

Large-scale automatic reconstruction of neuronal processes from electron microscopy images



Verena Kaynig^{a,*}, Amelio Vazquez-Reina^{a,b}, Seymour Knowles-Barley^c, Mike Roberts^a, Thouis R. Jones^{a,c}, Narayanan Kasthuri^c, Eric Miller^b, Jeff Lichtman^c, Hanspeter Pfister^a

^a School of Engineering and Applied Sciences, Harvard University, United States

^b Department of Computer Science at Tufts University, United States

^c Department of Molecular and Cellular Biology, Harvard University, United States

ARTICLE INFO

Article history:

Received 26 March 2014

Received in revised form 2 November 2014

Accepted 6 February 2015

Available online 2 March 2015

Keywords:

Large-scale neuron reconstruction

Serial section electron microscopy

Random forest

Segmentation fusion

Proofreading

ABSTRACT

Automated sample preparation and electron microscopy enables acquisition of very large image data sets. These technical advances are of special importance to the field of neuroanatomy, as 3D reconstructions of neuronal processes at the nm scale can provide new insight into the fine grained structure of the brain. Segmentation of large-scale electron microscopy data is the main bottleneck in the analysis of these data sets. In this paper we present a pipeline that provides state-of-the art reconstruction performance while scaling to data sets in the GB-TB range. First, we train a random forest classifier on interactive sparse user annotations. The classifier output is combined with an anisotropic smoothing prior in a Conditional Random Field framework to generate multiple segmentation hypotheses per image. These segmentations are then combined into geometrically consistent 3D objects by segmentation fusion. We provide qualitative and quantitative evaluation of the automatic segmentation and demonstrate large-scale 3D reconstructions of neuronal processes from a 27,000 μm^3 volume of brain tissue over a cube of 30 μm in each dimension corresponding to 1000 consecutive image sections. We also introduce Mojo, a proofreading tool including semi-automated correction of merge errors based on sparse user scribbles.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Brain imaging modalities such as diffusion tensor MRI or functional MRI provide important information about the brain and the connectivity between brain regions (Seung, 2012). However, at a resolution of a cubic millimeter per voxel they provide little data about connectivity between individual neurons. Information about the anatomy and connectivity of neurons can provide new insights into the relation between the brain's structure and its function (Marc et al., 2013; Helmstaedter and Mitra, 2012; Denk et al., 2012; Lee and Reid, 2011; Seung, 2009). Such information may provide insights into the physical underpinnings of common serious disorders of brain function such as mental illnesses and learning disorders Kuwajima et al. (2013b), Penzes et al. (2011). Furthermore, information about the individual strength of synapses or the number of connections between two cells has important implications for computational neuroscience and theoretical analysis of neuronal networks (Valiant, 2006). As the

resolution of light microscopy is generally limited by diffraction, electron microscopy (EM) is a better imaging modality to resolve the brain at the level of synapses and thus provides insight into the anatomy and connectivity of neurons at nm resolution. To reconstruct the neuronal circuit at the level of individual cells, the field of neuroanatomy faces the challenge to acquire and analyze data volumes that cover a brain tissue volume large enough to allow meaningful analysis of circuits and detailed enough to detect synapses and thus the connectivity structure of the circuit. Recently, significant progress has been made in the automation of sample preparation (Hayworth et al., 2006) and automatic image acquisition (Kuwajima et al., 2013a; Bock et al., 2011; Knott et al., 2008; Denk and Horstmann, 2004) for electron microscopy. These techniques allow neuroscientists to acquire large datasets in the GB-TB range. Briggman and Bock (2012) provide an overview of different sample preparation and electron microscopy techniques used for Connectomics. With a resolution of 5 nm per pixel, and a section thickness of 50 nm, one cubic millimeter of brain tissue results in 20,000 sections with 40 Gigapixels per image, leading to an image volume of 800 TB. For comparison, this volume corresponds to the size of one voxel in an fMRI data set. With data sets

* Corresponding author.

this size, manual analysis is no longer feasible, leading to new challenges in automated analysis and visualization.

In this paper we present a pipeline for semi-automated 3D reconstruction of neurons from serial section electron microscopy images. The pipeline is designed to address large data sets, while reducing user interaction to the initial training of a random forest classifier on manually annotated data and computer aided proof-reading of the automatic reconstruction output. Our experiments demonstrate that the proposed pipeline yields state-of-the-art reconstruction results, based on sparse annotations of only ten EM images (1024×1024 pixels). We provide quantitative evaluation for each step of the pipeline and an example of a reconstructed volume of $27,000 \mu\text{m}^3$, which to our knowledge is the largest volume of conventionally stained mammalian brain tissue reconstructed automatically (see Fig. 1).

Some of the work in this paper has been previously published (Kaynig et al., 2010a; Vazquez-Reina et al., 2011; Roberts et al., 2011). However, this is the first time we publish the complete reconstruction pipeline and its application to large data. Specifically the novel contributions in this paper are:

- We demonstrate that interactively training a random forest classifier for membrane detection not only reduces the manual annotation effort, but leads to significantly better cell region segmentations measured in terms of variation of information against manual annotated data.
- We combine the cell region segmentation of Kaynig et al. (2010a) with the segmentation fusion of Vazquez-Reina et al. (2011) into a consistent pipeline leading to long-range reconstructions of neuronal processes over $30 \mu\text{m}$ of brain tissue (up to 1000 image sections).
- We extend the segmentation fusion approach to allow for branching structures.
- We enable parallel processing of sub volumes via a pairwise matching scheme of segmented blocks into one consistent reconstruction volume.
- We provide large-scale reconstruction results covering a volume of $27,000 \mu\text{m}^3$. To our knowledge we are the first to achieve automatic reconstructions of individual spine necks in anisotropic serial section electron microscopy data prior to manual proofreading.
- Finally, we introduce Mojo, a semi-automated proofreading tool, utilizing sparse user scribbles as described by Roberts et al. (2011) to correct for merge errors in the 3D reconstruction.




Fig. 1. We propose a pipeline to automatically reconstruct neuronal processes from large-scale electron microscopy image data. The target volume consists of 1000 images with a size of 5120×5120 pixels, corresponding to $27,000 \mu\text{m}^3$ of mammalian brain tissue. With 8 bits per pixel, the full data volume is 25 GB in size.

2. Related work

Automated reconstruction of neuronal processes has received increased attention in recent years. With electron microscopy techniques acquiring large volumes automatically, automated analysis is becoming the major bottleneck in gaining new insights into the functional structure of the brain at nm scale. The task to reconstruct the full neuroanatomy including synaptic contacts is referred to as *connectomics* in the literature (Lichtman and Sanes, 2008). A number of software packages have been developed to aid the user in manual annotation of the images (Helmstaedter et al., 2011; Cardona et al., 2010; Fiala, 2005). A complete overview of the different tools and their strength and limitations is provided by Helmstaedter and Mitra (2012). In addition, semi-automatic methods have been developed to facilitate the manual segmentation process (Roberts et al., 2011; Sommer et al., 2011; Straehle et al., 2011; Chklovskii et al., 2010; Vazquez-Reina et al., 2009).

In the area of fully automatic neuron reconstruction, significant improvement has been made for the segmentation of isotropic image data using a special staining method to facilitate the segmentation (Andres et al., 2012b; Andres et al., 2012a; Turaga et al., 2010; Andres et al., 2008; Jain et al., 2007). While these methods yield good performance for long range reconstructions, they sacrifice the staining of biologically relevant internal cell structures like vesicles or mitochondria to simplify the segmentation problem. Without staining these cell organelles, identification of synapses relies entirely on geometrical features, like the apposition of spines and boutons. daCosta and Martin (2011) have shown that geometrical features are not sufficient to identify synapses. In this paper we provide long range reconstructions with conventional osmium stained images, preserving all structural information for biological analysis of the data, such as synapse identification or automatic mitochondria reconstruction (Lucchi et al., 2012; Giuly et al., 2012).

While isotropic volume data enables the direct use of 3D segmentation methods for reconstruction, the microscopy techniques for these volumes are either limited in resolution to 30 nm voxels (Denk and Horstmann, 2004) or in the field of view to $20 \mu\text{m}^2$ (Knott et al., 2008). Serial section imaging is the only technique to record data volumes of millions of cubic micrometers (Bock et al., 2011). The tissue sample is cut into ultra thin sections of 30 nm and each section is imaged with an electron microscope typically at a resolution of $3\text{--}5 \text{ nm}$ per pixel. The z resolution of the resulting data volume is limited to 30 nm leading to an anisotropic data volume. An interesting work by Hu et al. (2013) aims at enhancing the z resolution by leveraging tomographic projections, but acquiring the necessary tilt images so far has not been automated for large-scale image acquisition.

Automatic neuron reconstruction methods for anisotropic serial section data typically focus on segmenting 2D neuronal regions in the high resolution images (Knowles-Barley et al., 2011; Jurrus et al., 2010; Kaynig et al., 2010a) or on grouping 2D regions across multiple sections into 3D neuronal processes (Funke et al., 2012; Vitaladevuni et al., 2010; Kaynig et al., 2010b; Jurrus et al., 2008), though some work also addresses both steps, the region segmentation and the grouping across sections together (Vazquez-Reina et al., 2011; Chklovskii et al., 2010; Mishchenko, 2009). To our knowledge Chklovskii et al. (2010) describe the only pipeline so far that addresses large-scale reconstructions in the order of thousands of μm^3 . They divide the original large EM data volume into biologically relevant sub volumes of about $3000 \mu\text{m}^3$ which are then segmented and reconstructed. In this paper we demonstrate successful segmentation of a volume that is nine times larger than the result shown by Chklovskii et al. (2010). Our experiments

also demonstrate that the employed CRF framework yields better neuronal region segmentations than their use of watersheds, leading to a reduction in proofreading effort.

3. Overview

We now provide an overview of our reconstruction workflow (see Fig. 2), as well as evaluation metrics for neuron segmentation and the data sets used for all experiments throughout this paper.

3.1. Workflow

For automatic segmentation methods, serial section imaging is challenging, as the resulting image data is highly anisotropic. While the xy image resolution for each section is only limited by the resolution of the microscope, the z -resolution is limited by the section thickness of about 30 nm. For our pipeline we assume that the image data has been previously aligned. While registration and alignment for large electron microscopy stacks is a topic of ongoing research (Wang et al., 2014; Saalfeld et al., 2012), it is not the focus of this paper.

Fig. 2 provides an illustration of the entire workflow. The first part of the pipeline concentrates on the 2D segmentations of the high resolution section images. We first train a random forest classifier on interactive manual annotations for membrane detection. Then, we generate multiple segmentation hypothesis per section based on the classification output. Our experiments demonstrate that thresholding the membrane probability map at different intervals combined with anisotropic smoothing in a conditional random field (CRF) framework is superior to watershed segmentations of the membrane probability map (Kaynig et al., 2010a). We modified

the original anisotropic smoothing prior to emphasize the importance of the membrane probability map over the original gray value images, leading to an improvement in segmentation performance.

Subsequently, we leverage the previously obtained 2D segmentations and group these into geometrical consistent 3D objects using segmentation fusion (Vazquez-Reina et al., 2011). This step is especially challenging for large data sets, as geometrically consistency requires context information across multiple sections. We reduced the number of features used to measure region similarity to streamline the fusion computation. In addition, we extend the original segmentation fusion model (Vazquez-Reina et al., 2011) to allow for the reconstruction of branching structures. We evaluate the fusion step of the pipeline and compare bipartite matchings of globally optimal groupings of sub volumes with a greedy optimization scheme.

In the final step, the segmentation output has to be proofread by a user, to ensure correct geometries. As fully manual proofreading is labor-intensive and practically unfeasible for large volumes, we introduce Mojo, a semi-automatic proofreading tool, that leverages sparse user scribbles to correct merge errors in the automatic segmentation (Roberts et al., 2011).

3.2. Evaluation measure

There are two types of errors: split errors and merge errors. In 2D segmentation, a split error refers to a single region being split into two or more regions in the segmentation due to false positive cell boundary detections. A merge error is caused by a gap in the segmented cell boundaries, leading to separate regions being merged into one region in the automatic segmentation. Both errors

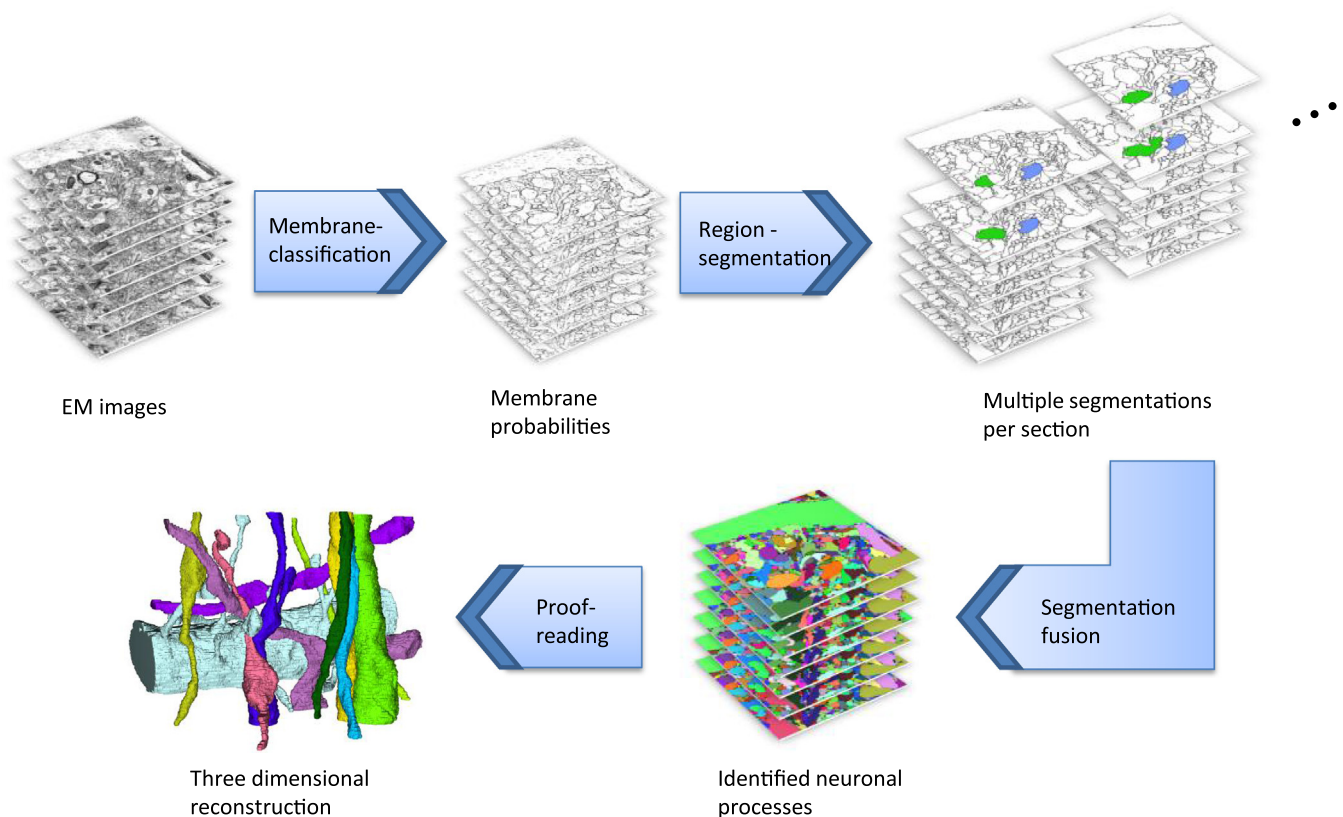


Fig. 2. Illustration of our complete workflow for large-scale neuron reconstruction. First, a random forest classifier is trained by interactive sparse annotations to detect membranes in the images. Then we generate multiple region segmentation hypotheses per section. Subsequently, we find the three dimensional reconstructions as geometrically consistent segmentations across multiple sections. Finally, a manual proofreading step ensures accurate reconstructions of neuronal processes of interest.

can also occur during region grouping in 3D. Missing a branch, for example, can lead to a split error, whereas merging branches incorrectly can merge two different neural processes into one object.

The quality of segmented regions does not directly correlate with boundary detection performance. A small number of false negatives in the boundary detection can lead to merged regions and thus large errors in the underlying segmentation, whereas false positives can be neglectable as long as they do not introduce a new region to the segmentation. To account for this difference, region-oriented segmentation tasks can be evaluated using clustering evaluation measures. The idea is that pixels with the same label form a cluster. The whole segmentation can then be compared to the pixel clustering obtained from the manual annotation. Rand Index (RI) and Variation of Information (VI) are two clustering evaluation measures that have been employed to evaluate region segmentation performance (Arbeláez et al., 2011). Both measures are invariant to label permutations. The main difference between VI and RI in the context of Connectomics is that VI is less sensitive to region sizes and rescaling, and therefore correlates better to proofreading effort than RI (Nunez-Iglesias et al., 2013). Thus, we measure the quality of our segmentation by comparing it to a manual annotation using variation of information.

In contrast to the Rand index, variation of information establishes a metric in segmentation space (Meila, 2007) with lower values capturing segmentations that are closer to each other. The variation of information metric is based on information theory and compares two segmentations S_1 and S_2 based on their entropy H and mutual information I :

$$VI(S_1, S_2) = H(S_1) + H(S_2) - 2I(S_1, S_2). \quad (1)$$

Entropy H measures the randomness of the segmentation, and mutual information I measures the information that the two segmentations share. Eq. (1) can be rewritten as $VI(S_1, S_2) = H(S_1|S_2) + H(S_2|S_1)$, thus, variation of information measures how much new information is obtained from one segmentation given that we have seen the other segmentation. All variation of information scores reported in the paper are computed using natural logarithms and hence are given in nats.

Variation of Information can be computed efficiently and is defined for arbitrary dimensions. Thus, the same evaluation criterion can be employed to evaluate the 2D region segmentations as well as the 3D region grouping step of our pipeline.

As drawing of cell boundaries requires more precision than clicking on objects, it is generally faster for a user to correct split errors than merge errors. Therefore, previous work on neuron segmentation has biased the output of the automatic reconstruction towards obtaining an over-segmentation of the data (Chklovskii et al., 2010). We follow a different approach in our work. Instead of biasing the pipeline towards split errors, we provide a 3D semi-automatic segmentation method in our proofreading tool to assist the user with the correction of merge errors. This allows us to focus on optimizing the overall error rate with an equal weighting of split and merge errors.

3.3. Data sets

To demonstrate the scalability of our reconstruction workflow we use a data set consisting of 1000 sections, with 5120×5120 pixels per image. The tissue is dense mammalian neuropil from layers 4 and 5 of the S1 primary somatosensory cortex of a 5 month old healthy C45BL/6J mouse. The images were taken at a resolution of 3 nm per pixel and downsampled by a factor of two, leading to a resolution in the image plane of 6 nm per pixel. The section thickness is 30 nm. The entire data set captures a tissue volume of $30 \times 30 \times 30 \mu\text{m}^3$. Our target is a volumetric reconstruction of

the 3D geometry of all neuronal processes in the data set. In contrast to center line tracings used in previous work by Helmstaedter et al. (2011), Bock et al. (2011) we concentrate on segmenting the whole volumetric geometry of the neurons. Center line tracings can be annotated faster by a user than complete volume reconstructions and are sufficient to evaluate the tracing of neuronal processes over long ranges in a given volume. In contrast, our complete volume segmentation enables us to also evaluate the automatic segmentation with respect to important fine structure details, such as spine necks (see Section 8), that are necessary to identify neuron connectivity.

In addition, we used a smaller volume consisting of 150 section images with 1024×1024 pixels per image. This volume has been densely annotated and thus captures the full range of different object sizes and variability on a small scale. Instead of carefully drawing cell boundaries, manual segmentation was performed by focusing on the regions corresponding to neuronal processes. Thus, small extracellular space between cells as well as thick or fuzzy membranes can lead to unlabeled pixels in the manual annotation. In order to preserve the duality between cell boundaries and annotated regions, we assign unlabeled pixels the label of the closest annotated region using seeded region growing. For our experiments we divided the data into three disjunct subsets for training, parameter validation and testing. All training and parameter validation is restricted to the first 75 images of the densely labeled data set. We use 10 images as training set, and 65 images as validation set. The remaining 75 images from the second half of the stack are used as test data, only after all parameters of the workflow have been fixed.

4. Region segmentation

While the texture characteristics of cell regions in electron microscopy images can vary significantly between different animal types and staining protocols, the basic appearance of the cell boundary membranes as thin, smooth, and elongated structures remains the same. Thus, instead of segmenting interior cell regions, we focus on segmenting the cell membranes to make our approach easily adaptable to a wide range of data.

4.1. Membrane classification

To learn the characteristics of membranes in the electron microscopy images, we train a random forest classifier based on sparse manual membrane annotations. Random forests combine the idea of bagging decision trees with random feature selection (Breiman, 2001). Each decision tree is built from a bootstrapped sample of the training data and at each node a random subset of the available features is selected to estimate the best split (Breiman, 2001). For prediction, the votes of all decision trees in the forest are accumulated. As each tree can be grown and queried independently, random forests are ideal for parallelization during training and prediction, as well as in an interactive training framework. In addition, random forests are robust against over-fitting, leading to good generalization performance with few manual annotations. The parameters to tune are the number of decision trees and the size of the feature subset used to determine the best split. We employ 300 trees and we set the number of features to the square root of the total number of features, which is the default suggested by Breiman (2001). To account for imbalanced training data, we follow the approach of Chen et al. (2004), and reduce the bootstrap sample for each tree to the size of the minority class.

The feature set extracted from the images is designed to capture the characteristics of membranes with little computational cost. Extracted features include the gray value, gradient magnitude,

Hessian eigenvalues, and difference of Gaussian for the image smoothed with Gaussian filters of different kernel sizes. In addition, we convolve the image with a steerable filter at different orientations. Each filter output serves as a feature, as well as the minimal, maximal, and average output of the steerable filter for different orientations at a pixel position.

Changes in the sample preparation process or different animal types can lead to significantly different data sets, requiring a retraining of the membrane classifier. Thus, our approach aims at minimizing manual interaction. We use an interactive training approach, similar to Sommer et al. (2011). The user provides sparse training annotations of membranes and the background class and interactively corrects the output of the classifier in a feedback loop.

There are two main benefits of this method. One benefit is that the annotation effort is efficiently guided towards challenging classifications and saves the user from annotating membranes that are already correctly classified. The second benefit is that users can weigh pixel errors in the classifier predictions implicitly by deciding to correct the segmentation or leave the prediction unchanged. The random forest classifier can only optimize for pixel misclassification and not for variation of information directly. While both evaluation measures are minimal for the correct segmentation, pixel misclassifications do not directly correspond to region segmentation performance. False positive membrane detections, for example on vesicles, can be neglectable as long as these do not lead to split errors in the segmented regions, false negative membrane detections can be neglectable as long as they do not introduce a gap in the outline of a neuronal region. The interactive training allows the user to provide a training set that is biased towards correcting pixel misclassifications which impact region segmentation performance.

Fig. 3 depicts an example of the interactive annotation. Our experiments demonstrate that in the context of small training samples, this interactive approach outperforms complete annotation of all membranes in the images (see Fig. 4).

4.2. Interactive training evaluation

To evaluate the interactive annotation against conventional training of fully annotated images, we chose a training set of five images out of the first 75 images of the 150 section data set. These images were manually selected to cover the variability in the data, from changes in image focus, contrast, or section thickness. The interactive training consists of multiple passes over the training images. Each pass consists of an annotation and a feedback phase. The user is presented with the current classification result




Fig. 3. An example of the interactive annotation workflow. Left: An original electron microscopy image, with overlaid membrane annotations (green) and background annotations (red). Right: The thresholded membrane probability map overlaid in red. The ellipses mark a split and merge error respectively, which could be corrected by additional annotations in the next iteration. Both images are 1.2 μm wide. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Evaluation of the interactive annotation approach using variation of information as the metric. Lower values correspond to better performance. After the third to fourth round of interactively correcting the classifier output (orange and green lines), the interactive approach outperforms training on fully annotated images (red). The blue line corresponds to the performance of the classifier trained on the last round of both sparse training sets together. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and then manually provides additional training samples to correct misclassifications. We repeated the interactive annotation procedure with a second data set of five training images. Fig. 4 demonstrates the median performance of the interactive sparse training in terms of variation of information. The performance is measured over the validation image data, which consists of the 65 images out of 75 that were not used for training annotations. To compare the performance of sparse interactive annotations to conventional batch training on fully annotated images, an expert labeled the center line of all membranes in all 10 images of both training sets. We chose the centerline to ensure that only true membrane pixels are in the training set. As membranes can appear fuzzy or smeared in the images, we neglect all pixels with a distance smaller or equal to five pixels from this center line, excluding possibly ambiguous examples from the training set. All remaining pixels are taken as background examples.

After the third to fourth pass over the images, the interactively trained classifier outperforms the classifier trained on fully annotated images. As demonstrated by the green and orange curves in Fig. 4, the performance of the second training set is similar to the first training set. Interestingly, the second pass on the second training set shows a significant degradation of the segmentation. In this step the annotator introduced background labels on mitochondria, leading the classifier to misclassify fuzzy membranes as background and thus to introduce gaps in the cell boundaries. In the next step these membrane misclassifications are corrected, leading to an improvement in the segmentation performance. The blue line corresponds to the performance of a classifier trained on the final ten training images of both interactive sparsely annotated sets. This is the classifier we use for the remaining steps of the pipeline.

4.3. 2D segmentation

The random forest classifier captures the main image characteristics of membranes with little manual annotation data. Previous work has shown that anisotropic smoothing of images is beneficial for the segmentation of membranes (Mishchenko, 2009). We follow the approach of Kaynig et al. (2010a), which combines the membrane probability output of the random forest classifier with an anisotropic smoothing prior for gap completion in a Conditional Random Field (CRF). In a CRF, the binary segmentation of all pixels as foreground or background is estimated by

maximizing the a posteriori probability of the labels y given the observed data x :

$$p(y|x) \propto \exp \left(\sum_{i,p \in P} \lambda_i F_{\text{state}_i}(y_p, x, p) + \sum_{j,p \in P, q \in N(p)} \lambda_j F_{\text{trans}_j}(y_p, y_q, x, p, q) \right). \quad (2)$$

F_{state_i} is a state feature function of the label $y_p \in \{0, 1\}$ at pixel $p \in P$, and the image intensity values x , and F_{trans_j} is a transition feature function of the labels y_p and their neighbored labels y_q in the 8-connected neighborhood $N(p)$. Intuitively, in our framework the state feature function estimates the probability of a single pixel as being foreground or background, whereas the transition feature function introduces dependencies between neighbored pixels, leading to smooth segmentations. Instead of maximizing the a posteriori probability of the labels y we minimize the negative logarithm, leading to the following energy term:

$$E(y) = \sum_{p \in P} E_{\text{rf}}(y_p) + \lambda_s \sum_{p \in P, q \in N(p)} E_s(y_p, y_q) + \lambda_{gc} \sum_{p \in P, q \in N(p)} E_{gc}(y_p, y_q). \quad (3)$$

The state function $F_{\text{state}_i}(y_p, x, p)$ corresponds to the data term $E_{\text{rf}}(y_p)$, which uses the output of the random forest classifier to specify the costs for label y_p being membrane or non-membrane. In Eq. (3) we omit the arguments for the observed data x and the pixel positions p, q to simplify the equation.

In addition, we include two smoothness terms which correspond to transition feature functions in Eq. (2). One is an isotropic smoothness term $E_s(y_p, y_q)$, which penalizes for discontinuities in the segmentation for neighboring pixels of similar intensities. This smoothness term is widely used in graph cut approaches (Boykov and Funka-Lea, 2006):

$$E_s(y_p, y_q) = \exp \left(-\frac{(x_p - x_q)^2}{2\sigma_s^2} \right) \cdot \frac{\delta_-(y_p, y_q)}{\text{dist}(p, q)}, \quad (4)$$

where x_p is the gray value of the image at pixel p and $\text{dist}(p, q)$ takes the Euclidean 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. Thus, the Kronecker delta function penalizes label changes, whereas the first factor of the energy term alleviates this penalty for strong changes of contrast in the image.

The second smoothness term $E_{gc}(y_p, y_q)$ enhances the coliniarity of segmented pixels:

$$E_{gc}(y_p, y_q) = | \langle v_p, u_{pq} \rangle | \cdot \exp \left(-\frac{(1 - x_m)^2}{2\sigma_{gc}^2} \right) \cdot \frac{\delta_-(y_p, y_q)}{\text{dist}(p, q)}, \quad (5)$$

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 . To measure the orientation of the membrane we use a steerable filter consisting of a straight line with a thickness comparable to the membrane thickness in the training images. The term $\langle v_p, u_{pq} \rangle$ is then estimated by the response to this filter at the orientation corresponding to u_{pq} . The value of x_m is the probability of pixel x being a membrane, and σ_{gc}^2 can be estimated as the variance of these probabilities. Thus, the difference $(1 - x_m)$ weighs the energy term according to the confidence of the classifier in x_m being a membrane. In contrast to Eq. (4), the factor $\delta_-(y_p, y_q)$ is not symmetric. Instead $\delta_-(y_p, y_q) = 1$ for $y_p = 1, y_q = 0$ and $\delta_-(y_p, y_q) = 0$ 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.

The smoothness terms E_s and E_{gc} are submodular, i.e., $E(0, 0) + E(1, 1) \leq E(1, 0) + E(0, 1)$, and thus the global minimum

of $E(y)$ can be efficiently found by max-flow/min-cut computation (Kolmogorov and Zabini, 2004; Boykov and Kolmogorov, 2004; Boykov and Funka-Lea, 2006).

For this purpose, we define a graph $G = (\mathcal{V}, \mathcal{E})$. The set of graph nodes \mathcal{V} consists of all pixels $p \in P$ and two additional terminal nodes s and t that represent foreground and background in the segmentation. The set of directed edges \mathcal{E} connects all pixels p to their neighbors $q \in N(p)$. Weights are assigned to these edges as specified by the smoothness terms E_s and E_{gc} . In addition, the set of edges \mathcal{E} connects each pixel to two additional terminal nodes s and t with weights specified by E_{rf} . 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 $G_{\text{cut}} = (\mathcal{V}, \mathcal{E} - \mathcal{C})$. The cut is optimal in the sense that the sum of all edge weights of all edges included in the cut is minimal.

The optimal labeling y corresponds to a binary segmentation of the image into membrane and non-membrane pixels. As we are ultimately interested in the region segmentation of neuronal processes, we identify neuronal regions as connected background components and then use seeded region growing to obtain a complete tessellation of the image into segments corresponding to neuronal processes.

4.4. Region segmentation evaluation

To evaluate the performance of our 2D segmentation step we set the isotropic smoothing weight $\lambda_s = 0.6$ and the anisotropic smoothing weight $\lambda_{gc} = 0.1$. The values for these parameters were optimized over the 65 images of our validation set. A visual example of the output is provided in Fig. 5. Fig. 6 provides a quantitative comparison of different 2D segmentation methods over 75 test images that were not used for any parameter tuning. For these evaluations, we normalized the gray value images to a range between zero and one prior to thresholding the gray value. The random forest results present thresholding of the respective membrane probability map. Watersheds and CRF results both use the probability map of the interactively trained random forest classifier. For watersheds we suppressed all minima in the probability map whose depth is less than the given threshold. For the CRF we adjusted the data term to reflect the given threshold on the probability map. A random forest classifier trained on fully




Fig. 5. Example segmentation of a test EM image. Top left: original image. Top right: manual annotation. Bottom left: random forest output. Bottom right: CRF segmentation. The isotropic smoothing closes small regions caused by extracellular space between cells, whereas the anisotropic term prevents shrinking bias for long elongated structures and enhances gap completion. All images are 1.9 μm wide.




Fig. 6. Evaluation of membrane segmentations on test data. The top line (violet) corresponds to the performance of thresholded gray value images. The orange line demonstrates the performance by a random forest classifier trained on ten fully contoured training images. The blue line shows the evaluation for a random forest trained on interactive sparse annotations. The red line corresponds to watershed segmentations of the random forest probability map. The green line corresponds to the performance of the CRF framework in our pipeline. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

contoured images (orange line in Fig. 6) gives a large improvement in segmentation performance over direct segmentation of gray values, but the random forest trained with sparse interactive annotations demonstrates a better generalization to the test data. We also compare the output of the CRF framework with gap completion against watershed segmentations, which have been widely used in previous work to generate segmentations of neuronal structures in EM images (Vazquez-Reina et al., 2011; Straehle et al., 2011; Chklovskii et al., 2010; Andres et al., 2008). While the optimal watershed segmentation performs 0.03 better in terms of variation of information than the best thresholded random forest output, the CRF framework yields an additional improvement of 0.09 over the best watershed segmentation.

5. Segmentation fusion

5.1. Region grouping across sections

The previous steps of the pipeline focus on the segmentation of neuronal processes in the 2D image plane to take advantage of the high resolution provided by the electron microscope. To extract the 3D geometry of neuronal processes, these regions need to be grouped across sections. We follow the segmentation fusion approach of Vazquez-Reina et al. (2011) that allows for globally optimal groupings of regions across sections. The term fusion refers to the option to pick the best choice of geometrically consistent region groupings out of a set of possible segmentations for each section. The fusion problem is formulated as the maximum a posteriori labeling over a set of binary indicator variables. Each indicator variable s_i corresponds to a possible 2D region of a neuronal process, and each indicator variable l_j to a 3D link between regions of adjacent sections. If an indicator variable is activated (e.g., $s_i = 1$), the correspondent region is assumed to be selected for the final segmentation, and similarly for a 3D link l_j . Thus, a labeling of the indicator variables corresponds to a 3D segmentation of the whole data volume.

Following the model of a CRF described in Eq. (2) the fusion problem is modeled as:

$$p(s, l | r) \propto \exp \left(\sum_{i=1}^n F_{\text{segment}}(s_i, r, i) + \sum_{j=1}^m F_{\text{link}}(l_j, r, j) \right) \psi(s, l). \quad (6)$$

The two functions $F_{\text{segment}}(s_i, r, i)$ and $F_{\text{link}}(l_j, r, j)$ are state functions for the indicator variables s_i and l_j , r refers to the set of all regions obtained from the 2D region segmentation, and n and m are the total number of indicator variables s_i and l_j . To ensure that any large segment from the region segmentations can compete equally against a set of smaller regions covering the same 2D area, both state functions take the size of the corresponding regions into account. In addition, links between regions are weighted according to the similarity of the linked regions, leading to the following definitions:

$$F_{\text{segment}}(s_i, r, i) = \text{size}(r_i), \quad (7)$$

$$F_{\text{link}}(l_j, r, j) = \theta(r_{j,a}, r_{j,b}) \cdot (\text{size}(r_{j,a}) + \text{size}(r_{j,b})). \quad (8)$$

$r_{j,a}$ and $r_{j,b}$ are the two regions connected by the 3D link l_j and $\theta(r_{j,a}, r_{j,b})$ measures the similarity between two regions. In the original fusion formulation, Vazquez-Reina et al. (2011) defined θ in terms of the cross-correlation and displacement between the pair of segments that re connected by the link in question. We instead define the region similarity θ in terms of the minimum relative overlap size of the two regions. This definition does not take texture similarity into account, but it is computationally faster than cross-correlation while providing equally good region similarity measurements for our EM data:

$$\theta(r_{j,a}, r_{j,b}) = \min(\text{overlap}(r_{j,a}, r_{j,b}), \text{overlap}(r_{j,b}, r_{j,a})), \quad (9)$$

$$\text{overlap}(r_{j,a}, r_{j,b}) = \frac{|r_{j,a} \cap r_{j,b}|}{|r_{j,a}|}. \quad (10)$$

By using the minimum of the relative overlap, θ is based on the relative overlap with respect to the larger region. This definition is useful, because if a large region is overlapped by two smaller regions by 40% and 60% respectively, we want the link to the region covering 60% of the overlap to outweigh the link to the region covering 40%.

The compatibility function $\psi(s, l)$ in Eq. (6) is defined over the indicator variables and assigns zero-mass to configurations of the CRF that are unrealistic or undesirable given our domain knowledge of the problem. Specifically, we want each pixel of the segmentation to belong to no more than one neuron, thus we want to prevent overlapping of activated segments s_i in the same section image: $\sum_{i \in o_k} s_i \leq 1$ for every set of overlapping segments o_k . Furthermore, the selection of activated segments and links should yield good 3D continuity through the stack. We achieve this by making the selection of segments and links dependent on each other and rewarding the activation of segments that are connected by strong links:

$$\left(\sum_{j \in \text{TOP}_i} l_j \right) \leq |\text{TOP}_i| \cdot s_i, \quad \left(\sum_{j \in \text{BOT}_i} l_j \right) \leq |\text{BOT}_i| \cdot s_i, \quad (11)$$

where TOP_i and BOT_i are the sets of activated links connecting to segment s_i from the sections immediately above or below, respectively. The main idea is to make the activation of segments and links depend on each other. Whenever a link is activated, the corresponding connected segments have to be activated as well. Compared to the original fusion formulation, our constraint does allow for multiple links connecting to the same segment. This relaxed version allows for branching of segmented structures and thus can adequately model the geometry of neuronal cells.

In our experiments, we noticed that if we allow segments to connect with any number of overlapping segments from adjacent sections, we have to be careful to not over-merge the segmentation. To prevent branching from connecting too many objects, we require links to only connect segments with significant overlap between sections. Links connecting sections that do not or only minimally overlap are pruned from the solution space. To obtain the maximum a posteriori (MAP) solution to the whole segmentation fusion problem (Eq. (6)), we solve the following binary linear programming problem:

$$\begin{aligned} \text{argmax}_{s,l} \quad & \sum_{i=1}^n F_{\text{segment}}(s_i, r, i) + \sum_{j=1}^m F_{\text{link}}(l_j, r, j) \\ \text{s.t. } \quad & s_i, l_j \in \{0, 1\}, \\ & \psi(s, l) = 1. \end{aligned} \quad (12)$$

We solve this problem using a general-purpose binary linear programming solver (IBM, 2014).

5.2. Segmentation fusion evaluation

There are two important aspects for the evaluation of segmentation fusion: the benefit of using multiple segmentations per section on the 2D segmentation, and the performance with respect to the 3D region grouping into geometrically consistent objects. (see Fig. 7). To evaluate the 3D region grouping performance we compare against greedy agglomerative clustering (Kaynig et al., 2010b) and the Sopnet framework developed by Funke et al. (2012). Segmentation fusion and the Sopnet framework both significantly outperform agglomerative clustering by finding the globally optimal grouping with respect to a large volume context. Segmentation fusion outperforms Sopnet with an improvement of 0.06 in terms of variation of information. Both approaches leverage multiple segmentation hypotheses per section and obtain the optimal grouping by solving an integer linear programming problem. The main difference is that within Sopnet a classifier is trained to score the similarity between regions, whereas segmentation fusion relies on region overlap and size alone (see Section 5.1). It is possible that Sopnet could benefit from a training set larger than 75 images, but generating such a large training set would require a considerable effort of manual annotation. In addition, the feature extraction and classification to obtain the region similarity adds significant computational overhead to the region grouping. Thus, segmentation fusion is the better approach for our large-scale reconstruction effort.

To gain more insight into the segmentation performance we evaluate 2D segmentations with respect to split and merge errors. Fig. 8 compares the segmentation fusion output with the best single 2D segmentation obtained by the gap completion CRF framework as described in Section 4.4. Over the whole test set of 75 images, fusion gave a significant improvement in the overall segmentation performance, increasing the percentage of correctly




Fig. 7. Evaluation of 3D reconstructions obtained by segmentation fusion compared to agglomerative clustering and Sopnet (Funke et al., 2012). As a baseline we provide the variation of information score for ungrouped data. Lower scores correspond to better region groupings.

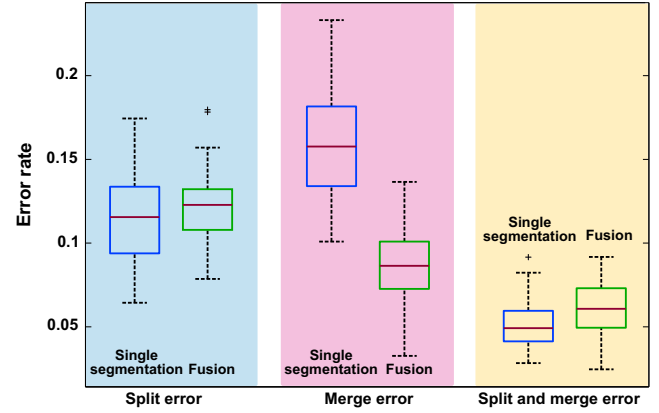


Fig. 8. Boxplots of the improvement of split and merge errors in the 2D segmentations. On each box, the central mark is the median and the edges of the box are the 25th and 75th percentiles. Outliers are marked with a + symbol. Compared to the best single segmentation from the gap completion CRF output, fusion reduces the overall error rate by 5% by leveraging multiple segmentations. The evaluation shows that this improvement is largely due to fewer merge errors.

segmented regions from 68%(±4%) to 73%(±3%). We define a correctly segmented ground truth region as having a reciprocal majority overlap of 60% or greater with a region from the automatic segmentation. Fig. 8 demonstrates that this improvement is mainly due to a correction of merge errors. While the split error rate is slightly increased by 1% the merge error rate is nearly halved, dropping from 15.8% for the single segmentation to 8.6% for the fusion result. Fig. 9 shows the segmentation performance with respect to region sizes. In total, error rates are lower for the larger regions than for the small regions. It is also noteworthy that the dominant error type changes from mainly merge errors for smaller regions to split errors for larger regions. For the most typical region size the errors are nearly balanced. Because smaller regions belong to thin flexible neuronal processes they change more prominently between adjacent sections than large regions belonging to thick neuronal processes. Thus, it is challenging to pick the right region segments based on the overlap criterion as described in (Eq. (10)).

6. Semiautomatic proofreading with Mojo

Manual proofreading is necessary in order to guarantee the correct topology of the neuron reconstruction. Fig. 10 shows an

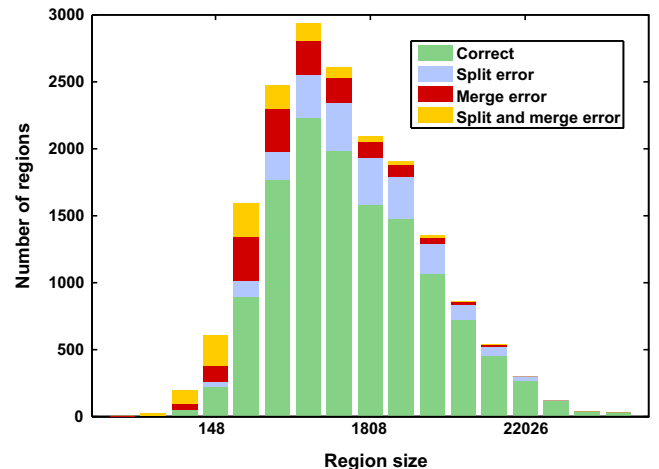


Fig. 9. Histogram of error rates for different region sizes. While over 70% of the regions are correctly segmented for most region sizes, smaller regions tend to be merged, whereas larger regions tend to be split. Overall large regions exhibit smaller error rates than small regions.



Fig. 10. Example region of an automatically segmented image (middle) compared to manual annotation (left). The original EM image is shown on the right. All images are 3 μm wide.

example segmentation of a 2D section compared to a manual annotation. While most regions are correctly segmented, some are split into several parts and need manual merging, while other regions span multiple objects and need to be manually split.

In order to minimize the user effort required to correct split and merge errors, we developed an interactive system called Mojo (see Fig. 11). The proofreading workflow in Mojo is as follows: The user is presented with a 2D view of a 3D EM image stack that allows zooming, panning, and scrolling through the out-of-plane dimension. Mojo presents the user with an interactive color overlay of the automatic reconstruction. Mojo provides two distinct interaction modes for correcting split errors and merge errors, respectively.

Correcting a split error requires the user to merge objects. To correct a split error in Mojo, the user clicks on a *source* object, and then clicks on one or more *target* objects. Mojo is responsible for re-labeling the target objects with the label from the source object.

Correcting a merge error requires the user to split an object into multiple sub-objects. The user begins by clicking on the object to be split. The user then roughly scribbles with a uniquely colored brush on each distinct sub-object within the object to be split. We use the interactive segmentation method of Roberts et al. (2011) to segment each sub-object. We chose to implement this segmentation method in Mojo because it produces highly accurate segmentations with minimal user effort, and provides the user with real-time feedback on the resulting segmentation while the user is scribbling. During each split operation, we constrain the

user scribbles and the resulting sub-objects to be entirely contained within the original object to be split. This allows the user to more easily segment each sub-object without disturbing neighboring objects.

7. Parallel implementation

Our pipeline has been designed to efficiently scale to large data sets in the GB-TB range. Scalability is an important aspect in the context of Connectomics. While a resolution of 5 nm is essential to allow for identification of biological structures like vesicles or synapses, whole neuronal cells expand over several μm of brain tissue. In the following sections we describe the run-time performance and scalability of the current implementation. We use the Harvard Research Computing cluster Odyssey for our computations. A typical node contains 2×4 – core Intel Xeon E5410 2.3 Ghz and 32 GB RAM. All evaluations concerning run time are given with respect to the current MATLAB implementation and include computational and i/o overhead to facilitate restarting of jobs on a computer cluster. The 2D segmentations are computed on 1024×1024 image tiles with a runtime of 9–13 h per job. For segmentation fusion we use $512 \times 512 \times 64$ sub-volumes and the average runtime is 2–3 h per job. In the pairwise matching step, sub-volumes are joined to form a consistently labeled volume. Every pair of adjacent sub-volumes (in x , y and z directions) is considered independently and in parallel. Winning groups of segments from the fusion step are merged or split based on the proportion of overlapping voxels inside the overlap region. Merge operations link segments to form a single object. Split operations assign all 2D segments in the overlap to just one group, creating two non-overlapping objects. This ensures a consistent labeling of voxels between adjacent sub-volumes. Runtime for each pair of sub-volumes is about 5–6 min. Finally, a single global consistency step is required to link objects over multiple sub-volume pairs and assign unique IDs. This step is performed by a single job and takes 1–2 h due to the amount of i/o operations required.

8. Large-scale reconstruction results

We successfully used our pipeline to reconstruct neuronal processes in a 27,000 μm^3 volume of brain tissue. This volume is more than 150 times larger than the manually annotated volume we used for quantitative evaluations, and it would take about 13,500 h to segment this volume manually, rendering a full quantitative evaluation of the large volume infeasible. To address this challenge and still provide a quantitative measure for the quality of our automatic reconstruction we measure the number of processes that enter and exit the volume correctly. Neuronal processes are typically longer than 30 μm and it is unlikely that a process ends inside the volume, thus it is desirable to have many processes that are correctly traced from one face to the other. This evaluation measure is challenging for large volumes, as the chance of introducing an error to the reconstruction of an object grows exponentially with the object length.

The following reconstruction results were obtained automatically, without any manual proofreading.

Fig. 12 shows a subset of the reconstructed processes. The visualization in Fig. 12 only includes processes that are traced from one face of the volume to another and that do not show obvious errors in the reconstructed 3D geometry. In total 93 objects satisfied these criteria. Note that although our data set is anisotropic, the reconstruction contains processes that run orthogonal as well as horizontal across the volume. The annotations mark an example of a correctly reconstructed branching structure (A) as well as several reconstructed spine necks (B).

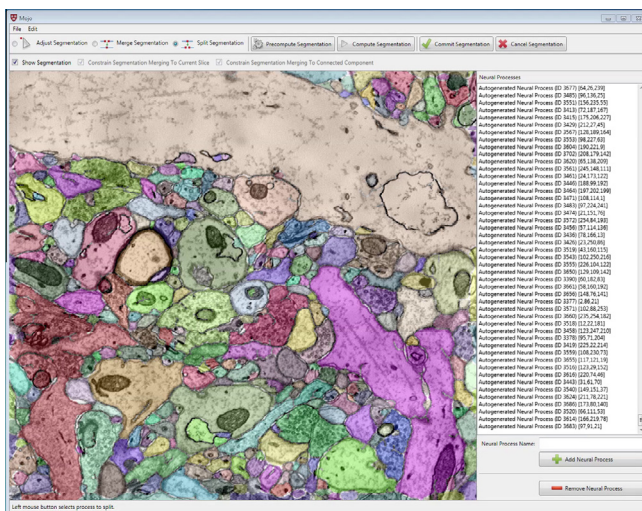


Fig. 11. Screenshot of our proofreading tool Mojo. The main area shows the segmentation as color-overlay on the original electron microscopy image. On the right side segmented processes can be selected by their name or ID. In the top menu the user can select split or merge error correction and if the correction should be applied in 3D or restricted to the currently shown section. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 12. Automatic large-scale reconstruction of 93 objects from the whole data volume without any manual proofreading. The reconstructions contain long-range tracings across the whole volume running orthogonal across all 1000 sections as well as longitudinal to the cutting plane. A correctly reconstructed branching structure is marked with (A). (B) marks automatically reconstructed spine necks of approximately 30 nm in diameter. The reconstructed volume size is $30.72 \mu\text{m} \times 30.72 \mu\text{m} \times 30 \mu\text{m}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Spines are important parts of neuronal processes in mammalian tissue. A spine typically consists of a thin spine neck and a thicker spine head at the end. Spine heads normally end in a synapse forming a connection to another neuronal process. Spine necks form the thinnest parts of a neuronal processes and can have a diameter of only 25 nm. Thus they can be hard to distinguish from extracellular space between cells and sometimes are also missed by human expert annotators. Fig. 13 shows an example image with annotations for a spine neck region and several regions corresponding to extracellular space. The small diameter of spine necks renders their automatic 3D reconstruction challenging. Differentiation of spine neck regions and extracellular space is often only possible by taking the broader 3D context into account. To gain more insight into the quality of our spine reconstructions we used a manually annotated part of a dendrite and cut out the corresponding area from the automatic reconstruction. Fig. 14 depicts the result. While none of the spines are correctly reconstructed




Fig. 13. Example image with an annotated spine neck (green) and multiple annotated extracellular space regions (purple). The small diameter of spine necks makes the automatic 3D reconstruction challenging. The width of the image is 1.2 μm . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 14. Fragmentation of a automatically reconstructed dendrite with respect to a manual reconstruction. The different colors correspond to different objects in the automatic segmentations. While manual proofreading is necessary to obtain the correct reconstruction, the automatic result includes spine necks traced over several sections. (A) and (B) only require merging of the identified spine neck with the spine head, whereas (C) and (D) exhibit further split errors, but still contain large continuous segments. (E) is an example of a fragmented spine neck, running longitudinal to the cutting plane. The bounding box of the shown dendrite is about $3 \mu\text{m} \times 6 \mu\text{m} \times 6 \mu\text{m}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 15. Examples of large-scale merge errors. The large red cell body is merged with neuronal processes marked (A), (B), and (C). The green object contains two branching structures, which are erroneously merged at location (D). The height of red cell is about 30 μm . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

automatically, proof reading the fragmentation of spines (A) and (B) is reduced to correcting one split error, merging the spine neck with the spine head. The spines (C) and (D) are more fragmented than (A) and (B), but contain tracings over several sections. (E) is




Fig. 16. Example segmentation of a region of interest from the reconstructed EM volume. The image shows an overlay of the segmentation in color on the gray value EM image. This result is prior to manual proofreading. The image width corresponds to 15.18 μm . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

heavily fragmented as it runs longitudinal across the sections and thus is harder to trace than orthogonal spine necks.

Region grouping with branching is essential not only to account for branching neuronal processes, but also to reconstruct spine necks along a dendrite. However, it can also lead to merge errors in the resulting segmentation. Fig. 15 shows examples of long-range merge errors. The reconstruction shown contains a correctly segmented cell body (red) including correct branches. The neuronal processes marked (A), (B), and (C) should be separate objects. The green structure is a merge of two neuronal processes. While both processes contain a correctly identified branching point, they are erroneously merged at location (D). These long-range merge errors are easy to detect for a human proof reader by looking at the 3D geometry of the reconstructed objects. Automatic identification and correction of the long-range 3D geometry is part of our future research.

Fig. 16 depicts an example segmentation of 2D image from our reconstructed volume.

9. Conclusions

In this paper we address the automatic reconstruction of neuronal processes at nm resolution for large-scale data sets. We demonstrate state-of-the art performance of our pipeline with respect to automatic dense reconstruction of neuronal tissue, and also for long range reconstructions covering neuronal processes over many μm . The workflow is designed to minimize manual effort and to be easy parallelizable on computer clusters and GPUs, with most steps scaling linearly with the number of processors. The electron microscopy data is available at <http://openconnecto.me/catmaid/?dataview=13#>, and the source code is available at <https://github.com/Rhoana>.

Future work concentrates on improving the performance of our pipeline, as well as facilitating the proofreading further. We are currently focusing our efforts on improving the runtime of the pipeline by optimizing code and removing MATLAB dependencies. We are also working on new algorithms to improve the overall

segmentation performance. With respect to proofreading, we are working on a web-based Mojo version that allows for collaborative proof reading of large data volumes.

Acknowledgment

The authors would like to thank Daniel Berger for providing manual segmentation, and Nancy Aulet for interactively annotating the training data. We also would like to thank Jan Funke for providing the Sopnet results, and Bjoern Andres for providing an efficient implementation of variation of information. This work has been partially supported by NSF Grants PHY 0938178, OIA 1125087, NIH Grant 2R44MH088088-03, Transformative R01 NS076467, the Gatsby Charitable Trust, Nvidia, Google, and the Intel Science and Technology Center for Visual Computing. The computations in this paper were run on the Odyssey cluster supported by the FAS Science Division Research Computing Group at Harvard University.

References

- Andres, B., Koethe, U., Helmstaedter, M., Denk, W., Hamprecht, F.A., 2008. Segmentation of sbfsem volume data of neural tissue by hierarchical classification. *Pattern Recognition*, vol. D. Springer, pp. 142–152.
- Andres, B., Koethe, U., Kroeger, T., Helmstaedter, M., Briggman, K.L., Denk, W., Hamprecht, F.A., 2012a. 3d segmentation of sbfsem images of neuropil by a graphical model over supervoxel boundaries. *Med. Image Anal.* 16 (4), 796–805.
- Andres, B., Kröger, T., Briggman, K.L., Denk, W., Korogod, N., Knott, G., Köthe, U., Hamprecht, F.A., 2012b. Globally optimal closed-surface segmentation for connectomics. In: *ECCV*, pp. 778–791.
- Arbeláez, P., Maire, M., Fowlkes, C., Malik, J., 2011. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.* 33 (5), 898–916.
- Bock, D., Lee, W., Kerlin, A., Andermann, M., Hood, G., Wetzell, A., Yurgenson, S., Soucy, E., Kim, H., Reid, R., 2011. Network anatomy and in vivo physiology of visual cortical neurons. *Nature* 471 (7337), 177–182.
- Boykov, Y., Funka-Lea, G., 2006. Graph cuts and efficient n-d image segmentation. *Int. J. Comput. Vis.* 70 (2), 109–131.
- Boykov, Y., Kolmogorov, V., 2004. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal.* 26 (9), 1124–1137.
- Breiman, L., 2001. Random forests. *Mach. Learn.* 45 (1), 5–32.

- Briggman, K.L., Bock, D.D., 2012. Volume electron microscopy for neuronal circuit reconstruction. *Current Opinion Neurobiol.* 22 (1), 154–161.
- Cardona, A., Saalfeld, S., Preibisch, S., Schmid, B., Cheng, A., Pulokas, J., Tomancak, P., Hartenstein, V., 2010. An integrated micro- and macroarchitectural analysis of the drosophila brain by computer-assisted serial section electron microscopy. *PLoS Biol.* 8 (10).
- Chen, C., Liaw, A., Breiman, L., July 2004. Using random forest to learn imbalanced data. Technical report, Department of Statistics, University of California, Berkeley.
- Chklovskii, D.B., Vitaladevuni, S., Scheffer, L.K., 2010. Semi-automated reconstruction of neural circuits using electron microscopy. *Current Opinion Neurobiol.* 20 (5), 667–675.
- daCosta, N.M., Martin, K.A.C., 2011. How thalamus connects to spiny stellate cells in the cat's visual cortex. *J. Neurosci. Official J. Soc. Neurosci.* 31 (8), 2925–2937.
- Denk, W., Horstmann, H., 2004. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol.* 2 (11), e329.
- Denk, W., Briggman, K.L., Helmstaedter, M., 2012. Structural neurobiology: missing link to a mechanistic understanding of neural computation. *Nature Rev. Neurosci.* 13, 351–358.
- Fiala, J.C., 2005. Reconstruct: a free editor for serial section microscopy. *J. Microsc.* 218 (Pt 1), 52–61.
- Funke, J., Andres, B., Hamprecht, F.A., 2012. Efficient automatic 3D-reconstruction of branching neurons from EM data. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1004–1011.
- Giuly, R.J., Martone, M.E., Ellisman, M.H., 2012. Method: automatic segmentation of mitochondria utilizing patch classification, contour pair classification, and automatically seeded level sets. *BMC Bioinform.* 13 (1), 29.
- Hayworth, K.J., Kasthuri, N., Schalek, R., Lichtman, J.W., 2006. Automating the collection of ultrathin serial sections for large volume tem reconstructions. *Microsc. Microanal.* 12 (S02), 86–87.
- Helmstaedter, M., Mitra, P.P., 2012. Computational methods and challenges for large-scale circuit mapping. *Current Opinion Neurobiol.* 22 (1), 162–169.
- Helmstaedter, M., Briggman, K.L., Denk, W., 2011. High-accuracy neurite reconstruction for high-throughput neuroanatomy. *Nature Neurosci.* 14 (8), 1081–1088.
- Hu, T., Nunez-Iglesias, J., Vitaladevuni, S., Scheffer, L., Xu, S., Bolorizadeh, M., Hess, H., Fetter, R., Chklovskii, D., 2013. Electron microscopy reconstruction of brain structure using sparse representations over learned dictionaries. *IEEE Trans. Med. Imaging* 32 (12), 2179–2188.
- IBM, 2014. Ilog cplex optimizer. <<http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>>.
- Jain, V., Murray, J.F., Roth, F., Turaga, S., Zhigulin, V., Briggman, K.L., Helmstaedter, M.N., Denk, W., Seung, H.S., 2007. Supervised learning of image restoration with convolutional networks. In: 2007 IEEE 11th International Conference on Computer Vision. IEEE, pp. 1–8.
- Jurrus, E., Whitaker, R., Jones, B.W., Marc, R., Tasdizen, T., May 2008. An optimal-path approach for neural circuit reconstruction. In: Proceedings/IEEE International Symposium on Biomedical Imaging: from nano to macro 2008 (4541320), pp. 1609–1612.
- Jurrus, E., Paiva, A.R.C., Watanabe, S., Anderson, J.R., Jones, W., Whitaker, R.T., Jorgensen, E.M., Marc, R.E., 2010. Detection of neuron membranes in electron microscopy images using a serial neural network architecture. *Med. Image Anal.* 14 (6), 770–783.
- Kaynig, V., Fuchs, T., Buhmann, J.M., 2010a. Neuron geometry extraction by perceptual grouping in sstem images. In: IEEE Conference on Computer Vision and Pattern Recognition 0, pp. 2902–2909.
- Kaynig, V., Fuchs, T.J., Buhmann, J.M., 2010b. Geometrical consistent 3d tracing of neuronal processes in sstem data. *Med. Image Comput. Comput.-Assist. Intervent.* 13 (Pt 2), 209–216.
- Knott, G., Marchman, H., Wall, D., Lich, B., 2008. Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. *J. Neurosci.* 28 (12), 2959–2964.
- Knowles-Barley, S., Butcher, N.J., Meinertzhagen, I.A., Armstrong, J.D., 2011. Biologically inspired em image alignment and neural reconstruction. *Bioinformatics* 27 (16), 2216–2223.
- Kolmogorov, V., Zabin, R., 2004. What energy functions can be minimized via graph cuts? *IEEE Trans. Pattern Anal.* 26 (2), 147–159.
- Kuwajima, M., Mendenhall, J.M., Lindsey, L.F., Harris, K.M., 2013a. Automated transmission-mode scanning electron microscopy (tSEM) for large volume analysis at nanoscale resolution. *PloS One* 8 (3), e59573, <<http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3608656&tool=pmcentrez&rendertype=abstract>>.
- Kuwajima, M., Spacek, J., Harris, K.M., 2013b. Beyond counts and shapes: studying pathology of dendritic spines in the context of the surrounding neuropil through serial section electron microscopy. *Neuroscience* 251, 75–89.
- Lee, W.-C.A., Reid, R.C., 2011. Specificity and randomness: structure-function relationships in neural circuits. *Current Opinion Neurobiol.* 21 (5), 801–807.
- Lichtman, J.W., Sanes, J.R., 2008. Ome sweet ome: what can the genome tell us about the connectome? *Current Opinion Neurobiol.* 18 (3), 346–353.
- Lucchi, A., Smith, K., Achanta, R., Knott, G., Fua, P., 2012. Supervoxel-based segmentation of mitochondria in em image stacks with learned shape features. *IEEE Trans. Med. Imaging* 31 (2), 474–486.
- Marc, R.E., Jones, B.W., Watt, C.B., Anderson, J.R., Sigulinsky, C., Lauritzen, S., 2013. Retinal connectomics: towards complete, accurate networks. *Prog. Retinal Eye Res.* 37, 141–162.
- Meila, M., 2007. Comparing clusterings – an information based distance. *J. Multivariate Anal.* 98 (5), 873–895.
- Mishchenko, Y., 2009. Automation of 3d reconstruction of neural tissue from large volume of conventional serial section transmission electron micrographs. *J. Neurosci. Methods* 176, 276–289.
- Nunez-Iglesias, J., Kennedyand, R., Paragand, T., Shi, J., Chklovskii, D.B., 2013. Machine learning of hierarchical clustering to segment 2d and 3d images. *PLoS ONE* 8 (8), e71715.
- Penzes, P., Cahill, M.E., Jones, K.A., VanLeeuwen, J.-E., Woolfrey, K.M., 2011. Dendritic spine pathology in neuropsychiatric disorders. *Nature Neurosci.* 14 (3), 285–293, <<http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3530413&tool=pmcentrez&rendertype=abstract>>.
- Roberts, M., Jeong, W.-K., Vázquez-Reina, A., Unger, M., Bischof, H., Lichtman, J., Pfister, H., 2011. Neural process reconstruction from sparse user scribbles. In: Medical Image Computing and Computer Assisted Intervention (MICCAI '11), pp. 621–628.
- Saalfeld, S., Fetter, R., Cardona, A., Tomancak, P., 2012. Elastic volume reconstruction from series of ultra-thin microscopy sections. *Nature Methods* 9 (7), 717–720.
- Seung, H.S., 2009. Reading the book of memory: sparse sampling versus dense mapping of connectomes. *Neuron* 62 (1), 17–29.
- Seung, S., 2012. Connectome: How the Brain's Wiring Makes Us Who We Are. Houghton, Mifflin Harcourt.
- Sommer, C., Straehle, C., Ullrich, K., Hamprecht, F.A., 2011. Ilastik: interactive learning and segmentation toolkit. *IEEE Int. Symp. Biomed. Imaging*, 230–233.
- Straehle, C.N., Köthe, U., Knott, G., Hamprecht, F.A., 2011. Carving: scalable interactive segmentation of neural volume electron microscopy images. *Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* 14 (Pt 1), 653–660.
- Turaga, S.C., Murray, J.F., Jain, V., Roth, F., Helmstaedter, M., Briggman, K., Denk, W., Seung, H.S., 2010. Convolutional networks can learn to generate affinity graphs for image segmentation. *Neural Comput.* 22 (2), 511–538.
- Valiant, L.G., 2006. A quantitative theory of neural computation. *Biol. Cybernet.* 95 (3), 205–211.
- Vazquez-Reina, A., Miller, E., Pfister, H., 2009. Multiphase geometric couplings for the segmentation of neural processes. In: Conference on Computer Vision and Pattern Recognition. IEEE, pp. 2020–2027.
- Vazquez-Reina, A., Huang, D., Gelbart, M., Lichtman, J., Miller, E., Pfister, H., 2011. Segmentation fusion for connectomics. *Int. Conf. Comput. Vis.*, 177–184.
- Vitaladevuni, S.N., Basri, R., 2010. Co-clustering of image segments using convex optimization applied to em neuronal reconstruction. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. IEEE, pp. 2203–2210.
- Wang, C.-W., Ka, S.-M., Chen, A., January 2014. Robust image registration of biological microscopic images. *Scientific reports* 4, 6050.