Published in final edited form as:

Neuroinformatics. 2013 January ; 11(1): 5–29. doi:10.1007/s12021-012-9149-y.

Semi-automated Neuron Boundary Detection and Nonbranching Process Segmentation in Electron Microscopy Images

Elizabeth Jurrus^{*,1,2}, Shigeki Watanabe³, Richard J. Giuly⁵, Antonio R. C. Paiva¹, Mark H. Ellisman⁵, Erik M. Jorgensen³, and Tolga Tasdizen^{4,1}

¹Scientific Computing and Imaging Institute, University of Utah

²School of Computing, University of Utah

³Department of Biology, University of Utah

⁴Department of Electrical Engineering, University of Utah

⁵National Center for Microscopy and Imaging Research, University of California, San Diego

Abstract

Neuroscientists are developing new imaging techniques and generating large volumes of data in an effort to understand the complex structure of the nervous system. The complexity and size of this data makes human interpretation a labor-intensive task. To aid in the analysis, new segmentation techniques for identifying neurons in these feature rich datasets are required. This paper presents a method for neuron boundary detection and nonbranching process segmentation in electron microscopy images and visualizing them in three dimensions. It combines both automated segmentation techniques with a graphical user interface for correction of mistakes in the automated process. The automated process first uses machine learning and image processing techniques to identify neuron membranes that delineate the cells in each two-dimensional section. To segment nonbranching processes, the cell regions in each two-dimensional section are connected in 3D using correlation of regions between sections. The combination of this method with a graphical user interface specially designed for this purpose, enables users to quickly segment cellular processes in large volumes.

1 Introduction

Neural circuit reconstruction is an important method for studying neural circuit connectivity and its behavioral implications. The differences between neuronal classes, patterns, and connections are central to the study of the nervous system and critical for understanding how neural circuits process information. The ability to reconstruct neural circuitry at ultrastructural resolution is also of great clinical importance. With each new dataset generated, new details of well-known brain areas are being revealed, promising new insights into the basic biology and disease processes of nervous systems. For instance, for the first time, scientists can study the structural integrity of the transition zone of the optic nerve from unmyelinated to myelinated in the nervous system. This transition zone is now identified as one of the sites of pathology in glaucoma (Gonzalez-Hernandez et al. 2008). Other retinal degenerative diseases, including retinitis pigmentosa and macular degeneration, result from a loss of photoreceptors. Photoreceptor cell stress and death induces subsequent changes in the neural circuitry of the retina resulting in corruption of the surviving retinal cell class circuitry. Ultrastructural examination of the cell identity and

* Corresponding author. 72 S Central Campus Drive, Salt Lake City, UT 84112 USA. liz@sci.utah.edu.

circuitry reveal substantial changes to retinal circuitry with implications for vision rescue strategies (Marc et al. 2008, 2007, 2003; Jones and Marc 2005; Jones et al. 2003; Jones, Watt, and Marc 2005; Peng et al. 2000). Ultrastructural evaluation of multiple canonical volumes of neural tissue are also critical towards the evaluation of differences in connectivity between different individuals.

Electron microscopy (EM) is a useful method for determining the anatomy of individual neurons and their connectivity because it has a resolution that is high enough to identify features, such as synaptic contacts and gap junctions. These features define connectivity, and therefore are required for neural circuit reconstruction. Manual analysis of this data is extremely time-consuming. Early work in mapping the complete nervous system of the relatively simple *C. elegans* took many years (White et al. 1986). Since then, several researchers have undertaken extensive EM imaging projects in order to create detailed maps of neuronal structure and connectivity (Fiala and Harris 2001; Briggman and Denk 2006a; Varshney et al. 2011). In comparison, newer imaging techniques are producing much larger volumes of very complex organisms, with thousands of neurons and millions of synapses (Briggman and Denk 2006b; Anderson et al. 2009). The complexity and size of these datasets, often approaching tens of terabytes, makes human segmentation of the complex textural information of electron microscopic imagery both a difficult and very time-consuming task. Moreover, population or screening studies are unfeasible since fully manual segmentation and analysis would require years of manual effort per specimen. As a result, research in new imaging techniques and protocols, as well as automation of the reconstruction process, are critical for the study of these systems.

To assist in neural circuit reconstruction, this paper presents a method for segmenting 3D nonbranching cellular processes in EM images and visualize the results. The segmentation of neurons combines both automated neuron segmentation techniques with a graphical user interface for correction of mistakes in the automated process. The automated process first uses machine learning and image processing techniques to segment the neurons in each 2D section and then connect them in 3D. The combination of this process with a graphical user interface specially designed for this purpose, enable users to quickly segment neuron cell processes in large volumes.

1.1 Imaging Methods

Serial-section Transmission Electron Microscopy (ssTEM) and Serial Block Face Scanning Electron Microscopy (SSBFSEM) are the two methods used for image acquisition in this paper. Compared with other state of the art methods, such as MRI (Xiao, Wang, and Felleman 2003) and scanning confocal light microscopy (Minsky 1961; Denk, Strickler, and Webb 1990; Egner and Hell 2005; Rust, Bates, and Zhuang 2006; Betzig et al. 2006), electron microscopy methods provide much higher resolution and remain the primary tool for resolving the 3D structure and connectivity of neurons.

One of the modalities chosen for reconstructing neuronal circuits at the individual cell level is serial-section transmission electron microscopy (ssTEM) (Anderson et al. 2009, 2011; Chklovskii, Vitaladevuni, and Scheffer 2010). Most importantly, through mosaicking of many individual images (Tasdizen et al. 2010; Saalfeld et al. 2010), ssTEM offers a relatively wide field of view to identify large sets of cells that may wander significantly as they progress through the sections. It also has an in-plane resolution that is high enough for identifying synapses. In collecting images through ssTEM, sections are cut from a specimen and suspended so that an electron beam can pass through it, creating a projection. The projection can be captured on a piece of film and scanned or captured directly as a digital image. An example ssTEM image is shown in Figure 2(a). An important trade-off occurs with respect to the section thickness. Thinner sections are preferable from an image analysis

point of view because structures are more easily identifiable due to less averaging. However, from an acquisition point of view, thinner sections are harder to physically handle and impose a limit on the area of the section that can be cut. Sections can be reliably cut at 30 – 90nm thickness with the current ssTEM technology. This leads to an extremely anisotropic z resolution, compared to 2 – 10nm in-plane. The *C. elegans* ventral nerve cord dataset used in this paper, for example, was imaged using ssTEM and has a resolution of $6\text{nm} \times 6\text{nm} \times 33\text{nm}$. This anisotropy poses two image processing challenges. First, the appearance of cell membranes can range from solid dark curves for neurons that run approximately perpendicular to the cutting-plane, to fuzzy grey swaths, commonly referred to as “grazed membranes,” for membranes that run more obliquely and suffer more from the averaging effect. This is demonstrated in Figure 1. Consequently, segmentations of neurons in these 2D images are difficult given the change in membrane contrast and thickness. Second, due to the large physical separation between sections, shapes and positions of neurons can change significantly between adjacent sections. An example of this is shown in Figure 3(a).

Another specimen preparation and EM imaging technique commonly used for neural circuit reconstruction is Serial-Block Face Scanning Electron Microscopy (SBF-SEM) (Denk and Horstmann 2004). In SBFSEM, successive slices are cut away and discarded, and the electron beam is scanned over the remaining block face to produce electron backscattering images. This method results in smaller deformations than ssTEM because the dimensions of the solid block remain relatively stable after slicing and thus deformation between sections is relatively small, usually eliminating the need for image registration between sections. The in-slice resolution (which is closer to 10nm) and signal-to-noise properties of SBFSEM are generally not as good as those of ssTEM, though. However, a specialized scanning electron microscope equipped with a high precision Gatan 3View ultramicrotome combined with an improved specimen staining protocol can produce high contrast images and increased detail of individual cells in the context of their surroundings (Deerinck et al. 2010). Specifically, by staining the tissue with a series of heavy metal stains, we were able to improve contrast and render the samples more conductive. The specimens were conductive enough to allow us to image at high vacuum, which results in images with improved resolution and signal/noise. This specimen preparation protocol was used to collect the neuropil of the molecular layer of the cerebellar cortex from an adult mouse, an exemplary image of which is shown in Figure 2(b). This image acquisition technique still results in anisotropic resolution, causing the separation between slices to be significant enough that positions of fine neurites and subcellular structures can shift and change significantly between sections (see Figure 3(b)). In the case of SBFSEM mouse neuropil dataset, the resolution is $10\text{nm} \times 10\text{nm} \times 50\text{nm}$ (SBFSEM).

1.2 Cellular Segmentation

There are two general approaches for neuron segmentation. One approach focuses first on the detection of neuron membranes in each 2D section. These boundaries can then be used to identify individual neurons, which are subsequently linked across sections to form a complete neuron (Jeong et al. 2010; Jurrus et al. 2008; Macke et al. 2008; Allen and Levinthal 1990). The other approach to neuron segmentation is to directly use the 3D characteristics of the data (Andres et al. 2008; Jain et al. 2007). Full 3D approaches are difficult due to the anisotropic nature of the data, however. As mentioned earlier, the large section thickness often causes features to shift significantly between sequential images, decreasing the potential advantages of a direct 3D approach.

A number of methods exist for neuron membrane detection in 2D sections. Image processing methods for finding membranes include edge detection (i.e., Canny), region growing methods, and intensity thresholding on enhanced membrane features, either through

Hessianbased diffusion (Tasdizen et al. 2005) or radon-like features (Kumar, Reina, and Pfister 2010). Accurate detection of neuron membranes using these methods alone is a difficult problem given the presence of intracellular structures. There are several methods that attempt to segment EM images of neural tissue using active contours, in both parametric and level set forms (Jurrus et al. 2009; Bertalmío, Sapiro, and Randall 2000; Vazquez, Sapiro, and Randall 1998; Vazquez-Reina, Miller, and Pfister 2009). 2D graph cuts can be used to segment images using region and boundary terms that separate intracellular structures from membranes (Vu and Manjunath 2008; Yang and Choe 2009). These can provide smooth, accurate segmentations of cells. However, when used alone they require a very specific energy minimization function so that neuron membranes are not confused with organelles making this method dependent on the type of cell being segmented. Additionally, their success can depend on their initialization (Vu and Manjunath 2008). Combined with machine learning methods (Kaynig, Fuchs, and Buhmann 2010), they have an improved detection accuracy and can be used more reliably in 3D.

Recent related work indicates that supervised machine learning methods are an effective approach for detection of neuron membranes in 2D and 3D (Jain, Seung, and Turaga 2010; Andres et al. 2008). Simple classifiers such as a single perceptron applied to a carefully chosen set of features have been shown to provide promising results in identifying membranes in EM images (Mishchenko 2008). Nevertheless, this method still needs significant post-processing to connect membranes and remove internal cellular structures. Similarly, Venkataraju *et al.* proposed using local context features computed from the Hessian matrix to train a boosted classifier to detect membranes, which highlights the importance of context for membrane detection (Venkatataju et al. 2009). Jain *et al.* use a multilayer convolutional ANN to classify pixels as membrane or non-membrane in specimens prepared with an extracellular stain (Jain et al. 2007; Turaga et al. 2009). The convolutional ANN has two important characteristics: it learns the filters for classification directly from data, and the multiple convolutions throughout the layers of the network account for an increasing (indirect) filter support region. The serial neural network architecture (Jurrus et al. 2010) used in this paper also takes advantage of context and samples the image pixels directly to learn membrane boundaries, but given the anisotropic data, focuses only on 2D sections. New cost functions used during training are being developed to take into account the topological constraints of neuron boundaries (Turaga et al. 2010; Jain et al. 2010). The results obtained with these methods demonstrate not only the complexity of the problem, but also the potential of supervised machine learning as a tool towards neuron segmentation.

One of the goals of this work is to combine machine learning and segmentation algorithms with three-dimensional rendering capabilities for users to better understand how to process the data and visualize the results. Towards this aim, there are several existing software efforts that incorporate many of the above algorithms specifically for reconstructing neural circuits from biological volumetric images. These tools provide an interface to the data and contour tools to segment structures in a stack of EM images. One of most widely used software tools is Reconstruct (Fiala, and Harris 2002; Fiala and Harris 2010) which enables users to view and outline structures of interest and then render them as 3D volumes. Combining Reconstruct with automated methods, such as the ones proposed by Mishchenko (Mishchenko et al. 2010) resulted in scientific discoveries regarding the predicted location of synapses within a neuron. IMOD (Kremer, Mastronarde, and McIntosh 1996) and TrakEM2 (Cardona et al. 2010) have also proved to be useful tools for mosaicking, segmenting, and rendering structures from a variety of biological volumetric image data. A more comprehensive set of tools for reconstruction is the Cell Centered Database (Martone et al. 2008) which performs not only annotation on EM images, but also provides data management and protein knowledge base interfaces. While these tools are critical in the

segmentation and reconstruction of EM data, the goal of this paper is to segment many structures from large sets of images through the design of automated memory efficient algorithms while also building a tool capable of streaming large datasets for volume viewing and interaction. Towards this goal, two software programs, the Serial Section Reconstruction and Tracing Tool (SSECRET) and NeuroTrace can segment large image databases (Jeong et al. 2009, 2010). SSECRET is an interface for slice-based viewing of large volumes using a client-server architecture to request only the data needed by the user. NeuroTrace incorporates 2D level set segmentation tools for segmenting individual sections, and then using those segmentations to identify long neuronal processes. These combined tools produce vital reconstruction data, however the interface to the data is specific to the implemented algorithms and still requires the user to initialize each neuron for segmentation. The software program designed for this paper similarly manages memory for large datasets, but also is designed to incorporate automated segmentation algorithms. In addition, it provides an interface specific to the segmentation method presented in this paper to make corrections, and most importantly, view the raw image data with its 3D segmentation.

2 Methods

The overall method proposed in this paper for reconstructing nonbranching neuron cell processes consists of two steps. First, neuron membranes are segmented in 2D and neuronal cross-sections are identified. Second, the regions are linked across all the sections to form 3D renderings of parallel processes. The initial neuron segmentation used for each 2D section builds upon previous work which uses a series of artificial neural networks (ANNs) to detect neuron membranes. To improve the membrane detection, that method is extended here by incorporating learned membranes from sequential sections into another ANN and applying tensor voting post-processing. Also drawing from previous work, we incorporate an optimal path algorithm to connect similar regions through the volume to form complete 3D segmentations. Furthermore, this paper combines all of the above techniques into an interactive tool, called the Neuron Reconstruction Viewer (NeRV), that lets the user view large EM datasets, evaluate the segmentations, and make corrections to both the 2D membrane detection and the joining of regions through the sections to segment a neuron in 3D.

2.1 2D Membrane Detection

The method developed here for neuron membrane detection extends previous work, which uses a series of ANN classifiers and image stencil neighborhood feature vectors to detect neuron membranes in 2D images (Jurrus et al. 2010). In that paper, membrane detection was limited to features within a 2D section. An example output from this algorithm is shown in Figure 4(b)–(d). This work is extended here to train on information from neighboring sections, using the confidence from sequential sections. Given the anisotropic nature of the data, sequential sections have very poor membrane correspondence. To account for this, classified results representing the membrane probability image are registered and a 3D stencil that spans 3 sections is formed for training. Finally, tensor voting, a method for closing remaining gaps, is used. This provides significantly improved segmentation results over the original method (Jurrus et al. 2010). The output from these additional steps is shown in Figure 4(e)–(f). Quantitatively, the improvement of these new methods can be seen in Figures 16, 17, 22, and 23.

2.1.1 Serial Neural Network Architecture—In previous work, a serial classifier architecture was implemented that used a series of classifiers, each operating on input from the previous classifier, to incrementally gain knowledge of a large neighborhood (Jurrus et al. 2010; Paiva, Jurrus, and Tasdizen 2010). This architecture is particularly useful for two

reasons. First, the data used for training requires no preprocessing with filter banks or statistics, and the classifier is trained directly on sampled image intensities. Second, by applying several classifiers in series, each classifier uses the classification context provided by the previous network to improve membrane detection accuracy. To initialize this architecture, the first classifier is trained only on image intensities. Each remaining classifier in the series then uses an input vector containing samples from the original image appended with the values from the output of the previous classifier, yielding a larger feature vector. While the desired output labels remain the same, each classifier is dependent on the information from the previous network and therefore must be trained sequentially. The output from each network is used to generate an image that represents the membrane probability map at that stage. Figure 5 demonstrates this flow of data between classifiers: ML is the classifier, I denotes the image, S represents the sampling of image intensities from the image using the stencil, and C denotes the output from the classifier, yielding the membrane detection.

Since the serial classifier architecture is not specific to any classifier and given the success of ANNs for membrane detection (Mishchenko 2008; Jain et al. 2007), the classifier chosen for this architecture is a multilayer perceptron (MLP) ANN (shown in Figure 6). An MLP is a feed-forward neural network which approximates a classification boundary with the use of nonlinearly weighted inputs. The output of each processing element (PE) (each node of the ANN) is given as (Haykin 1999; Principe, Euliano, and Lefebvre 2000)

$$y=f(\mathbf{w}^T \mathbf{x}+b), \quad (1)$$

where f is, in our case, the *tanh* nonlinearity, $\mathbf{w}$ is the weight vector, and b is the bias. The input vector $\mathbf{x}$ to PEs in the first hidden layer is the input feature vector discussed in more detail in the next section. For the PEs in subsequent layers, $\mathbf{x}$ contains the outputs of the PEs in the previous layer. ANNs are a method for learning general functions from examples. They are well suited for problems without prior knowledge of the function to be approximated (a.k.a., "black box models"). They have been successfully applied to robotics (Pomerleau 1993; Wells, Venaille, and Torras 1996) and face and speech recognition (Rabi and Lu 1998; Cottrell 1990), and are robust to noise.

To learn the weight vector and bias, back-propagation was used to minimize the minimum squared error (MSE) criterion (Haykin 1999; Principe, Euliano, and Lefebvre 2000). Back-propagation is a gradient descent procedure that maps the output layer error to the error at the output of each node, yielding a local update rule that depends only on the node's input and output error. It is obtained by direct application of the chain rule to the derivative of the criterion with regards to each one of the ANN parameters. Because back-propagation depends only on the local gradient and that information is available directly from the optimization criterion, there is no need to explicitly characterize the parameter space.

The serial classifier is trained by simply using raw image intensities. Training a classifier on raw image data yielded improved results over filter banks and neighborhood statistical information (Jurrus et al. 2009; Paiva, Jurrus, and Tasdizen 2010). The stencil, shown in Figure 7, can cover large areas representing the desired feature space, but samples it with a spatially adaptive resolution strategy. In this way, an ANN can be trained using a low dimensional feature vector without having to use the whole image patch. Pixels are selected close to the stencil center, along a radius, at a high resolution, and then further from the center at a more coarse resolution. This gives more detail for the training of our classifier around the feature of interest, while maintaining a large area in which to apply context. Since the number of weights to be computed in an ANN are dominated by the connection between the input and the hidden layers, reducing the number of inputs also reduces the

number of weights and helps regularize the learned network. Moreover, using fewer inputs generally allows for faster training. With this, one aims to provide the classifier with sparse, but sufficient context information and achieve faster training, while obtaining a larger context which can lead to improvements in membrane detection. This strategy, combined with the serial use of ANNs, grows the region of interest for classification within a smaller number of stages and without long training times. Example output from the serial classifier is shown in Figure 4(b)–(d).

The training data used to train the classifier is generated by hand by domain experts. For each dataset, a user annotated all the membrane boundaries with curves that were one pixel wide on a subset of the images. We dilated this boundary to cover the width of the neuron membrane and these pixels were used as positive training examples for our classifier. Unannotated pixels, which included intercellular features such as vesicles, mitochondria, and nucleus, were used as negative training examples. From these pixels, we randomly chose a balanced set of positive and negative training examples for our classifier. This is discussed in more detail for each dataset in Sections 3.1 and 3.2.

2.1.2 Sequential Section Serial Neural Network Architecture—Sequential sections from EM data often contain similar structures that we would like to use as context to improve the quality of the 2D segmentation. One way to do this would be with a stencil that spans multiple sections. However, the membrane locations between sections have poor correspondence. This is partly because of the anisotropic nature of the data, which often results in large movement of membranes between sections, and membranes sometimes do not run perpendicular to the cutting plane causing membranes to have low contrast and appear fuzzy. The differences between two sections is seen in Figure 8(a) which shows two sequential images with detected membranes overlaid with each other. Membranes in sequential sections are near each other, but they do not correspond well enough to use them directly in a 3D stencil that would span multiple sections. One way to resolve this problem is to perform a nonrigid registration across the whole volume to align as best as possible all the membrane boundaries. There are two problems with this approach. First, internal structures in the neurons complicate the registration process introducing possible errors to the segmentation. Second, this process introduces warping, changing the anatomy of the neurons. To account for this, we propose a novel approach which aligns sequential membrane probability map images between only two sections using a correlation-based nonlinear registration. We register only the membrane probability images because the classification process has removed many of the internal structures that would make an extremely fine-scale nonlinear registration on raw image data difficult. Also, we perform the registration between only two sections to keep the location of the neurons intact. Once registered, a 3D stencil that spans 3 adjacent sections samples the classification results from the previous stage and provides information to be used in the final classification step.

More specifically, after the membrane detection is complete for each section using the serial ANN architecture, images are registered in pairs to the center section and used as input to a new ANN. The serial ANN with the registration step and final ANN is depicted in Figure 9. The registration method proposed is a B-spline deformable registration (Ibanez et al. 2005). Given an image to be registered and a static template image, a nonlinear deformation can be generated which minimizes the mean squared difference energy, given by,

$$\int_{\Omega} (C^M \circ t(x) - C^S(x))^2 dx. \quad (2)$$

where Ω is the image domain and $t(x)$ is the deformation $\mathbb{R}^2 \rightarrow \mathbb{R}^2$, in our case given by a 2D tensor product B-spline transform of order 2 (Rueckert et al. 1999). C^M is the moving

classification image, and C^s is the static classification image. For our purposes, C_i (the center section) is the static image and C_{i-1} and C_{i+1} are the moving images.

Each section has its own set of neighboring registered sections. The change in the membrane locations after two images are registered is shown in Figure 8(b). Now that membranes are more carefully aligned across neighboring sections, a new stencil can be used to sample the 3D space. The 3D, three section, stencil is similar to the one shown previously in Figure 7. This stencil is used on the middle slice, while the stencil on the top and bottom slice have a shorter radius. The output from the ANN using this stencil is shown in Figure 4(e). Using information from the sequential sections, the ANN learns to identify membranes in C_i that were not previously detected, because the membranes were detected in C_{i-1} and C_{i+1} improving the overall segmentation. This helps specifically in cases where C_i contains grazed membranes, but C_{i-1} and C_{i+1} do not. A good example of this is shown in Figure 14, second column. Membranes in the raw image appear fuzzy and are not well detected after the serial ANN. Using information from the registered sequential sections strengthens these boundaries. In this way, the ANN also learns the possible shapes of membranes across several sections.

2.1.3 Tensor Voting—Tensor Voting (TV) is a method first proposed by Guy and Medioni (Medioni, Lee, and Tang 2000) for extraction or enhancement of local features (lines, curves or surface) extraction results. Local feature extraction by itself is often unreliable in noisy and complicated images. That is, the lines or curves are often noisy and interrupted. TV enhances or predicts local features by integrating clues from nearby features. In our case, the influence of nearby features is based on a given voting field designed to extract smooth curves.

In tensor voting, a 2-D tensor can be represented by a symmetric, positive semidefinite 2×2 matrix as follows:

$$\mathbf{T} = \begin{pmatrix} a_{xx} & a_{xy} \\ a_{xy} & a_{yy} \end{pmatrix} = \lambda_1 \mathbf{e}_1 \mathbf{e}_1^T + \lambda_2 \mathbf{e}_2 \mathbf{e}_2^T \quad (3)$$

where λ_1 and λ_2 are the eigenvalues ($\lambda_1 \geq \lambda_2 \geq 0$); $\mathbf{e}_1$ and $\mathbf{e}_2$ are the orthonormal eigenvectors. Graphical representation of this kind of tensor is ellipse, as shown in Figure 10(a). One common parameterization is to define a tensor with three parameters: orientation β , stickness s and ballness b .

$$\beta = \arccos(\mathbf{e}_x^T \mathbf{e}_1) \quad (4)$$

$$s = \lambda_1 - \lambda_2 \quad (5)$$

$$b = \lambda_2 \quad (6)$$

where $\mathbf{e}_x = (1, 0)^T$. A tensor with ballness equal to 0 is called stick tensor and represents a curve passing through that pixel. A tensor with stickness equal to 0 is called ball tensor. In our work the initial tensors are computed from the classifier output image as

$$\mathbf{T}_{x,y} = I_{x,y} \begin{pmatrix} \cos(\alpha_{x,y}) \\ \sin(\alpha_{x,y}) \end{pmatrix} (\cos(\alpha_{x,y}), \sin(\alpha_{x,y})) \quad (7)$$

where $I_{(x,y)}$ is the membrane strength at pixel (x,y) and $\alpha_{(x,y)}$ is the local orientation computed as the orientation of the eigenvector corresponding to the smaller eigenvalue of the local second order derivative matrix (Hessian).

In TV, the "vote" is a tensor calculated from the geometric relation between the voter and the votee. All the votes to a pixel will be summed and form the output tensor of that pixel. The voting field is shown in Figure 10(b). A typical voting method is to filter the tensor image with a nonholonomic filter, called voting field, aligned with the local tensor orientation. Since this is a computationally intensive method, it is typically applied in a sparse manner such that only a subset of the pixels in the image, those where curve features are detected, are allowed to cast votes. However, in problems where detection is hard, such as with neuron membranes, it can be advantageous to allow every pixel to cast votes proportional to their strength, as determined by our classifier, and to postpone detection until after this step. To achieve this in a computationally efficient manner, we used a rapid tensor voting algorithm which uses steerable filters as a basis for the voting field (Franken et al. 2006; Leng et al. 2011). In this algorithm, a set of basis voting fields are convolved with the image and then linearly combined to form the desired voting field at each pixel.

2.1.4 Region Segmentation—Given detected membranes in each 2D section, neurons can be segmented in each section using either a watershed segmentation (Gonzalez and Woods 1992; Ibanez et al. 2005) or a simple flood fill algorithm on the thresholded probability map. The flood fill algorithm operates on thresholded data and works best when a user has corrected segmentations with hand editing and wants a precise neuron membrane representation at every section. For larger problems, the watershed algorithm has the advantage in that it can close gaps automatically. For our method, we apply a watershed segmentation to the blurred output from the tensor voting and select the watershed depth that best segments the neuron regions. However, there are two trade-offs to consider when choosing to use the watershed for region segmentation. The first is a trade-off between the ability to close large gaps and the ability to segment smaller features. This is controlled by the parameter σ which is used to smooth the image as part of the watershed process. Large σ enables the watershed to close large gaps but also loses the ability to segment smaller features. Another trade-off of the watershed is that depending on the level (or depth) of the watershed, chosen by the user, over-segmentation can occur of regions, meaning areas that should be one whole region are instead two or three regions. The user has to balance these trade-offs when choosing the level to proper set of required parameters. This is discussed more in Section 3.2 when the watershed is applied to the mouse neuropil data.

2.2 Region Linking

In segmenting the structures relevant for the datasets described in this paper, we present a method that identifies only parallel processes through a stack of images. For this paper, neuron identification across a stack of EM images is formulated as an optimal path problem with a graph data structure (Jurrus et al. 2008). The vertices of the graph are defined as the regions obtained by 2D segmentation of the individual sections, as described in Section 2.1. Edges in the graph represent possible linkages between regions in neighboring sections. Linking together the neuron regions in the graph is performed using Dijkstra's shortest path algorithm. The resulting path through the graph is used to reconstruct the neuron in 3D.

2.2.1 Linking Method for Neuron Regions—Let $R_{s,i}$ be the i 'th region from the 2D segmentation in section s . A directed graph containing a set of nodes that correspond to the set of segmented regions in section s is constructed. The set of directed edges on the graph is between all nodes in adjacent sections. That is,

$$E = \left\{ \bigcup_{s, i, j=1}^{N, Q_s, Q_{s+1}} E_{s, i, j} \right\} \text{ where } E_{s, i, j} = [R_{s, i}, R_{s+1, j}], \quad (8)$$

N is the total number of sections, and Q_s denotes the number of segmented regions in section s .

A path through the graph is defined as a sequence of nodes connected by edges. We are interested in paths that span all sections $P = (R_{1, i_1}, R_{2, i_2}, \dots, R_{N, i_N})$, and the cost of the path is defined as the sum of the costs of the edges

$$K(P) = \sum_{s=0}^{N-1} W(E_{s, i_s, i_{s+1}}), \quad (9)$$

where $i_1, \dots, i_N$ is the set of indices that the path follows on each section; because of the directed nature of the graph, paths cannot cross back to previous sections.

For biologists, the identification of neurons between sections relies on texture, shape, and proximity. These properties motivate the construction of the edge cost as the negative of the log-product of the correlation between regions and a Gaussian penalty on in-section displacement. That is:

$$W(E_{s, i, j}) = -\log \left[C(R_{s, i}, R_{s+1, j}) \exp \left(\frac{-D(R_{s, i}, R_{s+1, j})^2}{\phi^2} \right) \right], \quad (10)$$

where $D(R_{s, i}, R_{s+1, j})$ is the Euclidean distance between region center of mass in the $x - y$ coordinates of the section. ϕ is the maximum distance we expect the neurons to move between sections. C is the maximum value of the normalized cross-correlation of the two segmented regions. Correlation is used most commonly in image processing and computer vision for locating or matching specific features across scenes. In this case, it is used to measure how well a region in section s matches with another region in section $s + 1$. The two section images are multiplied with the characteristic function of the regions (0 outside, 1 inside) corresponding to $R_{s, i}$ and $R_{s+1, j}$ to obtain the masked images $I_{s, i}$ and $I_{s+1, j}$, respectively. Then, the normalized cross-correlation between two vertices of the graph is computed as

$$C(R_{s, i}, R_{s+1, j}) = \frac{\sum_{x, y} I'_{s, i}(x - t_x, y - t_y) I'_{s+1, j}(x, y)}{\max_{t_x, t_y} \sqrt{\left(\sum_{x, y} I'_{s, i}(x, y)^2 \right) \left(\sum_{x, y} I'_{s+1, j}(x, y)^2 \right)}}. \quad (11)$$

For computational efficiency, the cross-correlation is computed in the Fourier Domain. The log is used so that the formulation is equivalent to a product through the sections, and the system avoids seeking out very good connections at the expense of very bad ones. Cell identity is lost if a connection between sections is not sufficiently strong. Finally, the log-product, which can be seen as an edge connection weight, is negated to create a cost function.

An important extension to this basic framework allows paths to skip sections, in order to avoid poor quality sections, which can happen regularly. To accomplish this, edges are added to the graph that allow connections up to M sections away:

$$E = \left\{ \bigcup_{k,s,i,j=1}^{M,N,Q_s,Q_{s+k}} E_{s,i,j,k} \right\} \text{ where } E_{s,i,j,k} = [R_{s,i}, R_{s+k,j}] \quad (12)$$

where k is the number of skipped sections. For the datasets in this paper, $M = 2$, thereby allowing connections between sections separated at most by a single intermediate section. This gives Dijkstra's algorithm a choice in calculating the best path in the case where an immediately adjacent section does not have the best match. This changes the construction of costs for these edges, because we want to avoid cost functions that favor skipping sections when there is sufficient data to support a path through a section. The function in Equation 10 is adjusted to penalize the correlation and distance terms for the skipped sections. Generally we have

$$W(E_{s,i,j}) = -\log \left[\alpha^{k-1} C(R_{s,i}, R_{s+1,j}) \exp \left(\frac{-D(R_{s,i}, R_{s+1,j})^2}{k\phi^2} \right) \right], \quad (13)$$

α is the typical normalized correlation penalty between a cell in two adjacent sections, which was found empirically to be about 0.6. The displacement Gaussian's variance is multiplied by k , allowing more spatial movement when a section is skipped. The effect of these changes is to normalize the correlation, but allow for more displacement between the skipped regions. Overall, this increases the edge cost for $k > 1$.

Dijkstra's algorithm, which finds a minimum distance path in a directed graph is used to find the optimal connectivity for each neuron (region) in the first section. Dijkstra is run with a zero cost for all the regions in the first section. The region with the best cost is found on the last section, and tracing the solution backwards results in the optimal path (best cell) for the whole data set. Of course in this solution, cells can share paths, which is not normally what we want for this particular application. To account for this, we enforce uniqueness iteratively, in a greedy optimization strategy. That is, we solve for the best path, remove those nodes from the graph, and repeat, producing a sequence of cells associated with a decreasing degree of evidence for connectivity.

One of the constraints required by this method is that the user know how far a neuron will likely go across an image volume. This linking algorithm is designed to identify as many neurons as possible that start on slice s and end on slice $s + n$. This limits this method to parallel processes. However, for certain datasets, such as the *C. elegans* where neurons rarely branch or terminate, scientists can potentially use this automatic linking algorithm to reconstruct as many paths, P , as possible of the ventral nerve cord.

2.2.2 Neuron Reconstruction Viewer—The automatic methods described up until this point all work fairly well on their own, but in the end, require the ability for viewing and editing of the segmentation results. The Neuron Reconstruction Viewer (NeRV) (shown in Figure 11) attempts to bridge these two requirements by providing a visual interface to large volumes of EM images and neuron segmentations, with the option to make corrections that will, in the long term improve the segmentation.

Primarily, NeRV is an interface for the user to view the raw image data and the 3D reconstruction. Interacting with the image data and the rendered neuron provides insight for the scientist on the arrangement of the neurons within the data. The pane on the left, in Figure 11, is mainly a slice viewer. The user can view the membrane detection, the region

segmentation, and the raw data all in one viewer. Spheres highlight the paths neurons take through the volume. The keyboard arrow keys or the slider in the middle lets the user scroll through the sections. The pane on the right, is a 3D viewer of the reconstructed neuron. Raw image data can be turned off and on in this view, and users can select other sections simply by clicking on the area of the neuron.

Users can interact with this using the graphical interface on the far left. First, users can correct segmentations to close gaps with a simple drawing tool, then recompute the regions and correlations to improve the optimal path calculation (as discussed in Section 2.2). Users can manually select regions in slices and create their own 3D renderings with the automatic path calculation. For precomputed and segmented neurons, a separate window allows users to select different neurons for viewing, deleting, or joining.

Figure 12 is an overview of all the software and data used to generate results in this paper. Before NeRV can load any image data, a sequence of command line tools needs to be executed to generate the membrane detection, segmentations, correlations, optimal paths, and isosurfaces. NeRV could easily interact with all these tools at each step; however, because of the time it takes to process the data, it is easier to do the early steps off line. However, once generated, NeRV has the ability to quickly edit membranes and view segmentations through streaming of the image data. Only the data that the user is viewing is loaded into memory. The 3D visualization and path editor give the user an opportunity to view full models of neurons and join paths, one of the most crucial steps in the reconstruction.

NeRV is built primarily using VTK (Schroeder, Martin, and Lorensen 2010) and Qt (Nokia). To handle large datasets, the VTK image data streamer is used to load only the images required for viewing and requested by the user. Since slices are loaded as needed, the memory of this system is limited only by the size of a single section. Other optimizations, such as down sampling and memory management, enable efficient building of the isosurfaces for the 3D reconstruction in the right pane.

3 Neuron Segmentation Results

Two EM datasets are segmented using the proposed methods. The first dataset is a stack of 400 sections from the ventral nerve cord of the *C. elegans* worm. The second dataset is a stack of 400 sections from the mouse neuropil. These datasets contain very different types of neural cells. Furthermore, the *C. elegans* data has a resolution of $6\text{nm} \times 6\text{nm} \times 33\text{nm}$ and each 2D section is 4008×2672 pixels, whereas the mouse neuropil data has a pixel resolution of $10\text{nm} \times 10\text{nm} \times 50\text{nm}$ and each 2D section is 4096×4096 pixels. Figure 2 shows images from each of these datasets. Note that the membranes in the mouse and worm images, shown in Figure 2, are very different. The section from the worm nerve cord (Figure 2(a)) has a low signal-to-noise ratio and the neuron membranes have varying thickness and contrast. While the membranes from the mouse neuropil (Figure 2(b)) are strong in contrast and have a high signal-to-noise ratio, they contain more variable internal structures. Both datasets contain grazed membranes, corresponding to neurons cut at nonperpendicular angles. This makes it difficult for even the human eye to identify all the membrane structures. More traditional statistics-based machine learning methods would require a specific filter design for each dataset. However, the use of stencils, rather than a predefined filter bank, means the proposed method can adapt to the idiosyncrasies of different samples and is successful in learning to detect neuron membranes in both datasets.

To segment the neurons in these datasets we focused on identifying parallel processes. The *C. elegans* data was ideal for this solution because nerves running along the ventral nerve

cord rarely branched or terminated. In contrast, the mouse neuropil contains a higher number of branching structures, although in our close examination of the data, most parallel processes found in this data branched very little, maybe two to three slices. In an effort to segment as many neurons as possible, we also restrict our segmentation to processes that span a specific number of sections.

The results presented in this paper were generated using 2 different computers. The first was on a desktop computer containing 8, 2.8 GHz Intel CPUs and 8G of Memory. The second machine was a 32 node, 2.93Ghz, shared memory computer containing 200G of memory. The raw *C. elegans* data, if loaded entirely into memory at once, requires 4.2G of memory, while the mouse neuropil data requires 25G. Because of these memory requirements, distribution of the processing was done across computers, in parallel, for the most efficient computation of results as possible. Details regarding the time for each computation are described in detail in the following sections.

3.1 Results for the *C. elegans* Ventral Nerve Cord

The nematode *C. elegans* is an important organism for neural circuit reconstruction because it is the only organism for which the connectivity has been determined (White et al. 1986; Varshney et al. 2011). Nevertheless, there are still numerous questions that require the determination of the connectivity, such as how genes regulate wiring (Jin, Hoskins, and Horvitz 1994) or how connectivity is altered to mediate different behaviors, for example, between males and females (White et al. 2007). In addition, reconstructions of the full nervous system reveal topological characteristics of the neurons that are important for studies of neuronal functions. The particular dataset used in this paper is from the ventral nerve cord of the *C. elegans* and is important for studying the interwoven topology of neurons making connections to local targets.

To segment the membranes in this dataset and create a 3D reconstruction, we first had to align all the ssTEM images into a volume. We performed a ridged alignment using a brute-force search for the unknown rotation and translation between adjacent pairs of sections (Tasdizen et al. 2010). This was a challenging task because there are significant changes between the sections resulting from slicing artifacts and missing sections. Approximately 10% of the images required user intervention to remove images of poor quality or realign sections that had little correspondence. Using the tools described in Tasdizen *et al.* hand alignment of two sections took just a couple of minutes. We did not perform a non-linear alignment on these sections because we wanted to maintain the shape of the neurons and prevent distortion.

For validation, we had experts segment 40 selected images from the first 400 sections. Each expert placed a one pixel wide line along the membranes of the neurons, which we dilated using a 5 pixel wide structuring element, to cover most of the membrane pixels. Before training, we performed Gaussian blurring of the EM images with a small $\sigma = 2$ to remove noise, and down sampled by 2 to reduce the computational complexity. Then we used a contrast limited adaptive histogram equalization (CLAHE) (Pizer et al. 1990) filter to enhance the contrast in the neuron membranes. For the training data, 30 images were randomly selected for training and the remaining 10 were used for validation. From those images, 1 million samples were randomly selected from the manually marked images. Because of the relatively small percentage of positive examples (representing membrane pixels), these 1 million samples were chosen to contain $\frac{1}{3}$ positive and $\frac{2}{3}$ negative examples. The stencil used to sample the image values had a radius of 10 and was similar to the one in Figure 7. The ANN we used was implemented in C++ and had one hidden layer of 20 nodes. To mitigate problems with local minima in training, each network was trained for 5 Monte

Carlo simulations using randomly initialized weights. Each stage of the serial ANN took between 9 and 12 hours to complete. The ANN for the sequential sections took about 22 hours to train. Applying the weights from the ANNs takes a total of 7 minutes per section. Finally, the tensor voting implemented in Matlab was completed in 6 minutes per section.

Figure 13(a) shows three selected sections from the *C. elegans* dataset. The final membrane detection with the proposed method is shown in Figure 13(e). The sequential section ANN uses information about membranes also detected in neighboring sections to improve the current segmentation. The tensor voting uses, as input, the final classification and closes remaining gaps. This is demonstrated in closer detail in Figure 14. Improved membrane detection is annotated with yellow circles. Most often a strong membrane in a neighboring section provides confidence for enhancing membranes with poor contrast in the current section. Using difference images, Figure 15 demonstrates the improvements between the different methods. Figure 16 gives a numerical evaluation of the improvement between the serial ANN and the sequential section ANN on the validated test images using ROC curves. While the impact from the tensor voting method to close gaps is demonstrated qualitatively, its true positive and false positive values do not change as much. This is partly due to the dilation that results from the tensor voting and that the addition of pixels for closing is small compared to the overall segmentation. Figure 15(b) demonstrates this dilation effect. To evaluate the effectiveness of the sequential section ANN and the tensor voting applied to the output of the serial ANN, Figure 17 uses the Rand index (Rand 1971) to validate our methods. For the *C. elegans* data, the value of using detected membranes from sequential sections to improve the classification is quite large, while the tensor voting contribute is smaller. Since we developed the tensor voting to work with both datasets, we suspect the lack of improvement for the *C. elegans* dataset is because the tensor voting was better tuned for the mouse neuropil image data.

Figure 18 shows the 3D reconstruction of 10 neurons through the first 300 sections of the *C. elegans* ventral nerve cord. Building this reconstruction was a two part process. First, we identified 6 significant breaks in the image volume where there was missing data due to lost or badly imaged sections. These were places where the data had significant changes and would cause the region linking algorithm to fail. As a result, neuron regions were linked only between the sections without breaks, producing six sets of paths that spanned the whole volume. To completely reconstruct these paths through the all 300 sections, NeRV was used to manually merge neurons in sequential sections, forming complete reconstructions through the volume. This final step, computing the isosurface for all the neurons, took approximately 5 minutes per neuron to complete. Figure 18 is the final output from this process.

To validate the accuracy of the 3D reconstruction automatically generated by the neuron region linking method, we manually inspected the top 34 total paths returned for a sequence of images. We picked the series of images 85 through 190 to be our test volume. This sequence of images had the longest continuous set of sections without significant breaks. Figure 19 shows the number of regions each neuron successfully tracked through the volume. Of these 34 paths, 19 (58%) went through the entire dataset with no errors. Of the remaining paths, 9 correctly linked neuron regions for over half of the sections and 6 paths correctly linked regions through less than half of the sections. Despite the continuity of the set of sections, significant changes between neuron shape and location still existed within this block of images. One example of this is shown in Figure 3. These cases sometimes caused neuron paths to sometimes follow intercellular spaces (also included in the graph) or other neuron regions. Overall, the automatic segmentation still has the potential to reduce the amount of time a scientist is required to perform annotation by half for this dataset.

3.2 Results on the Mouse Neuropil

Understanding the connectivity, types of connections, and roles of different cells in the mouse neuropil is an increasingly common area of study. The 3D organization of these structures provides insight into how the nervous system functions at very basic levels (Watanabe et al. 2010). In an effort to better understand and statistically quantify these structures, we are segmenting a large volume of the mouse neuropil.

The entire neuropil dataset is $4096 \times 4096 \times 400$. To train and validate our neural networks, a subset of this data ($700 \times 700 \times 70$) was manually segmented using Amira (Imaging) by an expert. From that set, 42 images were randomly selected and used for training in our classifier. The training set contained 4.5 million examples. To decrease training time, the ANN was trained first on 1 million examples for 50 iterations. The weights from this network were used to initialize the ANN for the 4.5 million training examples. The ANN contained one hidden layer of 10 nodes. The images required no preprocessing to remove noise or enhance contrast and were sampled with a stencil of radius 10.

Figure 20 shows the segmentation results on three images from different sections. Figure 24 demonstrates in more detail the gap closing that occurs when the sequential section ANN and tensor voting are used. Figures 21 and 22 give a quantitative comparison on the improvement of the different methods. While the improvement in the ROC curve for the sequential section ANN is evident, the improvement after the tensor voting is more difficult to distinguish. However, evaluation of the difference images and the Rand index, shown in Figures 21 and 23, shows a much greater improvement towards closing gaps and strengthening boundaries.

Final reconstruction of the volume on the entire dataset turned this task into a large data challenge, since the actual size of the full volume is much larger than the training data. First, neuron membranes needed to be detected in each $4K \times 4K$ section, which took about 40 minutes for each section, including applying all the weights in series from the ANNs and using tensor voting. We then used the watershed segmentation algorithm on the blurred output from the tensor voting. As discussed in Section 2.1.4, the user chose a level for the watershed that gave us a slightly over segmented set of regions, ensuring that gaps were closed. Using the watershed implementation from ITK (Ibanez et al. 2005), a level of .15 was used. These became the regions used in the correlations and linking methods. The computation of the watersheds was much faster, taking only about one minute per section. Calculating the correlations between every two sections took between 40 minutes and 1.5 hours, depending on the number of regions in each section. To reduce computational time and the required memory, correlations were only calculated for regions within β distance away. For this dataset $\beta = 150$. Once all the correlations are calculated, a graph can be constructed and Dijkstra can be used to find paths through the volume. The generic implementation of Dijkstra's algorithm has a complexity of $O(r^2)$, where r is the number of nodes in the graph. However, for this case, since the edges in the graph are limited to

connections between sections, the complexity is $O\left(\frac{r^2}{N}\right)$. The graph can be made even more sparse by limiting the number of edges to regions by $D(R_{s,i}, R_{s+1,j}) < d$, where d is the maximum distance a region is allowed to connect between region centers. For this dataset $d = 100$. This means the algorithm can scale more easily to larger graphs. To scale the path calculation even further, we divided the volume into smaller slabs and found paths through every 25 sections. Each path then took approximately 4 minutes to compute and paths were easily joined by matching paths with overlapping sections. To view in 3D, each path that spanned more than 300 sections was rendered in 45 minutes. We used, to our advantage, multiple processor machines to compute these results as efficiently as possible, in parallel.

NeRV easily handles the size of this data because it only loaded into memory what was requested by the user. Finally, users can easily select the neurons they want to view in the volume. The final 3D visualization of this dataset can be seen in Figure 25.

Our neuron linking method was validated for this dataset using the small training volume, discussed earlier, which is a subset of the main volume. We calculated every possible path that would pass from the top of the stack to the bottom of the stack. After we eliminated paths that attempted to link neurons that passed through the edge and outside of the volume subset, we had 56 possible paths. Of these, 14 were correctly segmented neurons through the whole stack, 12 were correctly segmented through more than half of the dataset, 23 were correctly segmented through less than half of the dataset, and finally 7 paths had no regions correctly linked. Figure 26 is a plot of these results, demonstrating the accuracy of the neuron linking through a small volume that is representative of the main dataset.

4 Conclusion and Future Work

This paper presents a method for the segmentation and extraction of neurons from electron microscopy images. Membranes in each 2D section are initially detected using a series of ANNs. The output image from this algorithm is used in a final ANN which uses registered images to improve the classification. Examining the detected membranes in sequential sections, above and below the current section, helps the classifier learn to detect membranes that are grazed or complicated by internal cellular structures. In the last step, tensor voting closes remaining gaps in the image. A software tool, NeRV, provides an interface for users to correct 2D segmentations and manually assign neuron paths through sections. Users can also view automatically generated paths and join sections. This aids the user in visualizing the reconstruction data.

Future work in this area includes extending this framework to detect synapses and vesicles. In addition, we would like to incorporate a multiscale context sampling method to train the series ANNs. One of the drawbacks of this method is that the classifier is very sensitive to the data it uses in training. As a result, applying this method to new image data without first training a classifier is not possible. Recent, related work in this field is available in a tool called Ilastic (Sommer et al. 2011) which computes segmentations based on user inputs. Building on these concepts, we would like to improve the performance and flexibility of our algorithms to enable the incorporation of interactive user input into our model. The path computation and linking of neuron membranes might also be more efficient to compute using a more flexible graph matching algorithm, incorporating recent work by Funke *et al.* (Funke et al. 2011). This would be beneficial because it could more easily handle neuron branching and termination. Finally, NeRV should be extended with more user interfaces for editing neuron segmentations, such as making corrections and handling branching.

We plan on extending this method on the full *C. elegans* ventral nerve cord dataset. This would reveal in 3D the physical layout of neurons and can be compared to the data from White *et al.* which is the current gold-standard for neuron connectivity in the *C. elegans*. Further work to segment the synapses in the neurons and the muscles that surround the nerve cord would provide insight into communication and wiring in the *C. elegans*. Likewise, a more thorough analysis of the 3D reconstructions in the mouse neuropil need to be completed so we can develop a better understanding of the types of connections present in this dataset.

5 Software Availability

NeRV and related command line tool software are available at <http://www.sci.utah.edu/~liz/nerv.html> for download. The mouse neuropil image data is available from ccdb.ucsd.edu as a microscopy product with the ID 8192.

Acknowledgments

This work was supported by NIH R01 EB005832 (TT), HHMI (EMJ), NIH NINDS 5R37NS34307-15 (EMJ) and 1R01NS075314 (MHE, TT) as well as NIH NCRR for support of the National Center for Microscopy and Imaging Research at UCSD, 5P41RR004050 (MHE). We are grateful to Nikita Thomas, Nels B. Jorgensen, Jeremy B. Thompson, and Blake Paulin for their help in imaging the *c. elegans* VNC and Eric Bushong and Thomas Deerinck for their work in preparing the examples from the mouse cerebellum.

References

- Allen, Barry A.; Levinthal, Cyrus. CARTOS II semi-automated nerve tracing: Three-dimensional reconstruction from serial section micrographs. *Computerized Medical Imaging and Graphics*. 1990; 14(5):319–329. [PubMed: 2224830]
- Anderson JR, Jones BW, Watt CB, Shaw MV, Yang JH, Demill D, Lauritzen JS, Lin Y, Rapp KD, Mastronarde D, Koshevoy P, Grimm B, Tasdizen T, Whitaker R, Marc RE. Exploring the retinal connectome. *Mol. Vis*. 2011; 17:355–379. [PubMed: 21311605]
- Anderson JR, Jones BW, Yang J-H, Shaw MV, Watt CB, Koshevoy P, Spal-tenstein J, Jurrus E, Kannan UV, Whitaker RT, Mastronarde D, Tasdizen T, Marc RE. A computational framework for ultrastructural mapping of neural circuitry. *PLoS Biology*. 2009; 7(3):e74.
- Andres, Björn; Köthe, Ullrich; Helmstaedter, Moritz; Denk, Winfried; Hamprecht, Fred A. Segmentation of SBFSEM volume data of neural tissue by hierarchical classification. In: Rigoll, Gerhard, editor. *Pattern Recognition*. Vol. Vol. 5096. Springer; 2008. p. 142-152. of *LNCS*
- Bertalmío, Marcelo; Sapiro, Guillermo; Randall, Gregory. Morphing active contours. *IEEE Trans. Pattern Anal. Mach. Intell*. 2000; 22:733–737.
- Betzig E, Patterson GH, Sougrat R, Lindwasser OW, Olenych S, Bonifacino JS, Davidson MW, Lippincott-Schwartz J, Hess HF. Imaging intracellular fluorescent proteins at nanometer resolution. *Science*. 2006; 313(5793):1642–1645. [PubMed: 16902090]
- Briggman, Kevin L.; Denk, Winfried. Towards neural circuit reconstruction with volume electron microscopy techniques. *Current Opinion in Neurobiology*. 2006a; 16(5):562–570. [PubMed: 16962767]
- Briggman, Kevin L.; Denk, Winfried. Towards neural circuit reconstruction with volume electron microscopy techniques. *Current Opinion in Neurobiology*. 2006b; 16(5):562–570. [PubMed: 16962767]
- Cardona, Albert; Saalfeld, Stephan; Preibisch, Stephan; Schmid, Benjamin; Cheng, Anchi; Pulokas, Jim; Toman áak, Pavel; Hartenstein, Volker. An integrated micro- and macroarchitectural analysis of the *Drosophila* brain by computer-assisted serial section electron microscopy. *PLoS Biol*. 2010; 8(10):e1000502. [PubMed: 20957184]
- Chklovskii, Dmitri B.; Vitaladevuni, Shiv; Scheffer, Louis K. Semi-automated reconstruction of neural circuits using electron microscopy. *Current Opinion in Neurobiology*. 2010; 20(5):667–675. [PubMed: 20833533]
- Cottrell, GW. Extracting features from faces using compression networks: face, identity, emotion and gender recognition using holons. Morgan Kaufmann; 1990. p. 328-337.
- Deerinck, Thomas J.; Bushong, Eric A.; Thor, Andrea; Ellisman, Mark H. NCMIR methods for 3D EM: A new protocol for preparation of biological specimens for serial block face scanning electron microscopy. *Microscopy and Microanalysis Meeting*. 2010
- Denk W, Horstmann H. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol*. 2004; 2(11)
- Denk W, Strickler JH, Webb WW. Two-photon laser scanning microscopy. *Science*. 1990; 248:73–76. [PubMed: 2321027]

- Egner A, Hell SW. Fluorescence microscopy with super-resolved optical sections. *Trends Cell Biol.* 2005; 15(4):207–215. [PubMed: 15817377]
- Fiala JC, Harris KM. Computer-based alignment and reconstruction of serial sections. *Microscopy and Analysis.* 2002; 87:5–8.
- Fiala JC, Harris KM. Extending unbiased stereology of brain ultrastructure to three-dimensional volumes. *J Am Med Inform Assoc.* 2001; 8(1):1–16. [PubMed: 11141509]
- Fiala, John C.; Harris, Kristen M. 2010. Synapseweb. <http://synapses.clm.utexas.edu/tools/reconstruct/reconstruct.stm>
- Franken E, Almsick M, Rongen P, Florack LMJ, Haar Romeny BM. An efficient method for tensor voting using steerable filters. *ECCV06.* 2006:228–240. IV.
- Funke, Jan; Andres, Björn; Hamprecht, Fred A.; Cardona, Albert; Cook, Matthew. Multi-hypothesis crf-segmentation of neural tissue in anisotropic em volumes CoRR. 2011 abs/1109.2449.
- Gonzalez, Rafael C.; Woods, Richard E. *Digital Image Processing.* Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.; 1992.
- Gonzalez-Hernandez, Marta; Pablo, Luis E.; Armas-Domingue, Karintia; Vega, Ricardo; Ferreras, Antonio; Rosa, Manuel. Structure-function relationship depends on glaucoma severity. *British Journal of Ophthalmology.* 2008
- Haykin, S. *Neural networks - A comprehensive foundation.* 2nd edition. Prentice-Hall; 1999.
- Ibanez, L.; Schroeder, W.; Ng, L.; Cates, J. *The ITK Software Guide.* second edition. Kitware, Inc.; 2005. ISBN 1-930934-15-7, <http://www.itk.org/ItkSoftwareGuide.pdf>
- Imaging, Visage. Visage imaging, amira. <http://www.amira.com>
- Jain, V.; Bollmann, B.; Richardson, M.; Berger, DR.; Helmstaedter, MN.; Briggman, KL.; Denk, W.; Bowden, JB.; Mendenhall, JM.; Abraham, WC.; Harris, KM.; Kasthuri, N.; Hayworth, KJ.; Schalek, R.; Tapia, JC.; Lichtman, JW.; Seung, HS. Boundary learning by optimization with topological constraints; *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on;* 2010. p. 2488-2495.
- Jain, V.; Murray, JF.; Roth, F.; Turaga, S.; Zhigulin, V.; Briggman, KL.; Helmstaedter, MN.; Denk, W.; Seung, HS. Supervised learning of image restoration with convolutional networks; *IEEE 11th International Conference on Computer Vision;* 2007. p. 1-8.
- Jain, Viren; Sebastian Seung, H.; Turaga, Srinivas C. Machines that learn to segment images: a crucial technology for connectomics. *Current Opinion in Neurobiology.* 2010; 20(5):653–666. [PubMed: 20801638]
- Jeong, Won-Ki; Beyer, Johanna; Hadwiger, Markus; Blue, Rusty; Law, Charles; Vazquez-Reina, Amelio; Clay Reid, R.; Lichtman, Jeff; Pfister, Hanspeter. Ssecret and neurotrace: Interactive visualization and analysis tools for large-scale neuroscience data sets. *IEEE Computer Graphics and Applications.* 2010; 30:58–70. [PubMed: 20650718]
- Jeong, Won-Ki; Beyer, Johanna; Hadwiger, Markus; Vazquez, Amelio; Pfister, Hanspeter; Whitaker, Ross T. Scalable and interactive segmentation and visualization of neural processes in EM datasets. *IEEE Transactions on Visualization and Computer Graphics.* 2009; 15:1505–1514. [PubMed: 19834227]
- Jin Y, Hoskins R, Horvitz HR. Control of type-D GABAergic neuron differentiation by *C. elegans* UNC-30 homeodomain protein. *Nature.* 1994; 372(6508):6780–6783.
- Jones BW, Marc RE. Retinal remodeling during retinal degeneration. *Exp. Eye Res.* 2005; 81:123–137. [PubMed: 15916760]
- Jones BW, Watt CB, Frederick JM, Baehr W, Chen CK, Levine EM, Milam AH, Lavail MM, Marc RE. Retinal remodeling triggered by photoreceptor degenerations. *J. Comp. Neurol.* 2003; 464:1–16. [PubMed: 12866125]
- Jones BW, Watt CB, Marc RE. Retinal remodelling. *Clin Exp Optom.* 2005; 88:282–291. [PubMed: 16255687]
- Jurrus E, Hardy M, Tasdizen T, Fletcher PT, Koshevoy P, Chien C-B, Denk W, Whitaker RT. Axon tracking in serial block-face scanning electron microscopy. *Medical Image Analysis.* 2009; 13(1): 180–188. [PubMed: 18617436]

- Jurrus E, Paiva ARC, Watanabe S, Anderson JR, Jones BW, Whitaker RT, Jorgensen EM, Marc RE, Tasdizen T. Detection of neuron membranes in electron microscopy images using a serial neural network architecture. *Medical Image Analysis*. 2010; 14(6):770–783. [PubMed: 20598935]
- Jurrus, E.; Paiva, ARC.; Watanabe, S.; Whitaker, R.; Jorgensen, EM.; Tasdizen, T. Serial neural network classifier for membrane detection using a filter bank; *Proc. Workshop on Microscopic Image Analysis with Applications in Biology*; 2009.
- Jurrus, E.; Whitaker, RT.; Jones, B.; Marc, R.; Tasdizen, T. An optimal-path approach for neural circuit reconstruction; *Proceedings of the 5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*; 2008. p. 1609-1612.
- Kaynig, Verena; Fuchs, Thomas; Buhmann, Joachim M. Neuron geometry extraction by perceptual grouping in ssTEM images. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*. 2010; 0:2902–2909.
- Kremer, James R.; Mastronarde, David N.; Richard McIntosh, J. Computer visualization of three-dimensional image data using imod. *Journal of Structural Biology*. 1996; 116(1):71–76. [PubMed: 8742726]
- Kumar, Ritwik; Reina, Amelio V.; Pfister, Hanspeter. Radon-like features and their application to connectomics. *MMBIA*. 2010
- Leng, Zhe; Korenberg, Julie R.; Roysam, Badri; Tasdizen, Tolga. Automatic markup of neural cell membranes using boosted decision stumps; *Proceedings of the 6th IEEE International Symposium on Biomedical Imaging*; 2011.
- Macke JH, Maack N, Gupta R, Denk W, Schölkopf B, Borst A. Contour-propagation algorithms for semi-automated reconstruction of neural processes. *Journal of Neuroscience Methods*. 2008; 167:349–357. [PubMed: 17870180]
- Marc RE, Jones BW, Anderson JR, Kinard K, Marshak DW, Wilson JH, Wensel T, Lucas RJ. Neural reprogramming in retinal degeneration. *Invest. Ophthalmol. Vis. Sci*. 2007; 48:3364–3371. [PubMed: 17591910]
- Marc RE, Jones BW, Watt CB, Strettoi E. Neural remodeling in retinal degeneration. *Prog Retin Eye Res*. 2003; 22:607–655. [PubMed: 12892644]
- Marc RE, Jones BW, Watt CB, Vazquez-Chona F, Vaughan DK, Organisciak DT. Extreme retinal remodeling triggered by light damage: implications for age related macular degeneration. *Mol. Vis*. 2008; 14:782–806. [PubMed: 18483561]
- Martone, Maryann E.; Tran, Joshua; Wong, Willy W.; Sargis, Joy; Fong, Lisa; Larson, Stephen; Lamont, Stephan P.; Gupta, Amarnath; Ellisman, Mark H. The cell centered database project: An update on building community resources for managing and sharing 3D imaging data. *Journal of Structural Biology; The 4th International Conference on Electron Tomography, The 4th International Conference on Electron Tomography*; 2008. p. 220-231.
- Medioni, Gerard; Lee, Mi-Suen; Tang, Chi-Keung. *Computational Framework for Segmentation and Grouping*. New York, NY, USA: Elsevier Science Inc.; 2000.
- Minsky, M. Microscopy apparatus. U.S. Patent. number 301467. 1961.
- Mishchenko, Yuriy. Automation of 3D reconstruction of neural tissue from large volume of conventional serial section transmission electron micrographs. *J Neurosci Methods*. 2008
- Mishchenko, Yuriy; Hu, Tao; Spacek, Josef; Mendenhall, John; Harris, Kristen M.; Chklovskii, Dmitri B. Ultrastructural analysis of hippocampal neuropil from the connectomics perspective. *Neuron*. 2010; 67:1009–1020. [PubMed: 20869597]
- Nokia. Qt. Cross-platform application and UI framework. <http://qt.nokia.com>
- Paiva, ARC.; Jurrus, E.; Tasdizen, T. Using sequential context for image analysis; In *Pattern Recognition (ICPR), 2010 20th International Conference on*; 2010. p. 2800-2803.
- Peng YW, Hao Y, Petters RM, Wong F. Ectopic synaptogenesis in the mammalian retina caused by rod photoreceptor-specific mutations. *Nat. Neurosci*. 2000; 3:1121–1127. [PubMed: 11036269]
- Pizer, SM.; Johnston, RE.; Erickson, JP.; Yankaskas, BC.; Muller, KE. Contrast-limited adaptive histogram equalization: speed and effectiveness; *Visualization in Biomedical Computing, 1990., Proceedings of the First Conference on*; 1990. p. 337-345.

- Pomerleau, Dean. Knowledge-based training of artificial neural networks for autonomous robot driving. In: Connell, J.; Mahadevan, S., editors. Robot Learning. Kluwer Academic Publishing; 1993. p. 19-43.
- Principe, JoséC.; Euliano, Neil R.; Curt Lefebvre, W. Neural and Adaptive Systems: Fundamentals through simulations. John Wiley & Sons; 2000.
- Rabi G, Lu SW. Visual speech recognition by recurrent neural networks. JEL. 1998; 7(1):61–69.
- Rand, William M. Objective criteria for the evaluation of clustering methods. Journal of the American Statistical Association. 1971; 66(336):846–850.
- Rueckert D, Sonoda LI, Hayes C, Hill DLG, Leach MO, Hawkes DJ. Non-rigid registration using free-form deformations: application to breast mr images. Medical Imaging, IEEE Transactions on. 1999; 18(8):712–721.
- Rust, Michael J.; Bates, Mark; Zhuang, Xiaowei. Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (storm). Nature Methods. 2006; 3(10):793–796. [PubMed: 16896339]
- Saalfeld, Stephan; Cardona, Albert; Hartenstein, Volker; Toman ák, Pavel. As-rigid-as-possible mosaicking and serial section registration of large ssTEM datasets. Bioinformatics. 2010; 26(12):i57–i63. [PubMed: 20529937]
- Schroeder, W.; Martin, K.; Lorensen, B. The VTK User's Guide. 11th edition. Kitware, Inc.; 2010. ISBN 1-930934-19-X, <http://www.vtk.org>
- Sommer, C.; Strahle, C.; Köthe, U.; Hamprecht, F. ilastik: Interactive learning and segmentation toolkit; Proceedings of the 6th IEEE International Symposium on Biomedical Imaging; 2011.
- Tasdizen T, Whitaker R, Marc R, Jones B. Enhancement of cell boundaries in transmission electron microscopy images. ICIP. 2005:642–645.
- Tasdizen, Tolga; Koshevoy, Pavel; Grimm, Bradley C.; Anderson, James R.; Jones, Bryan W.; Watt, Carl B.; Whitaker, Ross T.; Marc, Robert E. Automatic mosaicking and volume assembly for high-throughput serial-section transmission electron microscopy. Journal of Neuroscience Methods. 2010; 193(1):132–144. [PubMed: 20713087]
- Turaga, Srinivas C.; Briggman, Kevin L.; Helmstaedter, Moritz; Denk, Winfried; Seung, HSebastian. Maximin affinity learning of image segmentation. CoRR. 2009 abs/0911.5372.
- Turaga, Srinivas C.; Murray, Joseph F.; Jain, Viren; Roth, Fabian; Helmstaedter, Moritz; Briggman, Kevin; Denk, Winfried; Seung, HSebastian. Convolutional networks can learn to generate affinity graphs for image segmentation. Neural Computation. 2010; 22(2):511–538. [PubMed: 19922289]
- Varshney, Lav R.; Chen, Beth L.; Paniagua, Eric; Hall, David H.; Chklovskii, Dmitri B. Structural properties of the *Caenorhabditis elegans* neuronal network. PLoS Com-put Biol. 2011; 7(2):e1001066.
- Vazquez L, Sapiro G, Randall G. Segmenting neurons in electronic microscopy via geometric tracing. Proc. of ICIP. 1998:814–818.
- Vazquez-Reina A, Miller E, Pfister H. Multiphase geometric couplings for the segmentation of neural processes. Computer Vision and Pattern Recognition, IEEE Computer Society Conference on. 2009; 0:2020–2027.
- Venkatataju, Kannan U.; Paiva, A.; Jurrus, E.; Tasdizen, T. Automatic markup of neural cell membranes using boosted decision stumps; Proceedings of the 6th IEEE International Symposium on Biomedical Imaging; 2009. p. 1039-1042.
- Vu, Nhat; Manjunath, BS. Graph cut segmentation of neuronal structures from transmission electron micrographs; Image Processing, 2008 ICIP 2008 15th IEEE International Conference on; 2008. p. 725-728.
- Watanabe K, Takeishi H, Hayakawa T, Sasaki H. Three-dimensional organization of the perivascular glial limiting membrane and its relationship with the vasculature: a scanning electron microscope study. Okajimas folia anatomica japonica. 2010; 87(3):109–121. [PubMed: 21174940]
- Wells G, Venaille C, Torras C. Promising research: Vision-based robot positioning using neural networks. IVC. 1996; 14(10):715–732.
- White JQ, Nicholas TJ, Gritton J, Truong L, Davidson ER, Jor-gensen EM. The sensory circuitry for sexual attraction in *C. elegans* males. Curr. Biol. 2007; 17(21):1847–1857. [PubMed: 17964166]

- White JG, Southgate E, Thomson JN, Brenner FRS. The structure of the nervous system of the nematode *caenorhabditis elegans*. *Phil. Trans. Roy. Soc. London Ser. B Biol. Sci.* 1986; 314:1–340. [PubMed: 22462104]
- Xiao YP, Wang Y, Felleman DJ. A spatially organized representation of colour in macaque cortical area v2. *Nature*. 2003; 421(6922):535–539. [PubMed: 12556893]
- Yang, Huei-Fang; Choe, Yoonsuck. Cell tracking and segmentation in electron microscopy images using graph cuts; *Biomedical Imaging: From Nano to Macro*, 2009. ISBI '09 IEEE International Symposium on; 2009. p. 306-309.

Figure 1.
Diagram demonstrating the formation of fuzzy membranes in ssTEM images.

(a)

(b)

Figure 2.
(a) ssTEM image of the ventral nerve cord from the *C. elegans* (b) SBFSEM image of the mouse neuropil

Figure 3.
Example showing how much change frequently occurs in neuron profiles between sequential sections from an (a) ssTEM image of the ventral nerve cord from the *C. elegans* and (b) SBFSEM image of the mouse neuropil

Figure 4.

Output of the method on a test image from an SBFSEM dataset. (a) is the raw image, (b)–(d) are stages 1, 2 and 5 of the series ANN, (e) is the output from the sequential series ANN, (f) is the output from the tensor voting, (g) is the region segmentation after a simple flood fill, and (h) the gold standard, generated by an expert, for membranes and neuron regions.

Figure 6.
Artificial neural network diagram with one hidden layer.

Figure 7. Image neighborhood sampling technique: image pixels sampled using a stencil. For this example, the stencil contains a small number of samples, yet covers a larger area of the data. This is an efficient representation for sampling the image space.

Figure 8.
[color] Two sequential sections from the mouse neuropil with membranes detected after the serial ANN overlaid with each other with (a) no registration and (b) after the intensity-based nonlinear registration. Blue and yellow colors indicate membrane overlay mismatches and white indicates shared membranes.

Figure 9.

Diagram demonstrating the flow of data for the sequential section ANN architecture on a single image. I_i are the input images, *Serial ANN* is the diagram in Figure 5 collapsed, and C_i is the output of the classifier on image I_i . $C_{i-1} \circ t_{i-1}(x)$ is the registration of C_{i-1} to C_i and $C_{i+1} \circ t_{i+1}(x)$ is the registration of C_{i+1} to C_i . C_i^3 is the stack of all three registered images. S' is the 3D stencil used on the combined images as input to the classifier. C'_i is the final classification. Figure 10: Tensor and tensor voting field

(a) Ellipse representation of a tensor

(b) a typical tensor voting field

Figure 10.
Tensor and tensor voting field

Figure 11.
[color] Screen capture of NeRV displaying the automatic segmentation results on the *C. elegans* ventral nerve cord for a portion of the data.

Figure 12.

[color] Diagram demonstrating how the command line tools for segmenting EM images and NeRV interact with the EM images and data files. Directions of arrows indicate reads and writes. Double arrows designate data that is being streamed by VTK. The isosurfacing command line tool also takes, as input, the labeled images. "Binary Corr." are the correlations files for all the section to section edge weights, stored in binary for fast reading and writing of large files.

Figure 13.

[color] Output of the method on *C. elegans* test images. (a) is the raw image, (b) is the output from the final stage of the series ANN (Section 2.1.1), (c) is the output from the sequential section ANN (Section 2.1.2), (d) is the output after tensor voting (Section 2.1.3), and (e) is the segmentations of the neuron regions from a flood fill.

Figure 14.

[color]Example images demonstrating how the different elements of the proposed method contribute to strengthen undetected or grazed membranes, and close gaps on the *C. elegans* ventral nerve cord data. (a) is the raw image, (b) is the output from the final stage of the series ANN (Section 2.1.1), (c) is the output from the sequential section ANN (Section 2.1.2), and (d) is the final output after tensor voting. Yellow circles highlight improved membrane detection and gap closing that results from these methods. (f) is the expert annotated gold standard.

Figure 15. Difference images for the image in Figure 13, middle row, showing the change in membrane detection between (a) the series ANN and the sequential section ANN and (b) the sequential section ANN and tensor voting. White pixels indicate added membranes, black indicates removed membrane pixels, and grey is unchanged.

Figure 16.

These ROC curves show the improvement for the *C. elegans* data using the true positive and false negative membrane pixel classification rate with the use of the sequential section ANN (Seq. Sec. ANN) and the tensor voting, compared to using the serial ANN alone. Each curve the results over a series of threshold values from the output of the ANNs. The output from the tensor voting is binary and, thus, is represented by a single point.

Rand Curves for the C. elegans Testing and Training Data

Figure 17.

This plot shows the improvement of the Rand index, which measures the similarity between two segmentations after the new methods, proposed in this paper, are applied. For each point in this plot, the segmentation for the methods is compared to the segmentation from the truth data.

Figure 18.

[color]Two views of 10 neurons spanning 300 sections of the ventral nerve cord of the *C. elegans*. Neuron paths were generated automatically between six pairs of sections where known breaks in the image data existed. NeRV was used to connect paths between the breaks. Arrows identify discontinuities in neurons where some of these breaks occurred.

Figure 19.

This plot shows, in descending order, the number of regions correctly linked across slices for a continuous series of images from the *C. elegans* volume.

Figure 20.

[color] Output of the method on a three different test images from the mouse neuropil. (a) is the raw image, (b) is the output from the final stage of the series ANN (Section 2.1.1), (c) is the output from the sequential section ANN (Section 2.1.2), (d) is the output after tensor voting (Section 2.1.3) and (e) is the segmentations of the neuron regions from a flood fill.

(a)

(b)

Figure 21.

Difference images for the mouse neuropil section in Figure 20, middle row, showing the change in membrane detection between (a) the series ANN and the sequential section ANN and (b) the sequential section ANN and tensor voting. White pixels indicate added membranes, black indicates removed membrane pixels, and grey is unchanged.

ROC Curves for Mouse Neuropil Testing Data

Figure 22.

ROC curves showing the improvement of the membrane pixel classification rates on the mouse neuropil data using the two methods highlighted in this paper.

Rand Curves for Mouse Neuropil Testing and Training Data

Figure 23.

The Rand index is used to demonstrate the similarity between the segmentations produced after a flood fill of the output from the sequential section ANN and tensor voting when compared to the true neuron regions.

Figure 24.

[color] Example images demonstrating this method closing gaps on neuropil data. (a) is the raw image, (b) is the output from the final stage of the series ANN (Section 2.1.1), (c) is the output from the sequential section ANN (Section 2.1.2), and (d) is the final output after tensor voting. Yellow circles highlight membranes that were enhanced and detected using these methods. Finally, row (e) is the ground truth.

Figure 25.

[color]Two views of 14 fully automatically segmented parallel fibers spanning 400 sections of the mouse neuropil. The larger 3 structures, Purkinje cells, were segmented manually using the NeRV interface. Discontinuities in the neuron renderings indicate sections of the automatic algorithm that were skipped because of changing neuron regions.

Figure 26.

This plot shows, in descending order, the number of regions correctly linked, for each neuron, across the subset volume of the mouse neuropil.