

Automated tracing of myelinated axons and detection of the nodes of Ranvier in serial images of peripheral nerves

A. KRESHUK*, R. WALECKI*, U. KOETHE*, M. GIERTHMUEHLEN†, D. PLACHTA‡, C. GENOUD§, K. HAASTERT-TALINI|| & F.A. HAMPRECHT*

*Heidelberg Collaboratory for Image Processing (HCI), Interdisciplinary Center for Scientific Computing (IWR), University of Heidelberg, Heidelberg, Germany

†Department of Neurosurgery, University Medical Center, Freiburg, Germany

‡Department of Microsystems, Engineering, University of Freiburg, Freiburg, Germany

§Facility for Advanced Imaging and Microscopy, Friedrich Miescher Institute for Biomedical Research (FMI), Basel, Switzerland

||Institute of Neuroanatomy, Hannover Medical School, Hannover, Germany and Center for Systems Neurosciences (ZSN), Hannover, Germany

Key words. Axon, detection, nerve, Ranvier, segmentation, tracing.

Summary

The development of realistic neuroanatomical models of peripheral nerves for simulation purposes requires the reconstruction of the morphology of the myelinated fibres in the nerve, including their nodes of Ranvier. Currently, this information has to be extracted by semimanual procedures, which severely limit the scalability of the experiments.

In this contribution, we propose a supervised machine learning approach for the detailed reconstruction of the geometry of fibres inside a peripheral nerve based on its high-resolution serial section images. Learning from sparse expert annotations, the algorithm traces myelinated axons, even across the nodes of Ranvier. The latter are detected automatically.

The approach is based on classifying the myelinated membranes in a supervised fashion, closing the membrane gaps by solving an assignment problem, and classifying the closed gaps for the nodes of Ranvier detection.

The algorithm has been validated on two very different datasets: (i) rat vagus nerve subvolume, SBFSEM microscope, 200 × 200 × 200 nm resolution, (ii) rat sensory branch subvolume, confocal microscope, 384 × 384 × 800 nm resolution. For the first dataset, the algorithm correctly reconstructed 88% of the axons (241 out of 273) and achieved 92% accuracy on the task of Ranvier node detection. For the second dataset, the gap closing algorithm correctly closed 96.2% of the gaps, and 55% of axons were reconstructed correctly through the whole volume. On both datasets, training the algorithm on a small data subset and applying it to the full dataset takes a fraction of the time required by the currently used semiautomated protocols. Our software, raw data and ground truth

annotations are available at <http://hci.iwr.uni-heidelberg.de/Benchmarks/>. The development version of the code can be found at <https://github.com/RWalecki/ATMA>.

Introduction

The main application target of our work is the development of pattern recognition techniques to build realistic neural simulation models for the engineering of neural electrodes. Neural cuff electrodes are widely used in research and clinical practice as bipolar devices e.g. for vagal nerve stimulation, as they allow for chronic implantation without damaging the nerve's integrity. One disadvantage of these electrodes is that they can only record surface compound potentials, further localization of the signal source inside the nerve is not easily achievable. The same is true for the case of stimulation: a selective electrical stimulation of localized areas inside the nerve is not possible with simple bipolar stimuli. This lack of selectivity has recently been addressed by the introduction of a multichannel cuff electrode with 24 contacts, which allows for differentiated registration of the surface compound potential in order to localize the signal source and stimulate specific nerve areas and fascicles (Plachta *et al.*, 2013). Such selective registration and stimulation might for instance hold the key for selective vagal nerve stimulation against high blood pressure (Plachta *et al.*, 2014).

However, the more sophisticated nature of these electrodes requires much more detailed knowledge of the nerve anatomy. Usually, several cycles of *in vivo* experiments are required in order to adapt the electrode to the geometry of its target nerve.

A computer simulation of the nerve would greatly improve animal protection as well as allow for faster feedback cycles in the experiments. Such 'virtualization' of the nerve requires precise knowledge of the positions and morphology of the nerve fibres as well as the locations of the nodes of Ranvier

Authors A. Kreshuk and R. Walecki have contributed equally to this work.

Correspondence to: Anna Kreshuk, HCI/IWR, Speyerer Str. 6, Heidelberg, 69115, Germany. Tel: ++49-6221-545274; fax: ++49-6221-545276; e-mail: Anna.Kreshuk@iwr.uni-heidelberg.de

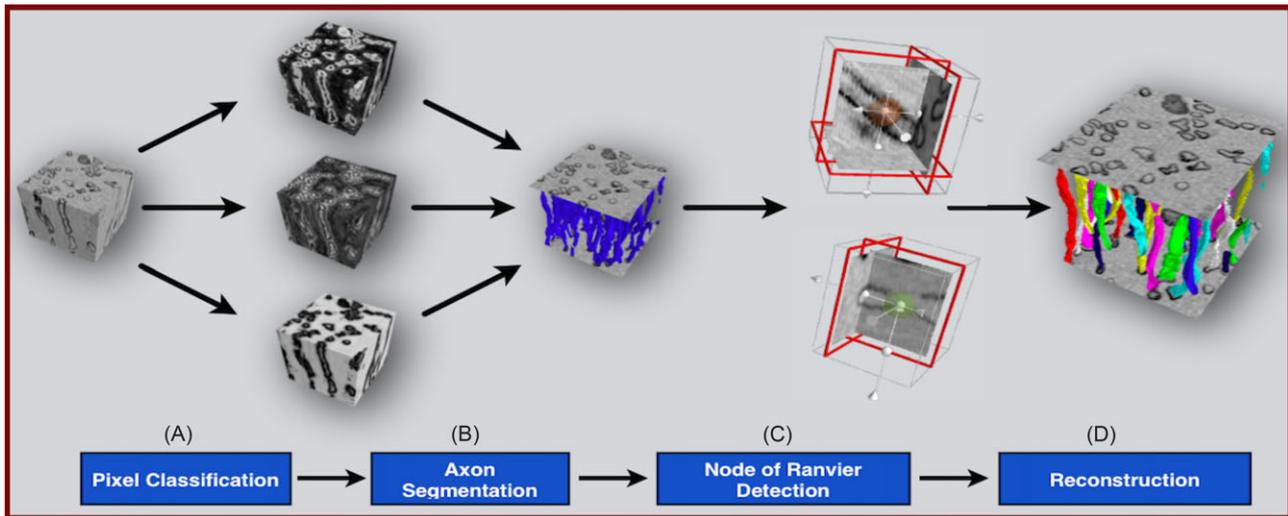


Fig. 1. Project pipeline. First, prediction maps are computed for three classes: myelinated membranes, axon interior and background (A). Parts of axons are then extracted by connected component analysis (B). Additional gap features are computed and used for the nodes of Ranvier detection (C). The gaps in the membranes are closed (D) and the axon tracing is completed.

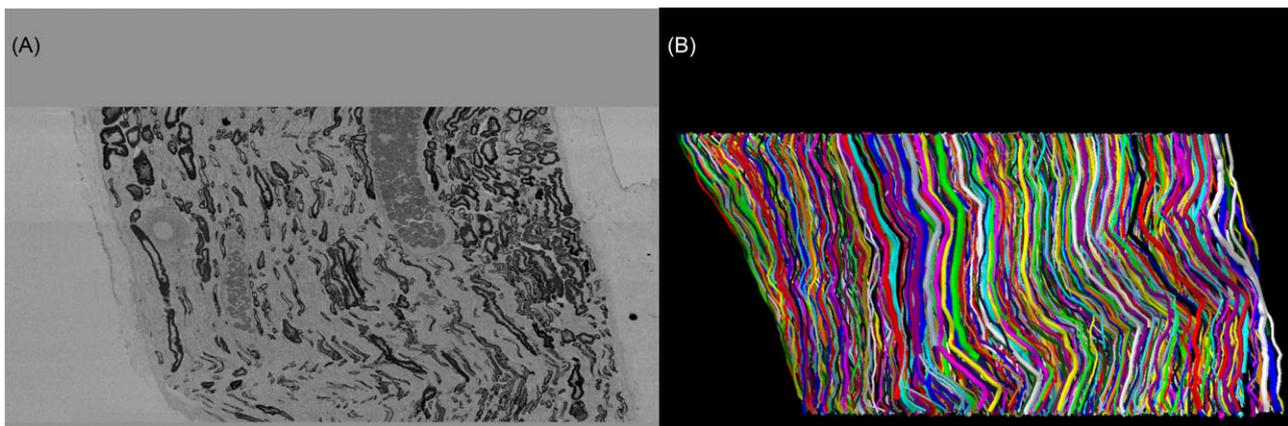


Fig. 2. (A) one slice of the raw data in the x - z projection. (B) Side view of the reconstructed axons. Only axons, which pass through at least 700 slices are shown for clarity.

– those parts of the fibre, where it is not covered by a myelin sheath. Gierthmuehlen *et al.* (2013) demonstrate that this information can be extracted from high-resolution volumetric images of the nerve, such as the ones shown in Figures 1 and 2(A). In Gierthmuehlen *et al.* (2013) a small cutaneous nerve was virtualized by semimanual analysis of a serial light microscopy image stack. The analysis included tracing 300 axons over a distance of 500 μm (625 images), which required 60 h of manual processing. Analysis of larger, more interesting nerves, such as the rat's vagus with hundreds or thousands of axons, would be almost infeasible by such semimanual means.

This is exactly the problem we aim to solve with this contribution. Our goals are (i) to provide the exact positions and trajectories of the myelinated axons in the nerve for the analysis of surface potentials and (ii) to localize the nodes of Ranvier for the selective stimulation of the nerve fibres by cuff electrodes.

Both are very important for realistic simulation of signal propagation: myelin sheaths serve as insulators, whereas nodes of Ranvier accelerate signal transmission by enabling saltatory conduction. Besides neural modelling, this work could also be useful for posttraumatic nerve regeneration studies, which need to differentiate between regrowing axons and axons with Ranvier nodes, which have already reconnected to the target tissue, but are not yet fully myelinated.

Over the last two decades, various methods have been proposed for semi- and fully-automated segmentation of myelin sheaths in two-dimensional (2D), in the context of neuromorphometry studies. Most of these methods start from thresholding the image by a global threshold (Vita *et al.*, 1992; Auer, 1994; Hunter *et al.*, 2007) and proceed to remove false positive detections by application of semantic rules on axon shape and size. Mathematical morphology operations are used to

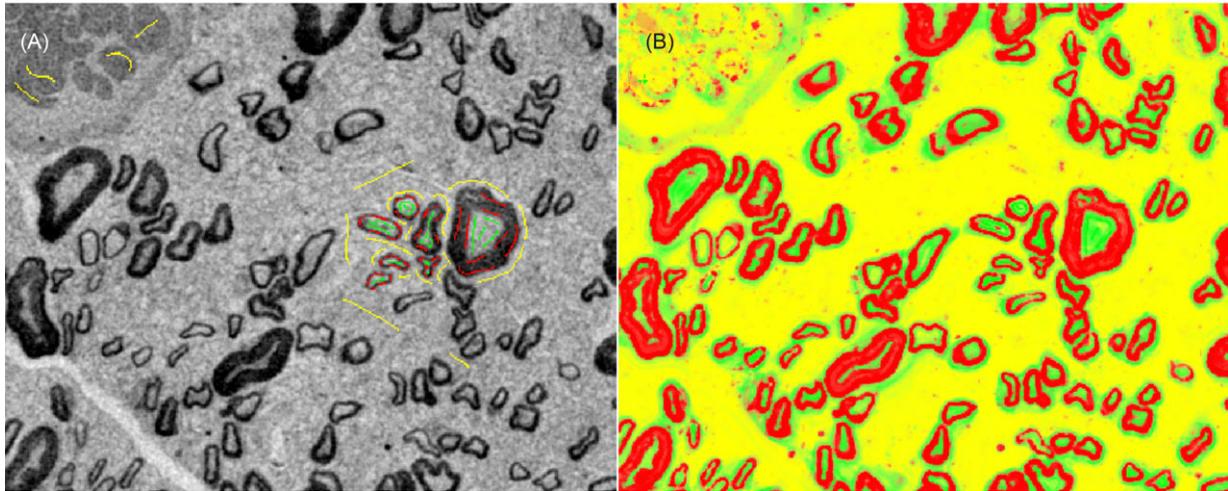


Fig. 3. (A) user annotations for three classes: myelinated membrane (red), axon interior (green) and background (yellow). (B) Probability map produced by the Random Forest classifier.

Table 1. 3D pixel features, used for the first step of the pipeline: pixelwise classification of the images into myelinated membranes, axon interior and background classes.

Feature	Gaussian σ	Channels
Gaussian smoothing	0, 1, 2, 3	1
Difference of Gaussians	1, 4, 6, 7	1
Eigenvalues of Hessian matrix	1, 2	3
Eigenvalues of structure tensor	1, 3	3
Total number of features		20

separate touching axons. Romero *et al.* (2000) use more sensitive locally adaptive thresholding as the first step. A comprehensive comparison of manual and semiautomated methods, developed up to 2007, has been performed by da Silva *et al.* (2007), which confirmed sufficient accuracy and reproducibility of semiautomated approaches.

A more recent contribution by Novas *et al.* (2013) proposes Arbib's colour space clustering for the initial segmentation, followed by Kumar's clump splitting for the separation of close axons. Although the methods above are applicable to stained tissue, Li *et al.* (2012) developed a protocol and a corresponding segmentation algorithm for unstained nerves, using molecular hyperspectral imaging. Use of scanning electron microscopy has been explored by More *et al.* (2011), who also propose a method for automated image stitching and segmentation, based on mathematical morphology.

Although the above-mentioned methods achieve fairly high accuracy and provide very significant speedups compared to purely manual axon segmentation, they work on individual 2D slices and, considering the necessary amount of user intervention, would not scale to the large three-dimensional (3D) datasets of our application domain.

Datasets of similar and much larger volume have been tackled as a part of automated neuron reconstruction efforts of recent years, aiming to segment all neurons in very high-resolution EM images of the brain (Jain *et al.*, 2011; Vazquez-Reina *et al.*, 2011; Andres *et al.*, 2012; Nunez-Iglesias *et al.*, 2013; Liu *et al.*, 2014; Funke *et al.*, 2014). Most of these methods use supervised machine learning to detect neural membranes and only require a small manually segmented dataset for training. The user is thus freed from the very nontrivial task of finding the right combination of thresholds and algorithm parameters, which would perform equally well throughout the volume. This feature makes such methods quite appealing also for our use case. On the other hand, all these methods have been developed to work on neuropil – volume of neural tissue, very densely packed with neurons and glia cells. Besides, they are tailored to images with specific staining, which gives contrast to all neural membranes, not only the myelinated ones. Consequently, these methods assume continuity of axons in the volume and cannot handle the gaps that are characteristic of the nodes of Ranvier in our target data.

We propose to base the myelinated axon segmentation on supervised pixel classification with 3D image features and sparse user labels, followed by the closing of gaps, which is formulated as a 2D assignment problem. The overall pipeline is summarized in Figure 1. The cost function of the assignment problem is learned from user annotations in the form of axon skeletons.

Assignment with a learned cost function has been introduced for multiple domains and applications, offline and online, with flat and structured learning (e.g. Li *et al.*, 2009; Lou & Hamprecht, 2011; Kim *et al.*, 2012; Yang & Nevatia, 2012). Specifically for neural data, it has been used by Funke *et al.* (2012) for connecting segmentations of individual serial sections into 3D objects. The general problem of multiple

object tracking has previously been formulated in the integer linear programming (ILP) framework by, for example Jiang *et al.* (2007), Zhang *et al.* (2008), Andriyenko & Schindler (2010), Berclaz *et al.* (2011). Our application is, however, considerably simpler, as we work with volumetric 3D data with appearances and disappearances, but no occlusions or divisions, where tracking reduces to a one-to-one assignment problem.

After solving the assignment problem and thus closing the axon gaps, we use the features of the closed gaps to detect the nodes of Ranvier by a Random Forest classifier (Breiman, 2001). Although we use machine learning for myelin segmentation, the gap closing procedure is agnostic to the segmentation method, and could also be used together with semiautomated neuromorphometry methods. Detection of the nodes of Ranvier partially relies on the probability map of axon interior. It could, however, be modified not to use this feature, if a nonprobabilistic method is used for axon segmentation.

The algorithm implementation is based on the *ilastik* learning and segmentation toolkit (Sommer *et al.*, 2011), the *vigra* image processing library (Koethe, 2008), the Gurobi solver (Gurobi Optimization, 2013) and the Mayavi data visualization toolkit (Ramachandran and Varoquaux, 2011). Apart from the Gurobi solver, all these tools are free and open source. The solver is free for academic use and for other uses it can simply be replaced by a different ILP solver or an efficient implementation of the Hungarian matching algorithm. Our software, along with the test dataset and the ground truth data, are freely available on our Website (<https://github.com/Rwalecki/ATMA>, <http://hci.iwr.uni-heidelberg.de/Benchmarks/>).

Methods

Our approach is comprised of pixel classification, followed by connected components analysis and reconnection of the remaining axon gaps, posed as an assignment problem. Pixel classification is based on very sparse user labelling and 3D pixel features. The costs in the assignment problem are learned from user annotations in the form of axon skeletons. Finally, the nodes of Ranvier are localized by another classification step.

Pixel classification

Volume pixels are classified into three classes: myelinated membranes, axon interiors and background. Very sparse user annotation of the images suffices for this step (see Fig. 3A for an example of user annotations). As a rule of thumb, users of *ilastik* usually continue to label until the interactive prediction no longer improves visually. Previous studies, performed on similar data with similar interactive classification methods (Kreshuk *et al.*, 2011; Kaynig *et al.*, 2013), have shown that the segmentation result is fairly stable to variations in user annotation.

Classification is performed based on pixel features, which represent intensity, edge and texture properties of the 3D pixel neighbourhood. The feature scales have been selected according to the variable importance of the features, computed as the Gini importance by the Random Forest classifier Breiman (2001). Variable importance computation, not yet present in *ilastik* version 1.1, has been performed in a script, using the *vigra* Random Forest library directly. The final list of features is summarized in Table 1. If the method is applied to data with different resolution, the scale of features has to be adjusted accordingly. Figure 3(B) shows the results of pixel classification.

Myelinated processes are large and smooth; neighbouring pixels of the volume are therefore likely to have the same label. We introduce the spatial regularization by filtering the myelin probability map with a 3D Gaussian filter (Hosni *et al.*, 2013). The smoothed probability map is then thresholded and connected component analysis is performed in each 2D binary image independently, assigning the interior of each connected myelin membrane to a connected component (axon slice). Axon slices larger than the user-defined maximum (see section 3.1c) are removed. The remaining axon slices are stacked in 3D to obtain a preliminary 3D axon reconstruction.

There are two ways to select the best values for the smoothing and thresholding parameters: interactively by visual inspection in a tool like *ilastik* (Sommer *et al.*, 2011) or by an optimization procedure on the user-provided ground-truth segmentation. In our experiments we used gradient descent optimization, maximizing the pixelwise segmentation *F*-score. More details on the parameter choice and the algorithm sensitivity can be found in section 3.2. An example slice with the resulting segmentation is shown in Figure 4.

In general, the aim of the pixel classification step is to produce a segmentation with as few erroneous axon merges as possible. The resulting gaps in the axons are reconnected in the next steps.

Assignment problem with a learned cost function

The segmentation obtained in the previous step does not fully reconstruct the axons. A single pixel error in the segmentation of a myelinated membrane can open its contour, which leads to its interior not being recognized as part of an axon. Besides, myelin sheaths are interrupted by nodes of Ranvier, which, since the nonmyelinated membranes are not stained, also introduce gaps into our segmentation (Fig. 5B). In this section, we propose to treat the problem of gap closing as a generalized assignment problem, both for gaps caused by pixel classification errors and by nodes of Ranvier. As a result of this step we expect to get fully reconstructed axons, whereas in the next section we introduce a method to distinguish between the erroneous gaps and the relevant gaps caused by nodes of Ranvier. We start by computing the skeletons from the centre points of each axon slice. Since our goal is reconstruction of fairly short pieces of peripheral nerves, we assume that

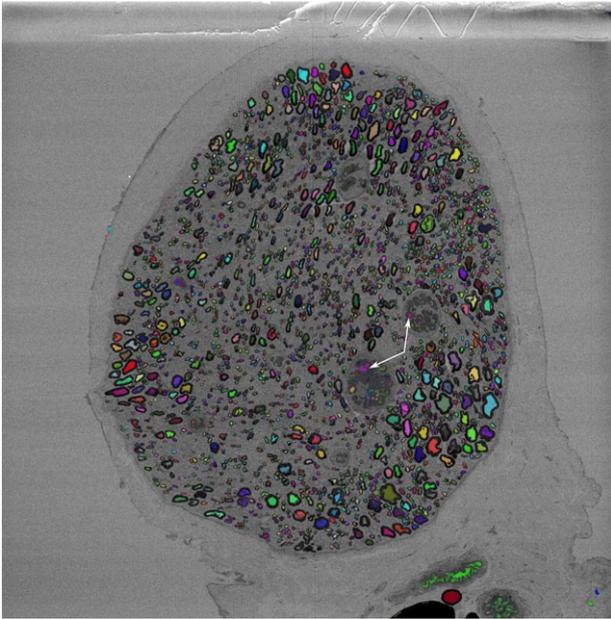


Fig. 4. Axon segmentation on a random image in the stack. In total, the algorithm segmented 2081 objects in this image, with 51 false positive detections (nonaxon objects) and 32 weakly stained axons missing. White arrows point to false positives, caused by stained structures other than myelinated axons.

axons in our target data do not branch, and the skeletons are represented by simple lines.

Closing the gaps in the axons can thus be formulated as pairwise matching of endpoints. To calculate the cost of pairwise assignments, we start by computing three attributes of the endpoints: position x , normalized orientation o and average axon thickness t . The position is the centre of mass of the end of the axon segment, defined as the last seven slices before the endpoint. Orientation is computed as the normalized shift over the last seven points of the axon segment skeleton. The axon segment thickness is obtained by calculating the average area of the last seven axon 2D slices. Although this feature is not rotation invariant, we do not expect the orientation of the axon to change significantly over the gap region, so our estimates of the axon areas at the two endpoints should match. After the endpoint attributes are computed, we use them to calculate the features of endpoint pairs (Fig. 5A), namely:

$$\begin{aligned} \text{Displacement: } & (x_{1i} - x_{1j}), (x_{2i} - x_{2j})(x_{3i} - x_{3j}), \\ \text{Distance: } & (x_{1i} - x_{1j})^2 + (x_{2i} - x_{2j})^2 + (x_{3i} - x_{3j})^2, \\ \text{Shape: } & t_i/t_j, \\ \text{Alignment: } & \vec{o}_i \cdot \vec{o}_j \text{ and } (\vec{r} \cdot \vec{o}_i) + (\vec{r} \cdot \vec{o}_j), \text{ where } \vec{r} = \frac{\vec{x}_i - \vec{x}_j}{\|\vec{x}_i - \vec{x}_j\|}. \end{aligned}$$

These features were used to train a Random Forest classifier to separate correct connections from incorrect ones. The training data for this step needs to be provided as a set of axon skeletons (see 3.1 for more details). Given this set, we introduce random gaps in the skeletons and treat the connec-

tions between pieces of the same skeleton as positive examples and the connections between pieces of different skeletons as negative examples.

The Random Forest produces an estimate of the probability p that two given endpoints should be connected. We then use $C = 1 - p$ as the cost of this connection in the assignment problem. We compute the full cost matrix of connections between the endpoints $C_{ij} = c(e_i, e_j)$, and solve the one-to-one assignment problem by the following integer linear program:

$$\begin{aligned} \text{minimize } & \sum_{i=1}^n \sum_{j=1}^n c_{ij} x_{ij}, \text{ subject to} \\ & \sum_{i=1}^n x_{ij} = 1, 1 \leq j \leq n, \\ & \sum_{j=1}^n x_{ij} = 1, 1 \leq i \leq n, \\ & x_{ij} \in \{0, 1\}. \end{aligned}$$

In the above formulation, we assume a perfect matching problem, where each endpoint should be connected with another one. Additionally, after the minimal cost solution of the linear program is computed, we remove matches over distances larger than a user-defined threshold (twice the length of an average Ranvier node, see section 3.1c). The above ILP can be solved by using the Hungarian method (Kuhn, 1955); however, in practice we found the Gurobi ILP solver to be faster than experimental code for the Hungarian method. Besides the speed issue, which probably results from a suboptimal Hungarian method implementation we used, the ILP solver and the Hungarian method give exactly the same results. A union-find algorithm is used to reconstruct full axons from pairwise endpoint matches.

Compensating for poor segmentation quality. In the experiments, where slices in the stack are imaged independently, artefacts cannot always be avoided, and some images may look significantly different from the rest of the stack. Automated algorithms, which rely on intensity-related features, are much more vulnerable to contrast fluctuations and histogram shifts than human experts, who segment based on higher level information. In our experiments with such nonhomogeneous image stacks, we observed a decrease in the quality of automated segmentation and an increase in the number of false merges. Unlike the false splits, this error cannot be corrected at the assignment stage. To compensate for this effect, we developed the following strategy:

- Run the segmentation and assignment steps of the algorithm once. Perform connected components analysis in the first and last image of the stack: collect all components that were separate in 2D, but received the same label in the assignment step. These are the potential false merges.
- Randomly select a merge and find the exact slices, where the two axons were segmented as one.

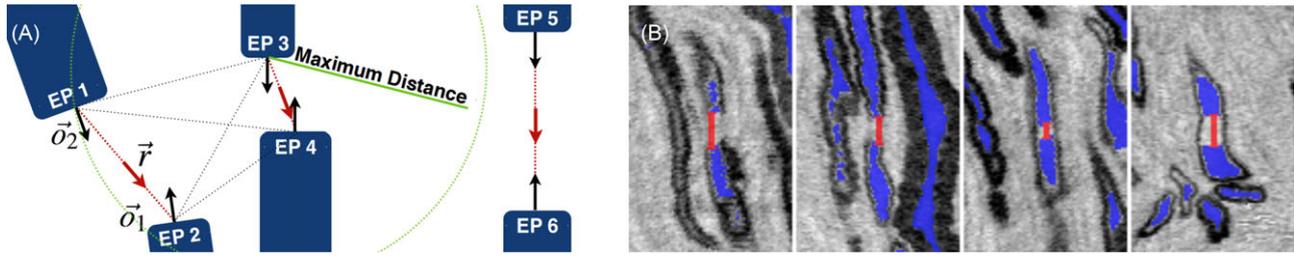


Fig. 5. (A) Gap closing costs. Blue tubes represent axons with endpoints EP 1–EP 6. Red lines are the cheapest connections according to our cost function. Black dotted lines show other connection candidates. A connection between EP 3 and EP 6 is not possible, since the distance is larger than the allowed threshold (twice the maximal expected length of a Ranvier node, section 3.1c). (B) Closed gaps: nodes of Ranvier and segmentation errors. First two images (left to right) show the nodes of Ranvier, whereas the last two images illustrate segmentation errors. The gaps are shown as thick red lines. Note the absence of a membrane around the gaps caused by the nodes of Ranvier.

- Remove these slices, rerun the assignment step of the algorithm and check how many axons can now be traced all the way through the stack. If more axons can be traced than before, keep the slices out of the stack, if not, put them back.
- Rerun the merge analysis and repeat the procedure until there are no merges left or no improvement in the number of fully traced axons can be reached by removing slices with merges.

Gap classification

Another classification step is used to separate the true axon gaps (nodes of Ranvier) from the segmentation errors (Fig. 5B).

For all gaps, closed in the previous step, i.e. for all the matched endpoint pairs, we look at the set of pixels X_{ij} on the straight line between the endpoints i and j . As the first feature, we compute for each point in X_{ij} its distance to the next myelinated membrane along the coordinate axes. This feature is similar to Ray features introduced in Smith *et al.* (2009), but much simpler and faster to compute. Briefly, we shoot six rays from each pixel along all coordinate axes and count the number of rays that reach a membrane within the average diameter of an axon, as predefined by the user (see section 3.1). The reasoning behind this feature is that true nodes of Ranvier are not surrounded by myelinated membranes, whereas gaps caused by segmentation errors usually have some pieces of membrane around them, not forming a closed contour. Besides Ray-like features, we compute the response of the Laplacian of Gaussian filter on the myelin probability map. This filter is widely used for detection of blobs and tubular structures (Lindeberg, 1993). Additionally, we use the axon interior probability map, produced at the pixel classification step. The following features are used for classification:

$$\text{Gap length: } (x_{1i} - x_{1j})^2 + (x_{2i} - x_{2j})^2 + (x_{3i} - x_{3j})^2,$$

$$\text{Axon thickness: } (t_i + t_j)/2,$$

$$\text{Alignment: } (\vec{r} \cdot \vec{o}_i) + (\vec{r} \cdot \vec{o}_j), \text{ where } r = \frac{\vec{x}_i - \vec{x}_j}{\|\vec{x}_i - \vec{x}_j\|},$$

$$\text{Ray features: } \min(Ry[X_{ij}]), \max(Ry[X_{ij}]), \text{mean}(Ry[X_{ij}]),$$

$$\text{LoG: } \min(LG[X_{ij}]), \max(LG[X_{ij}]), \text{mean}(LG[X_{ij}]),$$

$$\text{Axon probability: } \min(AP[X_{ij}]), \max(AP[X_{ij}]),$$

$$\text{mean}(AP[X_{ij}]).$$

The Random Forest classifier is then trained using labels from the third ground-truth dataset (section 3.1c). As the final result, we return a labelled volume, where a unique label is assigned to each axon, and a list of the locations of the nodes of Ranvier. The volumetric measurements of the segmented axons can then be analysed by Fiji (Schindelin *et al.*, 2012) or other image processing toolkits. Optionally, a list of the remaining nonmatched endpoints can also be returned for proofreading.

Results

In all of the following, we define an axon as fully traced through the volume if both of its endpoints are located within 50 slices of the first and last volume slice, respectively.

Data acquisition and ground truth annotation

Test datasets. The algorithm validation has been performed on two serial stacks of peripheral nerve images.

The first stack depicts the vagal nerve of an adult male rat. The resected nerves were fixed according to the procedure also used in Gierthmuehlen *et al.* (2013). Briefly, fixation in Karnovsky fixans (Karnovsky, 1965; Feria-Velasco & Karnovsky, 1970; 2% PFA, 2.5% glutaraldehyde in 0.2 M sodium cacodylate buffer, pH 7.3 for 24 h) prior to rinsing three times with 0.1 M sodium cacodylate buffer containing 7.5% sucrose. Postfixation was performed in 1% OsO4 for 1.5 h and myelin staining was performed for 24 h in 1% potassium dichromate, followed by 24 h in 25% ethanol and for

another 24 h in haematoxylin (0.5% in 70% ethanol) according to a modified Schultze protocol (Schultze, 1910). Finally, the samples were dehydrated and embedded in EPON (SERVA Electrophoresis GmbH, Heidelberg, Germany). Semithin ($1\ \mu\text{m}$) cross-sections were cut with glass knives (Ultramikrotome System, 2128 Ultratome[®], LKB, Bromma, Sweden) and mounted on uncoated glass slides followed by toluidine blue staining to further enhance the myelin staining (Huelsenbeck *et al.*, 2012).

The images were acquired on a serial block face scanning electron microscope Denk & Horstmann (2004), using the 3View microtome (GATAN) mounted on a QUANTA 200 VP-FEG (FEI). This technique was chosen for its high isotropic resolution and good intrinsic stack alignment. The block containing the vagal nerve perpendicular to the surface was cut out of the block and mounted on aluminium stubs for the sample holder of the 3View. The block was then trimmed on pyramidal shape and sputter-coated with gold. The surface of the block was imaged (6 kV beam, spot size 3, low vacuum mode with 0.25 Torr of water pressure in the chamber). The field of view was chosen to image the entire section of the vagal nerve. Each field of view represents a surface of $471 \times 471\ \mu\text{m}$ (3500×3500 pixels, $134\ \text{nm}$ pixel size, 5 ms per pixel). After imaging of the block face using the backscattered electron detector of 3View (GATAN), a diamond knife removed 200 nm from the surface of the block and the new surface was imaged using the same parameters. This cycle was repeated 7000 times. Images were then preprocessed to equalize the contrast and to be registered by TrakEM2 (Cardona *et al.*, 2012). The total size of the dataset measured 12 Gigabytes.

Since the images represent a peripheral nerve rather than a piece of brain tissue, all myelinated axons in the stack follow approximately the same direction (along the z -axis of the stack, see Fig. 2B). The study was conducted in strict accordance with the recommendations in the Guide for the Care and Use of Laboratory Animals of the National Institutes of Health. The protocol was approved through the Regierungspraesidium Freiburg (G13/44). The voxel size was set to $133 \times 133 \times 200\ \text{nm}^3$, with the total size of the dataset at $3995 \times 3995 \times 1000$ voxels.

To evaluate the generalizability of our approach to image stacks of lower resolution and quality, we also validated the algorithm on a cutaneous nerve branch dataset from Gierthmuehlen *et al.* (2013). The slices for this image stack were cut on an ultramicrotome at $0.8\ \mu\text{m}$ thickness and imaged individually on a confocal microscope at $384 \times 384\ \text{nm}$ resolution [see Gierthmuehlen *et al.* (2013) for more details]. The total size of this image stack is $1569 \times 1550 \times 625$ voxels. The dataset contains 338 myelinated axons. Besides the anisotropy of the resolution along the z -axis, there is also a strong fluctuation of contrast (see Fig. 8A). To partially compensate for these fluctuations, we have applied histogram homogenization from TrakEM2 (Cardona *et al.*, 2012) as pre-processing.

Ground truth. Out of the first image stack (the vagal nerve), we prepared three ground truth datasets for training and validation of the different parts of the algorithm. The first set consists of five fully labelled images, approximately 200×300 pixels each. In these images, we manually annotated axon membranes, axon interior and background. This dataset is used to study the sensitivity of the parameters of the pixel classification step.

The second ground truth dataset contains tracings of 166 axons in a $400 \times 400 \times 250$ data subvolume. The tracings have been obtained by skeletonization – the centres of each axon have been annotated in every tenth slice. This dataset is used for training the Random Forest to predict the likelihood of a connection between axon pieces.

The third ground truth dataset contains the locations of the nodes of Ranvier in a $700 \times 700 \times 700$ data subvolume. The following procedure has been used to create the dataset: first, the axons in the subvolume were segmented by the algorithm and the axon gaps were localized. Images of the gaps in three coordinate projections have been presented to a user, who classified them into nodes of Ranvier (208 cases) or segmentation errors (282 cases). Since all the gaps in this dataset have been re-examined by a human user, they can be closed correctly and the overall $700 \times 700 \times 700$ cube can serve for the validation of the gap-closing part of the pipeline.

For the second image stack, we prepared two more ground truth datasets, which consisted of 30 axons. First, the 30 axons were segmented in one image to create a dataset for the optimization of pixel classification parameters. Then, the same 30 axons were traced (by skeletons) through all images of the stack. Like the second vagus nerve ground truth dataset described above, this dataset has been used to train the Random Forest of the assignment cost function. Node of Ranvier detection has not been performed for the sciatic nerve dataset due to the difficulty of obtaining the training examples at this resolution.

To apply the algorithm to a new dataset, the user would have to provide a reasonable number of axon skeletons for training. Various tools, such as TrakEM2 (Cardona *et al.*, 2012) or Knossos (Helmstaedter *et al.*, 2013), have already been developed for this task. Pixel classification parameters, which we currently derive from the ground truth, can also be selected interactively in a tool like ilastik (Sommer *et al.*, 2011). Based on the skeleton training set, the algorithm will automatically choose images of axon gaps, which the user will have to classify into erroneous gaps and nodes of Ranvier. In our experience, the total training annotation time amounts to approximately 5 h.

User-defined parameters. Besides user annotations of the ground truth datasets, three parameters of the pipeline have to be defined in advance: the maximal axon cross-section area, the maximal connection distance (twice the maximal expected length of a Ranvier node) and the average axon diameter.

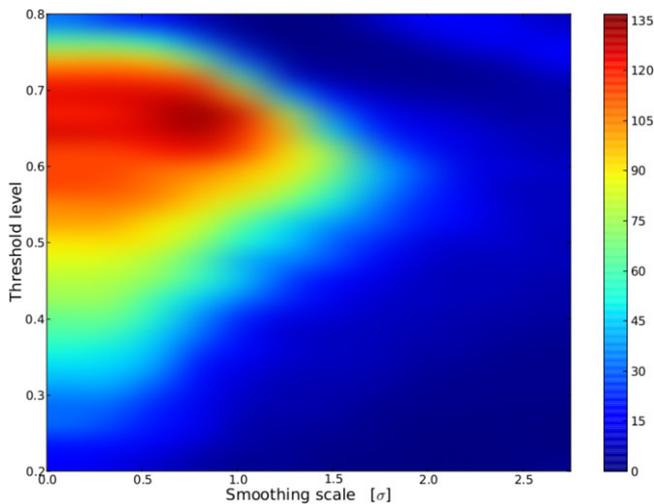


Fig. 6. Sensitivity of the algorithm to the selection of segmentation parameters. The colour bar shows the number of axons which could be traced continuously through an $800 \times 800 \times 800$ pixel subvolume.

These parameters depend on the resolution of the data and morphology of the nerve and are generally known ahead of time from neuroanatomy. For our test dataset the maximal axon cross-section area was set to 2000 pixels ($\sim 35 \mu\text{m}^2$), the maximal connection distance to 30 pixels ($\sim 4 \mu\text{m}$), and average axon diameter to 8 pixels ($\sim 1 \mu\text{m}$). For the second, lower resolution, image stack we changed the maximal connection distance parameter to 100 pixels, as we expected many more segmentation errors caused by lower image quality.

Segmentation

The first ground truth dataset (see 3.1.2) has been used to find the optimal values of the smoothing and thresholding parameters. Smoothing and thresholding are necessary to transfer from the pixel probability maps, produced by the Random Forest, to segmentations, used by the later steps of the algorithm. The maximum *F*-Score was 0.72 with $th = 0.5677$ and the smoothing sigma $\sigma = 1$ for the vagus nerve dataset. The corresponding values for sciatic nerve were estimated at $\sigma = 0.7$, $th = 0.31$. The sensitivity of the algorithm to these two parameters is illustrated in Figure 6, which shows the number of axons, traced fully through an $800 \times 800 \times 800$ pixel subvolume.

Full evaluation of the segmentation step would require pixelwise manual annotation of the volume. To avoid this extremely time-consuming step, we evaluated the segmentation on one slice and then counted the number of axons detected at other slices. Figure 4 shows all axons, detected in a random slice of the stack. This segmentation contains 2081 objects, with 51 false positive detections. Some of the false positives are located outside the nerve (lower right corner) and can be removed if a preprocessing step is used to localize the region

of interest. Some of the false positives inside the nerve belong to structures other than myelinated axons, shown by white arrows in Figure 4. Concerning false negative detections, we counted 32 axons which were missed by the algorithm on this slice.

Figure 7(A) shows the number of axons detected in each of the 1000 slices of the vagus nerve stack. The dips at both ends of the curve are expected, since we compute isotropic pixel features in 3D and at the ends of the stack the features do not see enough context along the *z*-axis. Otherwise, the number of axons segmented per slice stays fairly constant with a slight downward trend. This trend can at least partially be explained by the fact, that the myelin sheaths of some axons are truly not visible all the way through the volume. We found 10 such cases (myelinated axons ending inside the volume instead of continuously extending through it) by manual inspection of a random 100 axon endpoints that were left unmatched by our algorithm in the full $3995 \times 3995 \times 1000$ pixel dataset.

The lower curve in Figure 7(A) shows the number of axons detected in each of the 600 slices of the sciatic nerve dataset. Larger fluctuations in the segmentation quality are expected for this dataset, since the slices were imaged independently and the training annotation was kept very sparse (see also Fig. 8A for a side view of the image stack). Still, the majority of axons could be recovered in each slice (Fig. 8B displays the segmentation on a random slice inside the volume). The next section shows that, with the help of the compensation procedure, described in section 2.2.1, the tracing algorithm can successfully deal with lower quality segmentations of some slices.

Tracing

Vagus nerve dataset. The second ground truth dataset (axon skeletons, section 3.1.2) was used to train a Random Forest to predict which endpoints belong to the same axon. These predictions are used in the assignment cost function. We create artificial random gaps of variable length in the user-provided axon skeletons to obtain more training data. The maximum gap length was set to the value of the maximal possible connection distance (twice the maximal length of a Ranvier node), specified by the user. In total, we use 5100 pairs of endpoints for training: with 698 positive and 4402 negative connection examples.

After training, the gap closing procedure has been applied to a $700 \times 700 \times 700$ subvolume, not overlapping with the skeletonized second ground truth dataset. All gaps closed by the algorithm have been checked manually. Out of 502 gaps in total, 490 (97.6%) were closed correctly. The closed gaps have been manually classified into nodes of Ranvier and segmentation errors, thus creating the third ground truth dataset, described in section 3.1b. In total, we obtained 208 nodes of Ranvier and 294 segmentation error examples. In this dataset,

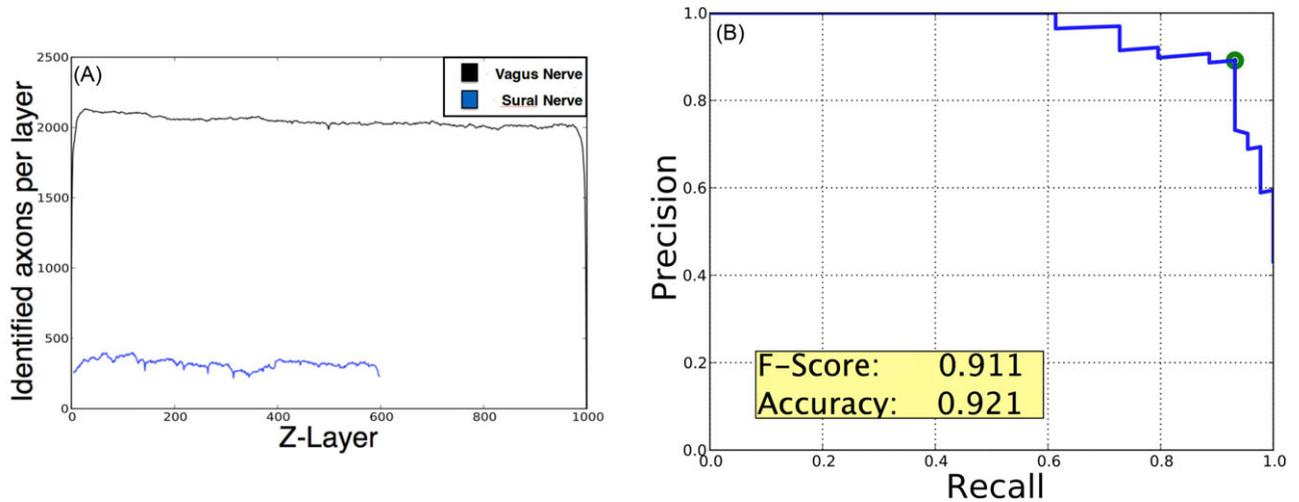


Fig. 7. Segmentation and gap classification performance. (A) The number of axon slices segmented by the algorithm in each image of the stack. (B) precision-recall curve for Ranvier node detection. The error estimates were obtained by 5-fold cross-validation on the third ground truth dataset (section 3.1c).

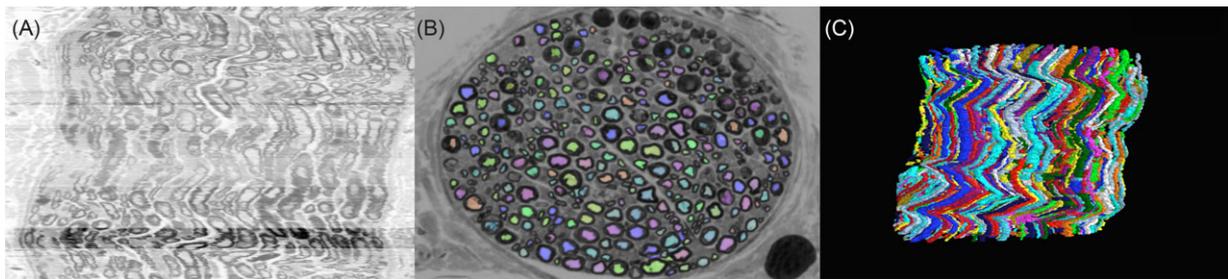


Fig. 8. From left to right: one slice of the rat sciatic nerve in the x - z projection, axon segmentation in a random image in the stack and a 3D rendering of the reconstructed axons.

241 out of 273 axons were traced through the whole volume (Fig. 9).

We compared the gap closing accuracy of our algorithm to greedily connecting the endpoints with the highest matching probability. Figure 10(A) shows this comparison, along with the baseline, obtained by connecting the endpoints based only on Euclidean distance between them. Figure 10(B) demonstrates an example of a connection, which was solved correctly by our algorithm and incorrectly by the greedy procedure. Such cases mostly occur in the 'crowded' regions of the volume, where many connections are possible and close in costs.

Cutaneous nerve branch dataset. First, we applied the slice removing procedure, described in section 2.2.1. In total, 17 slices of lower quality were removed and the performance of the assignment step improved significantly.

We traced 30 axons through the full volume to create the training set for the cost function Random Forest. The gap closing accuracy has been estimated by leave-one-out

cross-validation on the training set. Note that 96.2% of all gaps were closed correctly. Out of remaining 308 axons, we chose 100 of different sizes to cover all parts of the dataset. Fifty-five out of the 100 axons were traced correctly through the full volume (625 layers). A 3D rendering of all axons, reconstructed in this dataset, is shown in Figure 8(C).

Ranvier node detection

The closed gaps of the vagus nerve dataset were used to train another Random Forest to classify Ranvier nodes versus other membrane gaps. The 208 Ranvier nodes in the ground truth dataset have been labelled as positive and the 294 segmentation errors plus 12 wrong connections as negative. The resulting precision-recall curve for 5-fold cross validation is shown in Figure 7(B). The maximum F -Score (harmonic mean of precision and recall) is reached at the probability threshold level of 0.42, which corresponds to a true positive rate of 90.9% and a false negative rate of 6.9%.

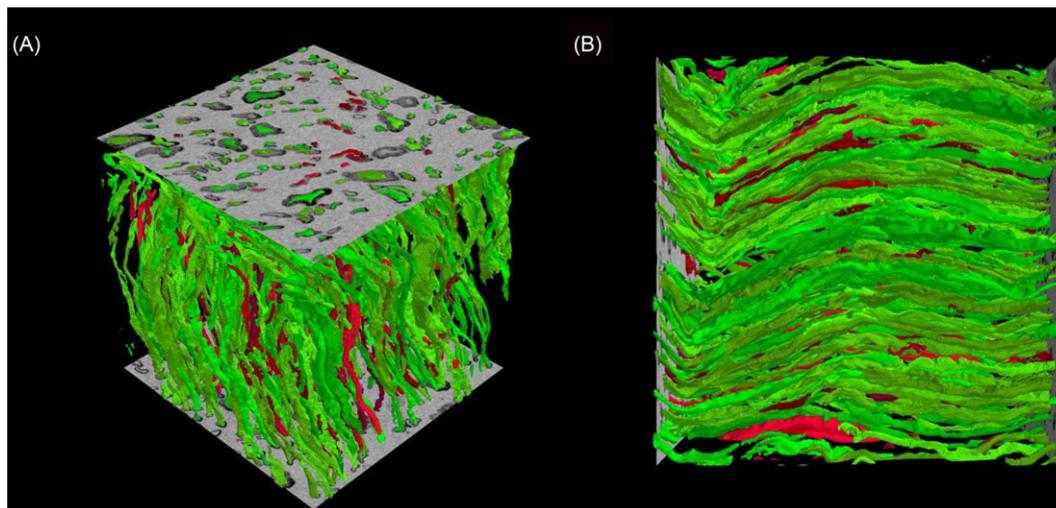


Fig. 9. Results on a randomly chosen subvolume. Axons which are fully traced through the dataset are shown in green. Those that have at least one open gap are shown in red.

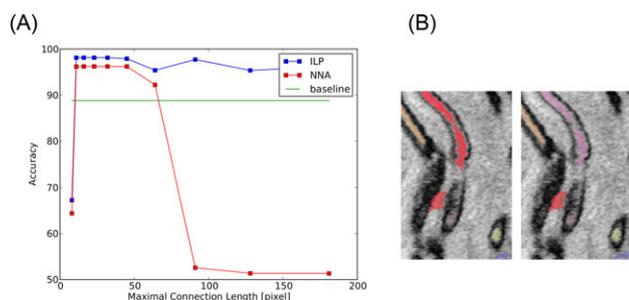


Fig. 10. (A) Accuracy of gap closing as a function of the maximal connection distance parameter. Accuracy is defined as the number of correctly closed gaps divided by the total number of gaps. Two methods are compared: globally optimal ILP solution and greedy connection based on the best matching score. The baseline corresponds to simply connecting the nearest neighbours by Euclidean distance. (B) An example of a gap, correctly connected in the ILP solution (right, the lower and upper part of the axon are purple), but incorrect in the greedy procedure results (left, the upper part of the red axon is connected to a segmentation error, also red).

Performance

On a 4-Core 2.80 Ghz machine, processing of the $700 \times 700 \times 700$ pixel dataset takes 99 min, with 24 min spent in segmentation. The most expensive step is the integer linear program used for the assignment problem. Seventy-one minutes were required to evaluate the assignment cost matrix. For comparison, we also solved this problem with the Hungarian Algorithm (Kuhn, 1955; experimental implementation, 104 min) and with simple greedy nearest neighbour assignment (6 min). Although faster, nearest neighbour assignment has lower precision than the proposed method due to its greedy nature.

The manual annotation of the training data from the sciatic nerve dataset took 5 h. For comparison, full manual

segmentation of this dataset in Gierthmuehlen *et al.* (2013) took 60 h, which corresponds to a 12-fold speedup for the automated method.

Since the maximal possible length of a gap in our algorithm is limited, the algorithm can easily be parallelized by cutting the dataset into blocks with a margin larger than the connection length limit (30 pixels in our case). Assuming uniform distribution of gaps in the volume, such parallelization can bring very considerable speedup to the solution of the matching problem. To obtain Figure 7(A), we processed the full $3995 \times 3995 \times 1000$ pixel dataset in 5003 blocks with 30 pixel halo, which took less than 6 h on the desktop machine described above.

Discussion

We have presented a method for automated tracing of myelinated axons and detection of the nodes of Ranvier in serial stacks of peripheral nerve images. After the training set is annotated, which, in our experience, takes 5 h, the algorithm works completely independently and can process very large stacks without any user involvement. Besides the training annotations, the algorithm is controlled by three user-defined parameters. These parameters are based on the biological properties of the studied nerve, so setting them does not require any image processing expertise.

The algorithm performance has been evaluated on two very different datasets: a high resolution, isotropic SBFSEM image volume and a lower resolution confocal microscope image stack. As expected, high and isotropic resolution, as well as intrinsic alignment of the SBFSEM images, allow for much better automated reconstruction. For this stack, 241 out of 273 axons could be traced through the 700 images of the test subvolume, which corresponds to a tremendous reduction of the

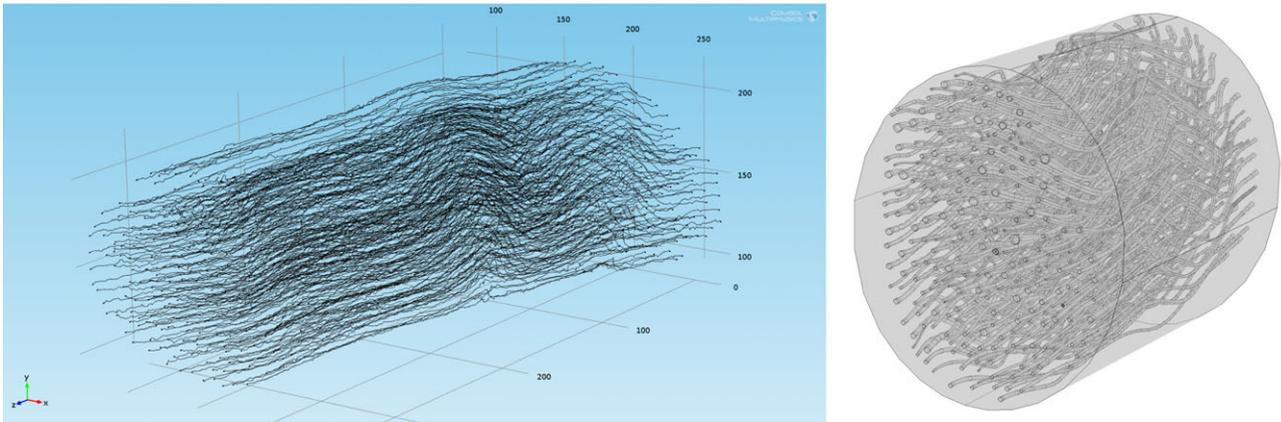


Fig. 11. Left: axon skeletons, as provided by the algorithm in a subvolume of the data. Right: nerve reconstruction, based on the skeletons and axon diameter values. Nodes of Ranvier are not shown.

manual annotation load. Even for the lower quality confocal stack, 55 out of 100 axons were traced fully. The algorithm also achieves very high accuracy in gap closing (97.6% for the SBFSEM stack, 96.2% for the confocal microscope stack), which implies, that even for the axons that were not fully traced, most ‘tracklets’ are correct and the human expert only needs to reconnect a few loose ends, as opposed to performing a full manual segmentation.

Figure 11 shows a part of the final algorithm results [skeleton traces through the volume, Fig. 11 (left)], which, combined with axon diameter and nodes of Ranvier positions, also provided by the algorithm, can be used to build a model of the nerve [Fig. 11 (right)].

It was only possible to evaluate the performance of the nodes of Ranvier detection on the high-resolution dataset. A fairly high accuracy (92.1%) has been achieved, however, we believe it can be further improved by, for example, incorporation of ‘Multiscale vessel enhancement’ features (Frangi *et al.*, 1998; Sato *et al.*, 1998). Another option would be to reformulate the problem in the framework of probabilistic graphical models and thus incorporate more prior information, such as the average expected distance between the nodes.

Next, we are planning to extend our algorithm to segment myelinated membranes in brain images, where the axons do not follow the same direction and the staining is not limited to myelin sheaths. By making both our algorithm and our test data with ground truth annotations publicly available, we hope to give a big boost to peripheral nerve simulation experiments and to bring the attention of the image processing community to this interesting data.

Acknowledgement

This work has been supported by the ‘Familie-Mehdorn-Stiftung’.

References

- Andres, B., Koethe, U., Kroeger, T., Helmstaedter, M., Briggman, K.L., Denk, W. & Hamprecht, F.A. (2012). 3D segmentation of SBFSEM images of neuropil by a graphical model over supervoxel boundaries. *Med. Image Anal.* **16**(4), 796–805.
- Andriyenko, A. & Schindler, K. (2010). Globally optimal multi-target tracking on a hexagonal lattice. *IEEE Eur. Conf. Comput. Vis.* **6311**, 466–479.
- Auer, R.N. (1994). Automated nerve fibre size and myelin sheath measurement using microcomputer-based digital image analysis: theory, method and results. *J. Neurosci. Methods* **51**(2), 229–238.
- Berclaz, J., Fleuret, F., Turetken, E. & Fua, P. (2011). Multiple object tracking using K-shortest paths optimization. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)* **33**(9), 1806–1819.
- Breiman, L. (2001). Random forests. *Machine Learning*, **45**, 5–32, Kluwer Academic, the Netherlands.
- Cardona, A., Saalfeld, S., Schindelin, J., *et al.* (2012). TrakEM2 software for neural circuit reconstruction. *PLoS One* **7**(6), e38011.
- da Silva, A.P.D., Jordao, C.E.R. & Fazan, V.P.S. (2007). Peripheral nerve morphometry: comparison between manual and semi-automated methods in the analysis of a small nerve. *J. Neurosci. Methods* **159**(1), 153–157.
- Denk, W. & Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol.* **2**(11), e329.
- Feria-Velasco, A. & Karnovsky, M.J. (1970). Optimal central nervous system preservation with glutaraldehyde perfusion for ultrastructural study [Article in Spanish]. *Arch. Invest. Med. (Mex)*, **1**(3), 201–220.
- Frangi, A.F., Niessen, W.J., Vincken, K.L. & Viergever, M.A. (1998). Multi-scale vessel enhancement filtering. *Med. Image Comput. Comput. Assist. Interv.* **1496**, 130–137.
- Funke, J., Andres, B., Hamprecht, F.A., Cardona, A. & Cook, M. (2012). Efficient automatic 3D-reconstruction of branching neurons from EM data. *IEEE Conf. Comput. Vis. Pattern Recognit.* 1004–1011, <http://dblp.uni-trier.de/rec/bibtex/conf/cvpr/FunkeAHCC12>.
- Funke, J., Martel, J., Gerhard, S., *et al.* (2014). Candidate sampling for neuron reconstruction from anisotropic electron microscopy volumes. *Med. Image Comput. Comput. Assist. Interv.* **8673**, 17–24.

- Gierthmuehlen, M., Freiman, T.M., Haastert-Talini, K., Mueller, A., Kaminsky, J., Stieglitz, T. & Plachta, D.T.T. (2013). Computational tissue volume reconstruction of a peripheral nerve using high-resolution light-microscopy and reconstruct. *PLoS One*, **8**(6), e66191. doi: 10.1371/journal.pone.0066191.
- Gurobi Optimization (2013). Gurobi Optimizer Reference Manual. <http://www.gurobi.com/documentation/6.0/refman/index.html>
- Helmstaedter, M., Briggman, K.L., Turaga, S.C., Jain, V., Seung, H.S. & Denk, W. (2013). Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature* **500**(7461), 168–174.
- Hosni, A., Rhemann, C., Bleyer, M., Rother, C. & Gelautz, M. (2013). “Fast Cost-Volume Filtering for Visual Correspondence and Beyond,” *Pattern Analysis and Machine Intelligence*, *IEEE Transactions on*, **35**(2), 504–511. doi: 10.1109/TPAMI.2012.156.
- Huelsenbeck, S.C., Rohrbeck, A., Handreck, A., et al. (2012). C3 peptide promotes axonal regeneration and functional motor recovery after peripheral nerve injury. *Neurotherapeutics* **9**, 185–192.
- Hunter, D.A., Moradzadeh, A., Whitlock, E.L., Brenner, M.J., Myckatyn, T.M., Wei, C.H., Tung, T.H.H. & Mackinnon, S.E. (2007). Binary imaging analysis for comprehensive quantitative histomorphometry of peripheral nerve. *J. Neurosci. Methods* **166**(1), 116–124.
- Kaynig, V., Vazquez-Reina, A., Knowles-Barley, S., et al. (2013). Large-scale automatic reconstruction of neuronal processes from electron microscopy images. arXiv:1303.7186v1[q-bio.NC].
- Karnovsky, M.J. (1965). A formaldehyde-glutaraldehyde fixative of high osmolarity for use in electron microscopy. *J. Cell. Biol.* **27**, 137–8a.
- Kim, S., Kwak, S., Feyereisl, J. & Han, B. (2012). Online multi-target tracking by large margin structured learning. *Asian Conf. Comput. Vis., Lecture Notes Comput. Sci.* **7726**(2013), 98–111.
- Koethe, U. (2000). Generische Programmierung für die Bildverarbeitung, Universität Hamburg, Hamburg 2000, ISBN: 3-8311-0239-2.
- Kreshuk, A., Straehle, C.N., Sommer, C., Koethe, U., Cantoni, M., Knott, G. & Hamprecht, F.A. (2011). Automated detection and segmentation of synaptic contacts in nearly isotropic serial electron microscopy images. *PLoS One* **6**(10), e24899.
- Kuhn, H.W. (1955). The Hungarian method for the assignment problem. *Naval Res. Logis. Quart.* **2**(1–2), 83–97.
- Li, Q., Chen, Z., He, X., Wang, Y., Liu, H. & Xu, Q. (2012). Automatic identification and quantitative morphometry of unstained spinal nerve using molecular hyperspectral imaging technology. *Neurochem. Int.* **61**(8), 1375–1384.
- Li, Y., Huang, C. & Nevatia, R. (2009). Learning to associate: hybrid-boosted multi-target tracker for crowded scene. *IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2953–2960. http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=5206735&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxppls%2Fabs_all.jsp%3Farnumber%3D5206735.
- Lindeberg, T. (1993). Detecting salient blob-like image structures and their scales with a scale-space primal sketch: a method for focus-of-attention. *Int. J. Comput. Vis.* **11**(3), 283–318.
- Liu, T., Jones, C., Seyedhosseini, M. & Tasdizen, T. (2014). A modular hierarchical approach to 3D electron microscopy image segmentation. *J. Neurosci. Methods* **226**, 88–102.
- Lou, X. & Hamprecht, F.A. (2011). Structured learning for cell tracking. *Adv. Neural. Inf. Process. Syst.* **24**, 1296–1304.
- More, H.L., Chen, J., Gibson, E., Donelan, J.M. & Beg, M.F. (2011). A semi-automated method for identifying and measuring myelinated nerve fibers in scanning electron microscope images. *J. Neurosci. Methods* **201**(1), 149–158.
- Novas, R.B., Paula, V., Fazan, S., Felipe, J.C., Preto, R. & Paulo, S. (2013). Fully-automatic tool for morphometric analysis of myelinated fibers. *IEEE Int. Symp. Comput.-Based Med. Syst.* 161–166. doi: 10.1109/CBMS.2013.6627782.
- Nunez-Iglesias, J., Kennedy, R., Parag, T., Shi, J. & Chklovskii, D.B. (2013). Machine learning of hierarchical clustering to segment 2D and 3D images. *PLoS One* **8**(8), e71715.
- Plachta, D.T.T., Gierthmuehlen, M., Cota, O., Boeser, F. & Stieglitz, T. (2013). Baroloop: using a multichannel cuff electrode and selective stimulation to reduce blood pressure. In *Proceedings of the 35th Annual International Conference of IEEE EMBS*, IEEE, Osaka, 755–758.
- Plachta, D.T.T., Gierthmuehlen, M., Cota, O., Espinosa, N., Boeser, F., Herrera, T.C., Stieglitz, T. & Zentner, J. (2014). Blood pressure control with selective vagal nerve stimulation and minimal side effects. *J. Neural Eng.* **11**(3), 036011. doi:10.1088/1741-2560/11/3/036011.
- Ramachandran, P. & Varoquaux, G. (2011). Mayavi: 3D visualization of scientific data. *Comput. Sci. Eng.* **13**(2), 40–51. doi: 10.1109/MCSE.2011.35.
- Romero, E., Cuisenaire, O., Denef, J.F., Delbeke, J., Macq, B. & Veraart, C. (2000). Automatic morphometry of nerve histological sections. *J. Neurosci. Methods* **97**(2), 111–22.
- Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G. & Kikinis, R. (1998). Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Med. Image Anal.* **2**, 143–168.
- Schindelin, J., Arganda-Carreras, I., Frise, E., et al. (2012). Fiji: an open-source platform for biological-image analysis. *Nat. Methods* **9**(7), 676–682.
- Schultze, O. (1910). Über die Anwendung von Aminosäure und eine neue Osmiumhämatoxylinmethode. *Z. Wiss. Mikrosk.* **27**, 465–475.
- Smith, K., Carleton, A. & Lepetit, V. (2009). Fast ray features for learning irregular shapes. *IEEE Int. Conf. Comput. Vis. (ICCV)*, 397–404, 10.1109/ICCV.2009.5459210.
- Sommer, C., Straehle, C., Koethe, U. & Hamprecht, F.A. (2011). Ilastik: interactive learning and segmentation toolkit. *2011 IEEE Int. Symp. Biomed. Imaging: From Nano to Macro* (1), 230–233. doi: 10.1109/ISBI.2011.5872394.
- Vazquez-Reina, A., Gelbart, M., Huang, D., Lichtman, J., Miller, E. & Pfister, H. (2011). Segmentation fusion for connectomics. *IEEE Int. Conf. Comput. Vis. (ICCV)*, 177–184. doi: 10.1109/ICCV.2011.6126240.
- Vita, G., Santoro, M., Trombetta, G., Leonardi, L. & Messina, C. (1992). A computer-assisted automatic method for myelinated nerve fiber morphometry. *Acta Neurol. Scand.* **85**(1), 18–22.
- Yang, B. & Nevatia, R., (2012). An online learned CRF model for multi-target tracking. *IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2034–2041. doi: 10.1109/CVPR.2012.6247907.
- Zhang, L., Li, Y. & Nevatia, R. (2008). Global data association for multi-object tracking using network flows. *IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 1–8. 10.1109/CVPR.2008.4587584.