:

$$E_{gc}(y_p, y_q) = | \langle v_p, u_{pq} \rangle | \cdot \exp \left(- \frac{(x_p - x_m)^2}{2\sigma_{gc}^2} \right) \cdot \frac{\delta_{\rightarrow}(y_p, y_q)}{dist(p, q)}, \quad (6)$$

where u_{pq} is a unit vector with the orientation of a straight line between pixels p and q , and v_p is a vector directed along the membrane. The length of v_p reflects the orientedness of the image at p . For this purpose we use a directed filter consisting of a straight line with a thickness equal to the average membrane thickness in the training images. $\langle v_p, u_{pq} \rangle$ is then estimated by the response to this filter oriented according to u_{pq} . The value of x_m is given as the average gray value of membrane pixels and σ_{gc}^2 can be estimated as the variance of these gray values. Thus, the difference $(x_p - x_m)$ weights the energy term according to the similarity of x_p to the typical gray value of a membrane.

In contrast to Equation 2 the factor $\delta_{\rightarrow}(y_p, y_q)$ is not symmetric. Instead $\delta_{\rightarrow}(y_p, y_q) = 1$ for $y_p = 1, y_q = 0$ and $\delta_{\rightarrow}(y_p, y_q) = 0$ for all other cases. This asymmetric definition ensures that E_{gc} only penalizes for cuts that violate the smoothness along the direction of membrane pixels. Although $\delta(y_p, y_q)$ is not symmetric, it is still regular and thus the global optimality of the resulting segmentation is assured (see also [5, 26]).

In addition we incorporate information from adjacent

sections into the segmentation using:

$$E_{na}(y_p, y_q) = m_q \cdot | < v_p, v_q > | \cdot \frac{\delta_{\leftarrow}(y_p, y_q)}{\text{dist}(p, q)}, \quad (7)$$

where m_q is the probability of pixel q being a membrane and v_p is the large eigenvector of the Hessian at pixel p multiplied by the corresponding eigenvalue. Thus, a high confidence in pixel q being a membrane is propagated to the next section if the corresponding region is similarly oriented. This has the benefit, that it is unlikely for false positive detections to be propagated to the next section, as they will not have a similar oriented correspondence in the other image. $\delta_{\leftarrow}(y_p, y_q)$ again is defined asymmetrically, such that only $E_{na}(0, 1)$ is penalized. In Equation (8) the corresponding sum runs over all neighbors $N_3(p)$, which are defined as neighbored pixels in adjacent sections (3 dimensional). To solve the correspondences between images we followed the nonlinear warping method described in [14].

From our experience, the use of gradient flux is likely to lead to false positive membrane segmentations due to texture in the images. Thus, we decided to omit gradient flux in the final energy term:

$$\begin{aligned} E(y) = & \sum_{p \in P} E_{rf}(y_p) + \lambda_{gc} \sum_{\substack{p \in P, \\ q \in N_2(p)}} E_{gc}(y_p, y_q) \\ & + \lambda_s \sum_{\substack{p \in P, \\ q \in N_2(p)}} E_s(y_p, y_q) + \lambda_{na} \sum_{\substack{p \in P, \\ q \in N_3(p)}} E_{na}(y_p, y_q). \end{aligned} \quad (8)$$

Although this energy term incorporates information from adjacent sections, the main focus of the segmentation is two dimensional. This is due to the fact that the resolution of TEM images is high (about 5nm per pixel), but along the vertical direction of the image stack, the resolution is limited by the section thickness of the sample. Even very skilled human operators can at best cut sections of 40nm thickness. Thus, resolution along the z direction is an order of magnitude worse than the resolution along the x-y plane (see also Figure 5). This strongly favors a 2D segmentation approach.

3. Experiments and results

We evaluate the proposed method on two different neuroanatomical data sets of ssTEM images. Data set 1 shows 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 40 images with 512x512 pixels, divided into two sub volumes of 10 and 30 sections. The resolution is 3.7 nm per pixel in the image plane and section thickness is 50 nm. Data set 2 was taken from layer 4 of Area 17 (primary visual cortex) of one adult cat. The

data set consists of 9 images with 4312x3018 pixels. Resolution is 1.38 nm per pixel in the image plane and section thickness is 40 nm. Both data sets resemble average image quality from neuroanatomy projects and were fully manually segmented by human experts using *TrakEM2* [7], a free plugin for *ImageJ* [2]. The samples for these data sets were not only taken from different brain types (insect and mammal), but also prepared with different staining procedures and recorded at different magnifications, leading to very different image characteristics. As can be seen in Figure 4, the membranes of data set 1 appear very dark in the images, but also fuzzy in a lot of areas. Data set 2 contains considerably more texture caused by sub cellular elements like vesicles, microtubules and mitochondria inside the cells. Despite these different challenges, the new approach yields good segmentations on both data sets, demonstrating the great robustness against varying image characteristics.

In addition the proposed framework was applied to satellite images of San Francisco. The extracted features and the classifier employed for the segmentation of streets are the same as for the membrane segmentation, as the focus of the evaluation is on the different graph cut energy terms and not the quality of the classifier.

For the evaluation of the perceptual grouping framework all data sets were split into training and test sets. For the drosophila data set, the small volume was used for training and the large volume for testing. For the cat data set only nine annotated images are available, therefore leave one out cross validation was used in this case. The random forest classifier ensemble consists of 500 trees. The trees were build with 10 out of 116 features randomly selected for each split. The plots in Figure 2 show the precision and recall of the segmentations on all test images. Here precision can be seen as the probability that a pixel classified as foreground by the automatic segmentation is also marked as foreground in the hand labels given. Recall corresponds to the probability that a foreground pixel is detected. For the membrane segmentation on both data sets the perceptual grouping framework was evaluated with $\lambda_s = 0.6$ and for the evaluation of E_{na} , $\lambda_{gc} = 1.6$. For the state of the art segmentation with gradient flux λ_{gf} was set to 5. For the San Francisco street data set the parameters are $\lambda_s = 0.8$ and $\lambda_{gf} = 10$. The street data set contains no 3d information, therefore E_{na} is not included in the evaluation. In all three data sets the good continuation energy term E_{gc} leads to a considerable improvement in recall. As can be seen in the example segmentations in Figure 2 the loss in precision is mainly caused by thicker membrane segmentations. For the 3d reconstruction of neuronal structures, high recall with closed contours is more desirable than a good precision, as long as no splitting errors are introduced. Therefore, we also evaluate the number of splitting and merging

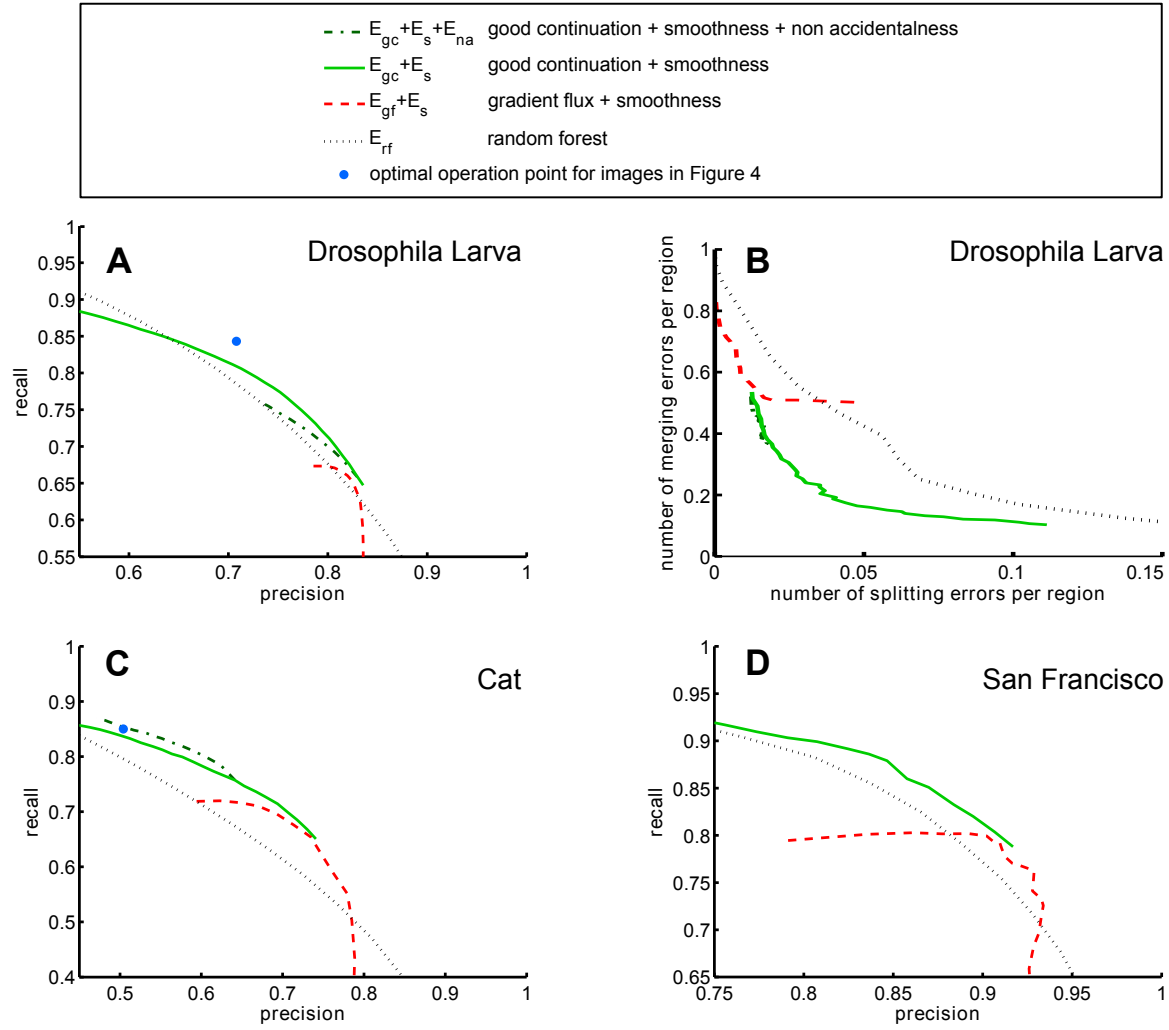


Figure 2. Comparison of the proposed framework $E_{rf} + E_{gc} + E_s + E_{na}$ against the state of the art gradient flux energy term $E_{rf} + E_{gf} + E_s$. The plots depict the precision and recall performance per pixel over all test images. The combination of random forests with perceptual grouping constraints yields a considerable improvement in recall. The split and merge error plot (B) demonstrates that the improved recall is caused by gap completion which is highly desirable for 3d reconstructions of neuronal structures.

errors per region for the drosophila larva data set (see plot B in Figure 2). The plot shows the number of splits and merges per region in the automatically obtained segmentation with respect to the manual ground truth. The splitting error counts the number of times a region from the ground truth segmentation is overlapped by more than one region from the automated segmentation. In order to be significant the split has to be bigger than one percent of the ground truth region. The merging error is the same in reverse. It counts how often a region of the automated segmentation is overlapped by more than one region in the ground truth. Thus, the error is increased if a segmented membrane is not closed and ground truth regions are merged in the auto-

mated segmentation. A low splitting and merging error per region preserves the duality between membranes and enclosed regions and thus enables automatic reconstructions of neuronal structures. Plot B in Figure 2 clearly demonstrates the substantial improvement in the segmentation by our good continuation term. The cat brain data set does not contain enough regions to provide meaningful results in terms of splitting and merging errors, due to the large size of the neuronal structures in these images. The term E_{na} that incorporates information from adjacent sections, is very beneficial for the cat data set and leads to an additional increase in recall. For the drosophila data set, the influence of adjacent sections is smaller than for the cat data set be-

cause the drosophila images change significantly between sections.

Example segmentations of test images are given in Figure 4. The segmentation is very good with respect to texture caused by vesicles and microtubules, but mitochondria still pose a challenge. They are not only surrounded by a membrane, but also very similar to small dendrites in shape, leading to false positive detections. A possible solution to this problem would be to include extra labels for mitochondria in the training set and either make the random forest classifier more sensitive to these structures or train a second classifier specifically for mitochondria. This is part of our future research. A segmentation result for the San Francisco data set is given in Figure 3. Shown are the segmentation results with 0.85 precision for smoothness combined with good continuation (green) or gradient flux (red). Black pixels were marked as streets by both methods. Although this image is from a completely different domain, the segmentation result shows the same characteristic for both methods as for the electron microscopy images. The good continuation constraint leads to thicker segmentations, but improves the segmentation by gap completion, whereas the gradient flux gives false positive responses at background pixels with high contrast.

The split and merge error of our cost function is low enough, to obtain fully automatically reconstructed dendrites over several sections. An example reconstruction is shown in Figure 5. The five dendrites are segmented over 30 sections. Regions were automatically grouped between sections by maximum overlap. This simple tracking method will fail if the structures of interest are not orthogonal to the cutting direction. Improvement of region tracking in more complex scenarios is the main focus of our future research. Also shown in this Figure are cutting planes through the image volume. The very low resolution of the volume in the direction orthogonal to the cutting plane is clearly visible. Because of this difference in resolution we decided to focus our segmentation on the image plane.

4. Conclusion

The framework introduced in this paper addresses one of the main bottlenecks for 3D reconstructions in neuroanatomy: the fully automated segmentation of membranes in ssTEM images. The architecture comprises a random forest for classifying single pixels, and novel energy terms for membrane segmentation with graph cut optimization. Large scale quantitative evaluation experiments demonstrated the algorithms performance on cat and drosophila larva brain.

In summary the proposed framework is characterized by the following benefits: (i) *local to global optimization*: a random forest classifier estimates the probability for a membrane locally, while a regular cost function guarantees a




Figure 3. Example segmentation at 0.85 precision for the San Francisco street data set. Green pixels are positive detections with the good continuation constraint, red pixels are positive detections by smoothness and gradient flux, black pixels were marked by both methods as streets. Segmentation by good continuation loses precision by thickening the detected streets, but gains additional recall by gap completion. Gradient flux loses precision by false positive detections at high gradient contours.

global optimum employing graph cuts. (ii) *good continuation*: novel energy terms allow for contour completion in situations where gradient flux based methods fail. (iii) *robustness*: the algorithm produces proper results even on different animal species. (iv) *consistency*: we have successfully reconstructed a 3D model of dendrites based on the consistent segmentation of an image stack with 30 slices. (v) *excellent performance*: the presented algorithm outperforms the state of the art on all quantitatively evaluated real world scenarios.

Acknowledgment

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.




Figure 4. Example images and segmentations from two data sets. **Upper row:** drosophila larva, **lower row:** cat. From left to right: original image, automatic segmentation with perceptual grouping constraints, manual labels. Most membranes are correctly segmented. The algorithm copes well with textured regions of vesicles and microtubuli. False positive detections are mainly caused by mitochondria

References

- [1] An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(9):1124–1137, September 2004. [2](#)
- [2] M. Abramoff, P. Magelhaes, and S. Ram. Image processing with imagej. *Biophotonics International*, 11(7):36–42, 2004. [4](#)
- [3] B. Andres, U. Köthe, M. Helmstaedter, W. Denk, and F. A. Hamprecht. Segmentation of sbfsem volume data of neural tissue by hierarchical classification. In *Proc. DAGM*, pages 142–152, 2008. [1](#)
- [4] I. Arganda-Carreras, C. O. S. Sorzano, R. Marabini, J. M. Carazo, C. O. de Solorzano, and J. Kybic. Consistent and elastic registration of histological sections using vector-spline regularization. In *CVAMIA*, pages 85–95, 2006. [1](#)
- [5] Y. Boykov and G. Funka-Lea. Graph cuts and efficient n-d image segmentation. *Int. J. Comput. Vision*, 70(2):109–131, Nov. 2006. [1](#), [2](#), [3](#)
- [6] L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, 2001. [3](#)
- [7] A. Cardona. Trakem2: an imagej-based program for morphological data mining and 3d modeling. In *Proc. ImageJ User and Developer Conference*, 2006. [4](#)
- [8] W. Denk and H. Horstmann. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *Curr. Opin. Neurobiol.*, 2006. [1](#)
- [9] P. Dollar, Z. Tu, and S. Belongie. Supervised learning of edges and object boundaries. In *Proc. CVPR*, June 2006. [3](#)
- [10] K. M. Harris, E. Perry, J. Bourne, M. Feinberg, L. Ostroff, and J. Hurlburt. Uniform serial sectioning for transmission electron microscopy. *J. Neurosci.*, 2006. [1](#)
- [11] B. Hoffpauir, B. Pope, and G. Spirou. Serial sectioning and electron microscopy of large tissue volumes for 3D analysis and reconstruction: a case study of the calyx of held. *Nature Protocols*, 2(1):9–22, 2007. [1](#)
- [12] V. Jain, J. Murray, F. Roth, S. Turaga, V. Zhigulin, K. Briggman, M. Helmstaedter, W. Denk, and H. Seung. Supervised learning of image restoration with convolutional networks. *Proc. ICCV*, pages 1–8, October 2007. [1](#)
- [13] E. Jurrus, R. Whitaker, B. W. Jones, R. Marc, and T. Tasdizen. An optimal-path approach for neural circuit reconstruction. *Proc. ISBI*, pages 1609–1612, 2008. [1](#)




Figure 5. Five dendrites fully automatically reconstructed over 30 sections from drosophila larva. The cutting planes surrounding the dendrites show the good resolution of the data volume in xy-direction in contrast to the low resolution in z-direction.

- [14] V. Kaynig, B. Fischer, and J. M. Buhmann. Probabilistic image registration and anomaly detection by nonlinear warping. In *Proc. CVPR*, pages 1–8, 2008. 1, 4
- [15] G. Knott, H. Marchman, D. Wall, and B. Lich. Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. *J. Neurosci.*, 2008. 1
- [16] V. Kolmogorov and Y. Boykov. What metrics can be approximated by geo-cuts, or global optimization of length/area and flux. *Proc. ICCV*, 1:564–571, Oct. 2005. 3
- [17] V. Kolmogorov and R. Zabini. What energy functions can be minimized via graph cuts? *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(2):147–159, Feb. 2004. 2
- [18] P. Koshevoy, T. Tasdizen, R. Whitaker, B. Jones, and R. Marc. Assembly of large three-dimensional volumes from serial-section transmission electron microscopy. In *Proc. MI-AAB*, pages 10–17, 2006. 1
- [19] J. Macke, N. Maack, R. Gupta, W. Denk, B. Schkopf, and A. Borst. Contour-propagation algorithms for semi-automated reconstruction of neural processes. *J. Neurosci. Methods*, 167(2):349–357, 2008. 1
- [20] Y. Mishchenko. Automation of 3d reconstruction of neural tissue from large volume of conventional serial section transmission electron micrographs. *J. Neurosci. Methods*, 176:276–289, 2009. 1
- [21] A. V. Reina, E. Miller, and H. Pfister. Multiphase geometric couplings for the segmentation of neural processes. In *Proc. CVPR*, June 2009. 1
- [22] S. C. Turaga, J. F. Murray, V. Jain, F. Roth, M. Helmstaedter, K. Briggman, W. Denk, and H. S. Seung. Convolutional networks can learn to generate affinity graphs for image segmentation. *Neural Comput.*, 22(2):511–538, 2010. 1
- [23] A. Vasilevskiy and K. Siddiqi. Flux maximizing geometric flows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(12):1565–1578, Dec 2002. 1, 2
- [24] N. Vu and B. Manjunath. Graph cut segmentation of neuronal structures from transmission electron micrographs. pages 725–728, Oct. 2008. 1
- [25] M. Wels, G. Carneiro, A. Aplas, M. Huber, J. Hornegger, and D. Comaniciu. A discriminative model-constrained graph cuts approach to fully automated pediatric brain tumor segmentation in 3-d mri. *Proc. MICCAI*, 2008. 3
- [26] J. Winn and J. Shotton. The layout consistent random field for recognizing and segmenting partially occluded objects. In *Proc. CVPR*, 2006. 3