

Multiclass U-Net Segmentation of Brain Electron Microscopy Data

Alexandra Getmanskaya¹, Nikolai Sokolov¹ and Vadim Turlapov¹

¹UNN, Gagarina av. 23, N.Novgorod, 603950, Russian Federation

Abstract

This work focuses on multi-class labeling and segmentation of electron microscopy data. The well-known and state-of-the-art EPFL open dataset has been labeled for 6 classes (instead of 1) and a multi-class version of the U-Net was trained. The new labeled classes are mitochondrion together with its border, mitochondrion's border (separately), membrane, PSD, axon, vesicle. Our labeling results are available on GitHub. Our study showed that the quality of segmentation is affected by the presence of a sufficient number of specific features that distinguish the selected classes and the representation of these features in the training dataset. With 6-classes segmentation, mitochondria were segmented with the Dice index of 0.94, which is higher than with 5-classes (without mitochondrial boundaries) segmentation (Dice index of 0.892).

Keywords

multi-class segmentation, electron microscopy, neural network, image segmentation, machine learning

1. Introduction

Information about the anatomy and connectivity of neurons can provide new insights into the relation between the brain's structure and its function [1]. Such information may also provide insights into the physical underpinnings of common serious disorders of brain function such as mental illnesses and learning disorders, which at present have no physical trace. Furthermore, information about the individual strength of synapses or the number of connections between two cells has important implications for computational neuroscience and theoretical analysis of neural networks [2].

Microscopy plays an indispensable role in biomedical research. Advanced technologies of microscopy like electron microscopy have opened up new eyesight for biomedical researchers. The image resolution is very high, which is why a cubic millimeter of brain tissue can take up more than 1000 terabytes. The resulting images are analyzed to identify individual cells. Segmentation is usually done by biologists manually. Processing one experiment takes up to six months of manual work.

Since the proposed method for segmentation of neural membranes using deep neural networks (DNN) [3], many DNN-based methods have been developed, such as U-Net [4], which is one of

GraphiCon 2021: 31st International Conference on Computer Graphics and Vision, September 27–30, 2021, Nizhny Novgorod, Russia

✉ getmanskaya.alexandra@gmail.com (A. Getmanskaya); nikcsokolof@mail.ru (N. Sokolov); vadim.turlapov@gmail.com (V. Turlapov)

ORCID 0000-0003-3533-1734 (A. Getmanskaya); 0000-0002-9055-2225 (N. Sokolov); 0000-0001-8484-0565 (V. Turlapov)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

the most successful in the field of biomedical data segmentation. The main idea of U-Net is to include local and more extensive contextual information (global context) from the input image. Based on this concept, a variety of DNN architectures have been proposed for volume data.

A similar idea is used by the deep contextual network [5], emphasizing the importance of capturing a sufficient receptive field. Both networks are based on fully convolutional networks [6], and feature maps obtained in the middle layer are propagated directly to deep layers using a skip connection. Drozdal et al. have also experimentally shown that the skip connection is essential for the segmentation of nerve cells [7]. Further research has developed deeper networks [8, 9] and mechanisms that take into account correlation in three-dimensional space. Xiao et al. [9] proposed a 3D U-Net with residual blocks. On the task of segmentation of mitochondria, they received the quality jaccard = 91.8%.

The next step was the use of three-dimensional convolutions in neural networks. V-Net[10], 3D U-Net[11], DeepMedic[12], HighRes3DNet[13] are different architectures using 3D convolutions.

In parallel with the development of three-dimensional architectures, methods of preprocessing input data and post-processing of segmentation results obtained by neural networks were proposed. Preprocessing and post-processing increases the neural networks segmentation quality about 5-9% each [14, 15, 16]. For example, Manca Žerovnik Mekuč et al. [17] used contrast enhancement based on the adaptive gamma correction with weight distribution (AGCWD) [18].

One of the limitations of 3D CNN is the significantly increased number of training parameters with a 3D convolution kernel, which leads to high computational costs and high GPU memory consumption.

Therefore, architectures that use 3D convolutions have been replaced by architectures that reduce the number of training parameters, adjusting the balance between the quality of networks with 3D convolutions and the speed of training two-dimensional convolutions.

Low-rank factorization of convolutional kernels [19, 20, 21] and hierarchical convolution (HVEC)[22] are the simple alternative to 3D convolution for exploring 3D spatial context.

Biological and medical data is characterized by a small amount of labeled data. And the publicly available electron microscopy data as a whole are presented in only a few volumes due to the laboriousness of preparing a tissue for an electron microscope and due to the need for post-processing of the obtained images. The laboriousness of annotation results in even fewer public labeled electron microscopy datasets.

We found eight publicly available electron microscopy data, six labeled in only 1 class (mitochondrion or membranes). And only in two volumes several classes are marked. Therefore, the vast majority of neural networks in electron microscopy are trained for only two classes. Therefore, it is important to create multi-class markup and examine the results of multi-class architectures using this markup.

2. Materials and methods

2.1. Open datasets

This section provides information on public datasets. The most widely used datasets for assessing mitochondrial segmentation have been provided by Lucchi et al. [23] and Xiao et al. [9].

Mouse neuropil. The mouse neuropil dataset [24] includes 400 images with a size of 4096×4096 with a resolution of $10 \times 10 \times 50$ nm / voxel. A subset of 70 images with a size of 700×700 was selected from the dataset and marked up by experts.

Mouse cerebral cortex dataset. Also known as AC4 dataset. The entire mouse cortical dataset is a stack of 1,850 images of 4096×4096 pixels with a resolution of $3 \times 3 \times 30$ nm / voxel. The images were compressed 2x in x - y, and two subsets of $1024 \times 1024 \times 100$ pixels were cropped out and used in the ISBI 2013 EM Challenge [25] as a training and testing set. For the training sample, 100 images were provided 2D and 3D marked by experts.

EPFL dataset (Lucchi dataset [23]). The dataset is a $5 \times 5 \times 5$ μm section taken from the CA1 hippocampus region of the brain, which corresponds to a volume of $1065 \times 2048 \times 1536$. The resolution of each voxel (vx) is approximately $5 \times 5 \times 5$ nm. Mitochondria were annotated in two parts of the dataset. Each chunk consists $1024 \times 768 \times 165$ image stack.

Lucchi ++ Mitochondrial Segmentation Dataset.

This dataset is based on the hippocampal EPFL dataset [23]. The difference lies in the more accurate mitochondrial markings, in which the senior biologist manually corrected the markings of the mitochondrial membranes using his own annotation software along with 2 neuroscientists.

Kasthuri ++ Somatosensory Cortex. Contains annotations of the mitochondria of the 3-cylinder volume of the mouse cortex Kasthuri et al. The tissue is a dense mammalian neuropil from layers 4 and 5 of the primary somatosensory cortex S1, obtained using serial sectional electron microscopy (ssEM). In this data, membrane inconsistencies in mitochondrial segmentation masks were corrected by experts similar to the Lucchi ++ dataset. The dimensions of the stack are $1463 \times 1613 \times 85\text{vx}$ and $1334 \times 1553 \times 75\text{vx}$ with a resolution of $3 \times 3 \times 30$ nm per voxel.

chm-supplemental data. Contains SBEM training data of the suprachiasmatic nucleus (SCN) of one 3-month-old mouse (images and labels of mitochondria, lysosomes, nuclei, isotropic nuclei and nucleoli). This dataset was used by Perez et al. [26].

UroCell [17]. The open volumetric EM dataset is the first of the urothelial cells. The tissue samples was taken from urinary bladders of 6–8 weeks old healthy male C57BL/6J mice. The dataset consists of 1056 consecutive layers of 1366×1180 pixels. Voxel sizes in the dataset are approximately $x = 16$ nm, $y = 16$ nm, $z = 15$ nm — nearly isotropic resolution in all three directions. Intracellular compartments (mitochondria and endolysosomes) are manually labeled in 5 sub-volumes of $256 \times 256 \times 256$ voxels. To increase variability, the selected annotated sub-volumes are taken from different parts of the entire volume and are therefore varied in terms of contrast, brightness, artifacts and content. The data is located on GitHub: <https://github.com/MancaZerovnikMekuc/UroCell>.

ISBI 2012 dataset [25]. The dataset is a set of 30 slices from the sequential transmission electron microscopy dataset of the ventral nerve circuit of the first stage larva of *Drosophila*, which was used in a competition held at the ISBI 2012. The competition task was to find the membranes of nerve cells. Membranes are marked in the images. The imaged volume measures $2 \times 2 \times 1.5$ μm , with a resolution of $4 \times 4 \times 50$ nm/pixel. The result resolution is the 4×4 nm per pixel.

You can see that six labeled open datasets in only 1 class. And only in two volumes several classes are marked. Therefore, the vast majority of neural networks in electron microscopy are trained for only two classes (object and background).

2.2. Dataset

The initial data was obtained from the site <https://www.epfl.ch/labs/cvlab/data/data-em/> and are called “EPFL dataset” or “Lucchi mitochondrial segmentation dataset”. This data originally contains masks only for mitochondria. Therefore, for the analysis and evaluation of multiclass segmentation algorithms, we manually labeled 14 layers (1024x768) for the following classes:

1. Mitochondrion together with the border
2. Mitochondrion’s border
3. Membranes
4. PSD
5. Axon sheaths
6. Vesicles

Precise marking of 1 layer by hand took about 5 hours. Our markup of the EPFL dataset is available at <https://github.com/GraphLabEMproj/unet>. We plan to continue working on the markup and mark up both of the available volumes. You can see an example of the image patch markup in Figure 1.




Figure 1: An example of our image markup: a) original image, b) membranes, c) mitochondrion with borders, d) mitochondrion’s borders, e) PSD, f) vesicles

2.3. Network architecture

U-Net is considered one of the standard CNN architectures for image segmentation tasks. The architecture consists of a constricting path to capture the global context and a symmetrical expanding path that allows precise localization. For the basis of the neural network, the U-Net project was taken <https://github.com/zhixuhao/unet>. In the original project, U-Net was used for the binary classification of membranes. In our research, we use U-Net for multi-class segmentation. We forked the original repository, all changes in the code and our markup of the Lucci data is available at <https://github.com/GraphLabEMproj/unet>.

Following the author of the code <https://github.com/zhixuhao/unet> in U-Net implementation, there are differences from the classical U-Net network [4]:

- The network input is an image reduced to the size 256x256x1.
- The network output is 256x256xnum_classes, where num_classes is the number of classes.
- The sigmoid activation function ensures that the mask is in the range [0, 1].

Also, we added batch normalization after each convolution and the ReLU activation layers.

3. Experiments and results

3.1. Evaluation criteria

For evaluation, we use the Dice similarity coefficient (DSC) and Jaccard index, both commonly used in the field of medical image segmentation. If we define TP to be the number of true positive voxels (correctly identified target class), FP the number of false positive voxels (target class identified on background), FN the number of false negative voxels (missed target) and TN the number of true negative voxels (correctly classified background), we can define the metrics as follows. The Dice similarity coefficient measures the similarity between annotations and predictions and is defined as:

$$DSC = \frac{2TP}{2TP + FP + FN}$$

The Jaccard index measures the same as Dice coefficient:

$$JAC = \frac{TP}{TP + FP + FN}$$

The Dice values, like Jaccard's, range from zero to one. Unlike Jaccard, the corresponding difference function is not a correct distance metric, since it does not satisfy the triangle inequality. Jaccard and Dice are equivalent in the sense that one can be expressed through one another:

$$DSC = \frac{2JAC}{1 + JAC}$$

Since in our study we consider multi-class segmentation, we are interested in multi-class metrics. Since the Jaccard distance (or Dice) compares 2 sets, in the case of a multiclass classification, the result will be a vector of Jaccard (or Dice) distances for each class. When training a neural network, a function that returns a scalar value is used to calculate the error. Therefore, for multiclass segmentation, it is necessary to convolve the distance vector. To convolve a vector into a scalar we use linear convolution:

$$W_{scalar} = \sum_{i=1}^N \lambda_i W_i, \lambda_i \geq 0, \sum_{i=1}^N \lambda_i = 1$$

where λ_i is the weight coefficient, and W_i is the value of the distance coefficient for the i -th class. W_{scalar} - scalar value or convolution of the distance vector. N is the number of classes.

In this work, the linear convolution Dice coefficients with weight coefficients λ_i equal to $1/N$ is used for the error function.

Two metrics used in ISBI 2012 are: Maximal foreground-restricted Rand score after thinning (V_{rand}) and maximal foreground-restricted information theoretic score after thinning (V_{info}). For a detailed description of the metrics, please refer to [25].

Also, for membrane detection used metric Rand error (RE): $1 -$ the maximal F-score of the foreground restricted rand index (Rand 1971), a measure of similarity between two clusters or segmentations. For the EM segmentation evaluation, the zero component of the original labels (background pixels of the ground truth) is excluded.

Table 1
Evaluation of the results

Label	Class number					
	6		5		binary	
	Jaccard	Dice	Jaccard	Dice	Jaccard	Dice
mitochondrion	0.887	0.940	0.806	0.892	0.913	0.954
mitochondrial borders	0.673	0.805	-	-	-	-
membranes	0.681	0.810	0.737	0.849	-	-
PSD	0.729	0.843	0.625	0.769	-	-
vesicles	0.618	0.764	0.616	0.762	-	-
axon	0.773	0.872	0.825	0.904	-	-

3.2. Experiments

We cut marked up 7 slices into 256x256 images, and obtain as the result 301 patches. The data augmentation was used for model training. In the training set, were used 266 patches, and 35 in the verification set, the batch size was equal to 7. We test the model on 1 slice (12 patches). We used random sub-sampling cross-validation. We use Adam as an optimizer with a learning rate set to 10^{-4} . Each training experiment was run for 100 epochs. Experiment’s learning curves you can see in Figure 2.

Experiment 1. Segmentation classes (5): mitochondrion together with the border, membranes, PSD, axon sheaths, vesicles.

Experiment 2. Segmentation classes (6): mitochondrion together with the borders, mitochondrion’s border, membranes, PSD, axon, vesicles.

Added 1 more class of mitochondrial borders.

Experiment 3. Segmentation class: Mitochondrion together with the border. Segmentation into class 1 mitochondrion is used.




Figure 2: Experiment’s learning curves: a) 5-class learning curve, b) 6-class learning curve, c) binary mitochondrial segmentation learning curve

As we can see from the comparison table (See table 1), multi-class segmentation of mitochondria is not much inferior in quality to binary segmentation.

The mitochondrial border class is a subclass of “mitochondrion” and further emphasizing the

border improves the segmentation results of the unifying class.

Of all classes, vesicles show the worst results due to their small size and due to single vesicles that can be confused with noise (Figure 3). Perhaps the use of three-dimensional convolutions will smooth this effect. Possibly combining adjacent vesicles into a "vesicle region" will improve the results.

The network was trained in unbalanced classes, because the size of the compartments and their occurrence in the layer differ tens of times.

Despite the fact that the smallest in area classes are PSD and axons, their recognition is higher than that of vesicles.




Figure 3: Images: a) original vesicles, b) ground truth, c) ground truth and 5-class prediction, d) ground truth and 6-class prediction

4. Discussion

In the discussion section, we present two tables: Table 2: "Comparison mitochondrion segmentation results" and Table 3: "Comparison membrane segmentation results". We have placed here the most illustrative results on the segmentation of two classes: mitochondria and membranes using binary and multiclass models.

Table 2

Comparison mitochondrion segmentation results

Method	Classes	Article	Dataset	Jaccard	Dice
Cheng et al. (3D) [27]	1	M. Yuan et al [22]	EPFL	0.889	0.941
3D U-Net [11]	1	M. Yuan et al [22]	EPFL	0.878	0.935
Cheng et al. (2D) [27]	1	M. Yuan et al [22]	EPFL	0.865	0.928
U-Net [4]	1	M. Yuan et al [22]	EPFL	0.844	0.915
U-Net	6	ours	EPFL	0.83	0.906
V-Net	2	M. Žerovnik Mekuč et al [17]	UroCell	-	0.898
U-Net	5	ours	EPFL	0.812	0.895
U-Net	1	ours	EPFL	0.789	0.880
HighRes3DNet	2	M. Žerovnik Mekuč et al [17]	UroCell	-	0.883
HighRes3DNet	1	M. Žerovnik Mekuč et al [17]	UroCell	-	0.862
Lucchi et al. [28]	1	M. Yuan et al [22]	EPFL	0.755	0.86
U-Net	2	M. Žerovnik Mekuč et al [17]	UroCell	-	0.855




Figure 4: Mitochondria segmentation problem: a) prediction 5-class mitochondria, b) original image, c) ground truth

We test our models on full test EPFL volume and use this values instead of Table 1 results. We cannot directly compare the results from the table 2 because our models were trained on a significantly reduced version of the EPFL dataset. But we can put forward several hypotheses that need to be tested. The worst results were obtained in the layers containing the axon, parts of the mitochondria cut by the border of the image, fuzzy membranes (Figure 4). The network confused the axon inner region with the mitochondria. We assume the causes are 1) the small number of axons represented in the training dataset; 2) we marked only the boundaries of the axon without the inner region, relying on their brightness as a sufficient feature.

When we added the “mitochondrion’s border” class, the recognition of mitochondria in the 6-class model increased from 0.895 to 0.906 (Table 2), while the recognition of the border

decreased from 0.849 to 0.81 (Table 1). But it was the only one negative issue only. It means that it is possible to pick up classes to get segmentation comparable to the segmentation of a binary model. A multi-class 3D model gives an increase in accuracy as well as a binary model.

Table 3
Comparison membrane segmentation results

Method	Classes	Article	Dataset	Jaccard	Dice	Rand Error	V_{rand}
U-Net	5	ours	EPFL	0.737	0.849	0.134	0.992
FusionNet	1	Quan et al. [29]	ISBI 2012	-	-	-	0.978
M^2FCN	1	Shen et al. [30]	ISBI 2012	-	-	-	0.978
U-Net	6	ours	EPFL	0.681	0.81	0.267	0.952
U-Net	1	Ronneberger et al. [4]	ISBI 2012	-	-	0.0382	0.973
Drozdal et al	1	Drozdal et al. [7]	ISBI 2012	-	-	-	0.969

Table 3 presents the results for two different datasets (EPFL and ISBI 2012) but one membrane class. And if you visually compare the two datasets, you can see that the ISBI 2012 has much more noise in the data, it is worse aligned in brightness and there are more membranes cut defects. The EPFL dataset quality allows to see that the membrane consists of two adjacent boundaries of different cells.

5. Conclusion

In this paper, we presented a new multi-class labels for the state-of-the-art EPFL dataset, which includes such classes as mitochondrion including the borders, mitochondrial boundaries, membranes, PSD, axon, vesicles.

We present the results of a multi-class segmentation of brain electron microscopy data using slightly modified U-Net with tiling data-layers onto the 256x256 fragments saving the initial resolution. The Dice index of mitochondrial segmentation from the network into 5 classes showed 0.82 instead of 0.954 for binary recognition.

The research has shown that increasing the number of classes does not necessarily have a negative impact on the quality of segmentation. The quality of segmentation is affected by the presence of a sufficient number of specific features that distinguish the selected classes and the representation of these features in the training dataset. With 6-class segmentation, mitochondria were segmented with a Dice index of 0.94, which is higher than with 5-class segmentation (0.892).

References

- [1] M. Helmstaedter, P. P. Mitra, Computational methods and challenges for large-scale circuit mapping, *Current Opinion in Neurobiology* 22 (2012) 162 – 169. URL: <http://www.sciencedirect.com/science/article/pii/S0959438811002133>. doi:<https://doi.org/10.1016/j.conb.2011.11.010>, neurotechnology.
- [2] L. G. Valiant, A quantitative theory of neural computation (2005).

- [3] D. C. Cireşan, L. M. Gambardella, A. Giusti, J. Schmidhuber, Deep neural networks segment neuronal membranes in electron microscopy images, in: IN NIPS, 2012, pp. 2852–2860.
- [4] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, 2015. [arXiv:1505.04597](https://arxiv.org/abs/1505.04597).
- [5] H. Chen, X. Qi, J.-Z. Cheng, P. Heng, Deep contextual networks for neuronal structure segmentation, in: AAAI, 2016.
- [6] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, 2015. [arXiv:1411.4038](https://arxiv.org/abs/1411.4038).
- [7] M. Drozdal, E. Vorontsov, G. Chartrand, S. Kadoury, C. Pal, The importance of skip connections in biomedical image segmentation, 2016. [arXiv:1608.04117](https://arxiv.org/abs/1608.04117).
- [8] A. E. Fakhry, T. Zeng, S. Ji, Residual deconvolutional networks for brain electron microscopy image segmentation, *IEEE Transactions on Medical Imaging* 36 (2017) 447–456.
- [9] C. Xiao, J. Liu, X. Chen, H. Han, C. Shu, Q. Xie, Deep contextual residual network for electron microscopy image segmentation in connectomics, in: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), 2018, pp. 378–381. doi:10.1109/ISBI.2018.8363597.
- [10] F. Milletari, N. Navab, S.-A. Ahmadi, V-net: Fully convolutional neural networks for volumetric medical image segmentation, 2016 Fourth International Conference on 3D Vision (3DV) (2016) 565–571.
- [11] Özgün Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, O. Ronneberger, 3d u-net: Learning dense volumetric segmentation from sparse annotation, 2016. [arXiv:1606.06650](https://arxiv.org/abs/1606.06650).
- [12] K. Kamnitsas, C. Ledig, V. Newcombe, J. P. Simpson, A. D. Kane, D. Menon, D. Rueckert, B. Glocker, Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation, *Medical Image Analysis* 36 (2017) 61–78.
- [13] W. Li, G. Wang, L. Fidon, S. Ourselin, M. J. Cardoso, T. Vercauteren, On the compactness, efficiency, and representation of 3d convolutional networks: Brain parcellation as a pretext task, in: M. Niethammer, M. Styner, S. Aylward, H. Zhu, I. Oguz, P.-T. Yap, D. Shen (Eds.), *Information Processing in Medical Imaging*, Springer International Publishing, Cham, 2017, pp. 348–360.
- [14] T. Liu, C. Jones, M. Seyedhosseini, T. Tasdizen, A modular hierarchical approach to 3d electron microscopy image segmentation, *Journal of Neuroscience Methods* 226 (2014) 88–102.
- [15] J. Liu, W. Li, C. Xiao, B. Hong, Q. Xie, H. Han, Automatic detection and segmentation of mitochondria from sem images using deep neural network, 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) (2018) 628–631.
- [16] I. Oztel, G. Yolcu, I. Ersoy, T. White, F. Bunyak, Mitochondria segmentation in electron microscopy volumes using deep convolutional neural network, in: 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2017, pp. 1195–1200. doi:10.1109/BIBM.2017.8217827.
- [17] M. Žerovnik Mekuč, C. Bohak, S. Hudoklin, B. H. Kim, R. Romih, M. Y. Kim, M. Marolt, Automatic segmentation of mitochondria and endolysosomes in volumetric electron microscopy data, *Computers in Biology and Medicine* 119 (2020) 103693. URL: <https://www.sciencedirect.com/science/article/pii/S0010482520300792>. doi:<https://doi.org/>

10.1016/j.combiomed.2020.103693.

- [18] S.-C. Huang, F. Cheng, Y. Chiu, Efficient contrast enhancement using adaptive gamma correction with weighting distribution, *IEEE Transactions on Image Processing* 22 (2013) 1032–1041.
- [19] F. Chollet, Xception: Deep learning with depthwise separable convolutions, 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 1800–1807.
- [20] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, 2016. doi:10.1109/CVPR.2016.308.
- [21] S. Xie, R. B. Girshick, P. Dollár, Z. Tu, K. He, Aggregated residual transformations for deep neural networks, 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 5987–5995.
- [22] Z. Yuan, X. Ma, J. Yi, Z. Luo, J. Peng, Hive-net: Centerline-aware hierarchical view-ensemble convolutional network for mitochondria segmentation in em images, *Computer methods and programs in biomedicine* 200 (2021) 105925.
- [23] A. Lucchi, K. Smith, R. Achanta, G. Knott, P. Fua, Supervoxel-based segmentation of mitochondria in em image stacks with learned shape features, *IEEE Transactions on Medical Imaging* 31 (2012) 474–486. doi:10.1109/TMI.2011.2171705.
- [24] T. Deerinck, E. Bushong, V. Lev-Ram, X. Shu, R. Tsien, M. Ellisman, Enhancing serial block-face scanning electron microscopy to enable high resolution 3-d nanohistology of cells and tissues, *Microscopy and Microanalysis* 16 (2010) 1138–1139. doi:10.1017/S1431927610055170.
- [25] I. Arganda-Carreras, S. C. Turaga, D. R. Berger, D. Cireşan, A. Giusti, L. M. Gambardella, J. Schmidhuber, D. Laptev, S. Dwivedi, J. M. Buhmann, T. Liu, M. Seyedhosseini, T. Tasdizen, L. Kamensky, R. Burget, V. Uher, X. Tan, C. Sun, T. D. Pham, E. Bas, M. G. Uzunbas, A. Cardona, J. Schindelin, H. S. Seung, Crowdsourcing the creation of image segmentation algorithms for connectomics, *Frontiers in Neuroanatomy* 9 (2015) 142. URL: <https://www.frontiersin.org/article/10.3389/fnana.2015.00142>. doi:10.3389/fnana.2015.00142.
- [26] A. J. Perez, M. Seyedhosseini, T. Deerinck, E. Bushong, S. Panda, T. Tasdizen, M. Ellisman, A workflow for the automatic segmentation of organelles in electron microscopy image stacks, *Frontiers in Neuroanatomy* 8 (2014).
- [27] H.-C. Cheng, A. Varshney, Volume segmentation using convolutional neural networks with limited training data, in: 2017 IEEE International Conference on Image Processing (ICIP), 2017, pp. 590–594. doi:10.1109/ICIP.2017.8296349.
- [28] A. Lucchi, Y. Li, P. Fua, Learning for structured prediction using approximate subgradient descent with working sets, in: 2013 IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 1987–1994. doi:10.1109/CVPR.2013.259.
- [29] T. M. Quan, D. G. C. Hildebrand, W.-K. Jeong, Fusionnet: A deep fully residual convolutional neural network for image segmentation in connectomics, *Frontiers in Computer Science* 3 (2021) 34. URL: <https://www.frontiersin.org/article/10.3389/fcomp.2021.613981>. doi:10.3389/fcomp.2021.613981.
- [30] W. Shen, B. Wang, Y. Jiang, Y. Wang, A. Yuille, Multi-stage multi-recursive-input fully convolutional networks for neuronal boundary detection, 2017 IEEE International Conference on Computer Vision (ICCV) (2017) 2410–2419.