



# Detecting Synapse Location and Connectivity by Signed Proximity Estimation and Pruning with Deep Nets

Toufiq Parag<sup>1</sup>(✉), Daniel Berger<sup>2</sup>, Lee Kamentsky<sup>1</sup>, Benedikt Staffler<sup>3</sup>,  
Donglai Wei<sup>1</sup>, Moritz Helmstaedter<sup>3</sup>, Jeff W. Lichtman<sup>2</sup>,  
and Hanspeter Pfister<sup>1</sup>

<sup>1</sup> SEAS, Harvard University, Cambridge, MA, USA  
[toufiq.parag@gmail.com](mailto:toufiq.parag@gmail.com)

<sup>2</sup> MCB, Harvard University, Cambridge, MA, USA

<sup>3</sup> Max Planck Institute for Brain Research, Frankfurt, Germany

**Abstract.** Synaptic connectivity detection is a critical task for neural reconstruction from Electron Microscopy (EM) data. Most of the existing algorithms for synapse detection do not identify the cleft location and direction of connectivity simultaneously. The few methods that compute direction along with contact location have only been demonstrated to work on either dyadic (most common in vertebrate brain) or polyadic (found in fruit fly brain) synapses, but not on both types. In this paper, we present an algorithm to automatically predict the location as well as the direction of both dyadic and polyadic synapses. The proposed algorithm first generates candidate synaptic connections from voxelwise predictions of signed proximity generated by a 3D U-net. A second 3D CNN then prunes the set of candidates to produce the final detection of cleft and connectivity orientation. Experimental results demonstrate that the proposed method outperforms the existing methods for determining synapses in both rodent and fruit fly brain. (Code at: <https://github.com/paragt/EMSynConn>).

## 1 Introduction

Connectomics has become a fervent field of study in neuroscience and computer vision recently. The goal of EM connectomics is to reconstruct the neural wiring diagram from Electron Microscopic (EM) images of animal brain. Neural reconstruction of EM images consists of two equally important tasks: (1) trace the anatomical structure of each neuron by labeling each voxel within a cell with a unique id; and (2) find the location and direction of synaptic connections among multiple cells.

The enormous amount of EM volume emerging from a tiny amount of tissue constrains any subsequent analysis to be performed (semi-) automatically to acquire a comprehensive knowledge within a practical time period [1, 2]. Discovering the anatomical structure entails a 3D segmentation of EM volume. Numerous studies have addressed this task with many different approaches, we

refer interested readers to [3–9] for further details. In order to unveil the connectivity, it is necessary to identify the locations *and* the direction of synaptic communications between two or more cells. Resolving the location of synaptic contact is crucial for neurobiological reasons, and, because the strength of connection between two cells is determined by the number of times they make a synaptic contact. The direction of the synaptic contact reveals the direction of information flow from presynaptic to postsynaptic cells. By defining the edges, the location and connectivity orientation of synapses complete the directed network of neural circuitry that a neural reconstruction seek to discover. In fact, discovering synaptic connectivity was one of the primary reasons to employ the immensely complex and expensive apparatus of electron microscopy for connectomics in the first place. Other imaging modalities (e.g., light microscopy) are either limited by their resolution or by a conclusive and exhaustive strategy (e.g., using reagents) to locate synapses [10–12].

In terms of complexity, identification of neural connectivity is as challenging as tracing the neurons [13]. With rapid and outstanding improvement in automated EM segmentation in recent years, detection of synaptic connectivity may soon become a bottleneck in the overall neural reconstruction process [14]. Although fewer in number when compared against those in neurite segmentation, there are past studies on synaptic connectivity detection; we mention some notable works in the relevant literature section below. Despite many discernible merits of previous works, very few of them aim to identify both the location and direction of synaptic junctions. Among these few methods, namely by [13–15], none of them have been shown to be generally applicable on different types of synapses typically found on different species of animals, e.g., dyadic connections in vertebrate (mouse, rat, zebra finch, etc.) and polyadic connections in non-vertebrate (fruit fly) brain<sup>1</sup>. Apart from a few, the past approaches do not benefit from the advantages deep (fully) convolutional networks offer. Use of hand crafted features could stifle the utility of a method on widely divergent EM volumes collected from different animals with different tissue processing and imaging techniques.

In this paper, we propose a general method to automatically detect both the 3D location and direction of both dyadic (vertebrate) and polyadic (fruit fly) synaptic contacts. The proposed algorithm is designed to predict the proximity (and polarity, as we will explain later) of every 3D voxel to a synaptic cleft using a deep fully convolutional network, namely a U-net [16]. A set of putative locations, along with their connection direction estimates, are computed given a segmentation of the volume and the voxelwise prediction from the U-net. A second stage of pruning, performed by a deep convolutional network, then trims the set of candidates to produce the final prediction of 3D synaptic cleft locations and the corresponding directions.

The use of CNNs makes the proposed approach generally applicable to new data without the need for feature selection. Estimation of the location and

---

<sup>1</sup> Although, there are examples of polyadic connections in mouse cerebellum between mossy fibers and granule cells.

connectivity at both voxel and segment level improves the accuracy but do not require any additional annotation effort (no need for labels for other classes such as vesicles). We show that our proposed algorithm outperforms the existing approaches for synaptic connectivity detection on both vertebrate and fruit fly datasets. Our evaluation measure (Sect. 3), which is critical to assess the actual performance of a synapse detection method, has also been confirmed by a prominent neurobiologist to correctly quantify actual mistakes on real dataset.

## 1.1 Relevant Literature

Initial studies on automatic synaptic detection focused on identifying the cleft location by classical machine learning/vision approaches using pre-defined features [17–22]. These algorithms assumed subsequent human intervention to determine the synaptic partners given the cleft predictions. Roncal et al. [23] combine the information provided by membrane and vesicle predictions with a CNN (not fully convolutional) and apply post-processing to locate synaptic clefts. To establish the pre- and post-synaptic partnership, [24] augmented the synaptic cleft detection with a multi-layer perceptron operating on hand designed features. On the other hand, Kreshuk et al. [15] seek to predict vesicles and clefts for each neuron boundary by a random forest classifier (RF) and then aggregate these predictions with a CRF to determine the connectivity for polyadic synapses in fruit fly. Dorkenwald et al. [13] utilize multiple CNNs, RFs to locate synaptic junctions as well as vesicles, mitochondria, and to decide the dyadic orientation in vertebrate EM data. SynEM [14] attempts to predict connectivity by classifying each neuron boundary (interfaces) to be synaptic or not using a RF classifier operating on a confined neighborhood and has been shown to perform better than [13] in terms of connectivity detection.

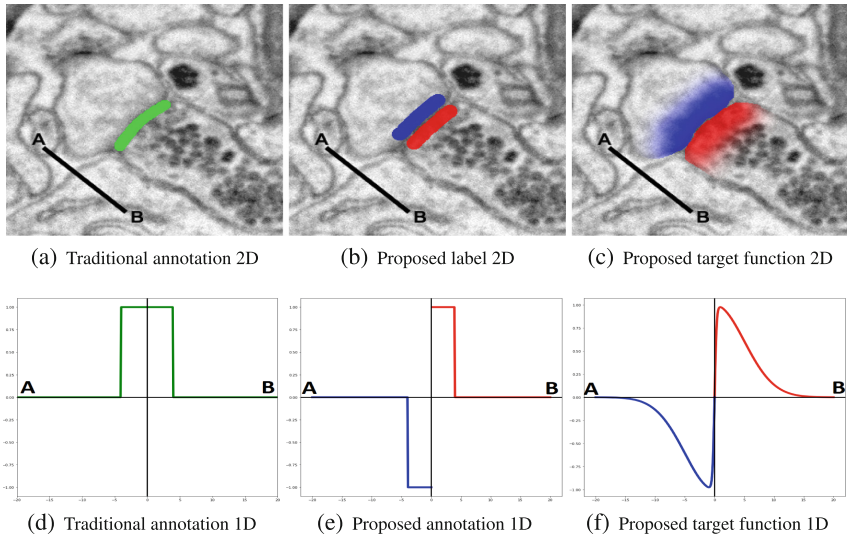
## 2 Method

The proposed method is designed to first predict both the location and direction of synaptic communication at the voxel level. Section 2.1 illustrates how this is performed by training a deep encoder-decoder type network, namely the U-net, to compute the proximity and direction of connection with respect to synapses at every voxel. The voxelwise predictions are clustered together after discretization and matched with a segmentation to establish synaptic relations between pairs of segment ids. Afterwards, a separate CNN is trained to discard the candidates that do not correspond to an actual synaptic contact, both in terms of location and direction, as described in Sect. 2.2.

### 2.1 Voxelwise Synaptic Cleft and Partner Detection

In order to learn both the position and connection orientation of a synapse, the training labels for voxels are modified slightly from the traditional annotation. It is the standard practice to demarcate only the synaptic cleft with a single strip

of id, or color, as the overlaid color shows in Fig. 1(a). In contrast, the proposed method requires the neighborhood of the pre- and post-synaptic neurons at the junction of synaptic expression to be marked by two ids or colors as depicted in Fig. 1(b). To distinguish the partners unambiguously, these ids can follow a particular pattern, e.g., pre-synaptic partners are always marked with odd id and post-synaptic partners are annotated with even ids. Such annotations inform us about both the location and direction of a connection with practically no increase in annotation effort. Note that, although we explain and visualize the labels in 2D and in 1D for better understanding, our proposed method learns a function of the 3D annotations.



**Fig. 1.** The traditional and proposed annotation of synapses and the signed proximity function that the proposed method estimates are shown in 2D in (a), (b), (c) and in 1D in (d), (e), (f) respectively for illustration purposes. The U-net learns the signed proximity functions in 3D. (Color figure online)

For voxelwise prediction of position and direction, ids of all pre- and post-synaptic labels (red and blue in Fig. 1(b)) are converted to 1 and  $-1$  respectively; all remaining voxels are labeled with 0. However, our approach *does not* directly learn from the discrete labels presented in Fig. 1(b). Instead, we attempt to learn a smoother version of the discrete labels, where the transition from 1 and  $-1$  to 0 occurs gradually, as shown in Fig. 1(c). The dissimilarities among these three types of annotations can be better understood in 1D. The Figs. 1(d), (e), and (f) plot the labels perpendicular to the black line drawn underneath the labels in images in the top row of Fig. 1. The proposed approach attempts to learn a 3D version of the smooth function in Fig. 1(c) and (f). Effectively, this target function computes a signed proximity from the synaptic contact, where the sign

indicates the connectivity orientation and the absolute function value determines the closeness or affinity of any voxel to the nearest cleft. Mathematically, this function can be formulated as

$$\text{proximity}_{\text{signed}}(x) = \left\{ \exp \frac{-d(x)^2}{2\sigma^2} \right\} \left\{ \frac{2}{1 + \exp(-\alpha d(x))} - 1 \right\}, \quad (1)$$

where  $d(x)$  is the signed distance between voxel at  $x$  and the synaptic cleft, and  $\alpha$  and  $\sigma$  are parameters that control the smoothness of transition. We solve a regression problem using a 3D implementation of U-net with linear final layer activation function to approximate this function. There are multiple motivations to approximate a smooth signed proximity function. A smooth proximity function as a target also eliminates the necessity of estimating the abrupt change near the annotation boundary and therefore assists the gradient descent to approach a more meaningful local minimum. Furthermore, some recent studies have suggested that a smooth activation function is more useful than its discrete counterparts for learning regression [25,26].

Our 3D U-net for learning signed proximity has an architecture similar to the original U-net model in [16]. The network has a depth of 3 where it applies two consecutive  $3 \times 3 \times 3$  convolutions at each depth and utilizes parametric leaky ReLU [27] activation function. The activation function in the final layer is linear. The input and output of the 3D U-net are  $316 \times 316 \times 32$  grayscale EM volumes and  $228 \times 228 \times 4$  proximity values, respectively. A weighted mean squared error loss is utilized to learn the proximity values during training.

## 2.2 Candidate Generation and Pruning


For computing putative pairs of pre- and post-synaptic partners, we first threshold the signed proximity values at an absolute value of  $\tau$  and compute connected components for pre and post-synaptic sites separately. Given a segmentation  $S$  for the EM volume, every pre-synaptic connected component  $e$  is paired with one or more segment  $s_{i_e} \in S$ ,  $i_e = 1, \dots, m_e$  based on a minimum overlap size  $\omega$  to form pre-synaptic site candidates  $T_{e,i_e}$ . Similarly, post-synaptic site candidates  $T_{o,j_o}$  are generated by associating each post-synaptic connected component  $o$  is with a set of segments  $s_{j_o} \in S$ ,  $j_o = 1, \dots, n_o$ . The set  $\mathcal{C}$  of candidate pairs of synaptic partners are computed by pairing up pre-synaptic candidate  $T_{e,i_e}$  with post-synaptic  $T_{o,j_o}$  wherever segment  $s_{j_o}$  is a neighbor of  $s_{i_e}$ , i.e.,  $s_{i_e}$  shares a boundary with segment  $s_{j_o}$ .

$$\mathcal{C} = \left\{ \{T_{e,i_e}, T_{o,j_o}\} \mid s_{j_o} \in \text{Nbr}(s_{i_e}), \forall e, o, i_e, j_o \right\}. \quad (2)$$

A 3D CNN is utilized to distinguish the correct synaptic partner pairs from the false positive candidates, i.e., to produce a binary decision for each candidate. The groundtruth labels for training this second convolutional network were computed by matching the segmentation  $S$  with the groundtruth segmentation of the volume  $G$ . The pruning network is constructed with 5 layer convolutions

of size  $3 \times 3 \times 3$  followed by two densely connected layers. The input to the 3D deep convolutional pruning (or trimming) network comprises  $160 \times 160 \times 16$  subvolumes of the grayscale EM image, the predicted signed proximity values from 3D U-net and the segmentation masks of  $s_{i_e}, s_{j_o}$ , extracted from a 3D window centered at the closest point between the connected components  $e$  and  $o$ , as shown in Fig. 2.

It is worth mentioning here that we have contemplated the possibility of combining the voxelwise network and the candidate trimming network to facilitate end-to-end training, but did not pursue that direction due to the difficulty in formulating a differentiable operation to transform voxelwise signed proximities to region wise candidates. We have, however, attempted to employ a region proposal network based method, in particular the mask R-CNN [28], to our problem. The proposal generating network of mask R-CNN method resulted in a low recall rate for locating the synapses in our experiments (leading to low recall in the final result after trimming). We observed Mask R-CNNs to struggle with targets with widely varying size in our dataset. Furthermore, we had difficulty in merging many proposals [28] produced for one connection, leading to lower final precision rate as well.



**Fig. 2.** Candidate pruning by 3D CNN. The EM image, proximity prediction and segmentation mask on one section of input subvolume are shown in (a), (b), (c) respectively. The cyan and yellow segmentation masks are provided as separate binary masks, shown here in one image (c) to save space. (Color figure online)

### 3 Experiments and Results

The deep nets for this work were implemented in Theano with Keras interface<sup>2</sup>. For both the network training, we used rotation and flip in all dimensions for data augmentation. For the candidate trimming network, we also displaced the center of the window by a small amount to augment the training set. The parameters for the target signed proximity and the candidate overlap calculation remained the same as  $\alpha = 5, \tau = 0.3$  (absolute value),  $\omega = 100$ . The parameter  $\sigma$  was set

<sup>2</sup> Code at: <https://github.com/paragt/EMSynConn>.

to 10 for the mouse and rat dataset but 14 for the TEM fly data to account for the difference in z-resolution.

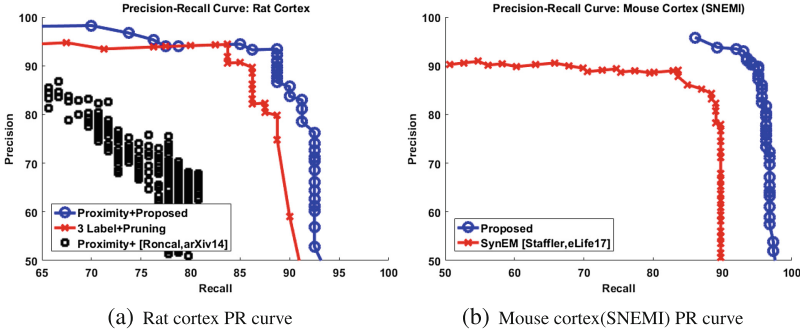
**Evaluation:** It is critical to apply the most biologically meaningful evaluation formula in order to correctly assess the performance of any given method. Distance based methods [14, 18], for example, are unrealistically tolerant to false positive detections nearby. On the other hand, pixelwise error computation [29] is more stringent than necessary for extracting the wiring diagram – two detections with 50% and 60% pixelwise overlap need not be penalized differently for connectomics purposes. Measures computed solely on overlap [23] becomes ambiguous when one prediction overlaps two junctions. We resolve this ambiguity by considering a detection be correct only if it overlaps with the span of synaptic expression (as delineated by an expert) *and* connects two cells with correct synaptic orientation, i.e., pre and post-synaptic partners. All the precision and recall values reported in the experiments on rat (Sect. 3.1) and mouse (Sect. 3.2) data are computed with this measure.

### 3.1 Rat Cortex

Our first experiment was designed to determine the utility of the two stages, i.e., voxelwise prediction and candidate set pruning, of the proposed algorithm. The EM images we used in this experiment were collected from rat cortex using SEM imaging at resolution of  $4 \times 4 \times 30$  nm. We used a volume of 97 images to train the 3D U-net and validated on a different set of 120 images. The candidate pruning CNN was trained on 97 images and then fine tuned on the 120 image dataset. For testing we used a third volume of 145 sections. The segmentation used to compute the synaptic direction was generated either by the method of [30].

Figure 3(a) compares the precision recall curves for detecting both location and connectivity with two variants. (1) 3 Label + pruning - where the proximity approximation is replaced by 3-class classification among pre-, post-synaptic, and rest of the voxels (Fig. 1(b)). (2) Proximity + [Roncal, arXiv14] - where the proposed pruning network is replaced by VESICLE [23] style post-processing. For the proposed (blue o) and 3 Label + pruning (red x) technique, each point on the plot correspond to a threshold on the prediction of the 3D trimming CNN. For the Proximity + [Roncal, arXiv14] technique (black o), we varied several parameters of the VESICLE post-processing.

This experiment suggests that the pruning network is substantially more effective than morphological post-processing [23]. The proposed signed proximity approximation yields 3% more true positives in the initial candidate set than those generated by the multiclass prediction. As discussed earlier in Sect. 2.1, a smooth target function places the focus on learning difficult examples by removing the necessity to learn the sharp boundaries. Empirically, we have noticed the training procedure to spend a significant number of iterations to infer the sharp boundary of a discrete label like Figure 1(b) and (c). Furthermore, we observed that the wider basin of prediction can identify more true positives and offers



**Fig. 3.** Precision recall curves for *synapse location and connectivity* for rat cortex (a) and mouse cortex (SNEMI) (b) experiments. Plot (a) suggests combining voxelwise signed proximity prediction with pruning performs significantly better than the versions that replaces one of these components with alternative strategies. Plot (b) indicates significant improvement achieved over performance of [14] on the same dataset. (Color figure online)

more information for the following pruning stage to improve the F-score of our method to 91.03% as opposed 87.5% of the 3 Label + pruning method.


### 3.2 Mouse Cortex Data (SNEMI)

We experimented next on the SEM dataset from Kasthuri et al. [31] that was used for the SNEMI challenge [32] to compare our performance with that of SynEM [14] (which was shown to outperform [13]). The synaptic partnership information was collected from the authors of [31] to compute the signed proximities for training the 3D U-net. Similar to the rat cortex experiment, we used 100 sections for training the voxelwise U-net and candidate pruning CNN and used 150 sections for validation of the voxelwise proximity U-net. The test dataset consists of 150 sections of size  $1024 \times 1024$  that is referred to as AC3 in [14]. The segmentation used to compute the synaptic direction was generated either by the method of [30].

The authors of [14] have kindly provided us with their results on this dataset. We read off the cleft detection and the connectivity estimation from their result and computed the error measures as explained in Sect. 3. The precision recall curve for detecting both cleft and connectivity is plotted in Fig. 3(b). In general, the SynEM method performs well overall, but produces significantly higher false negative rates than the proposed method. Among the synapses it detects, SynEM assigns the pre and post-synaptic partnership very accurately, although it used the actual segmentation groundtruth for such assignment whereas we used a segmentation output. Visual inspection of the detection performances also verifies the lower recall rate of [14] than ours. In Fig. 4 we show the groundtruth, the cleft estimation of SynEM (missed connection marked with red x) and the connected components corresponding to the predictions of the proposed method,



both computed at the largest F-score. The direction of synaptic connection in image of Fig. 4(c) is color coded – purple and green indicates pre and post-synaptic partners respectively.



**Fig. 4.** Qualitative results on mouse cortex data [31]. Left to right, groundtruth annotation, output of SynEM [14] (red x marks missed location), and that of the proposed method. In (c), purple and green indicates pre and post-synaptic partner respectively. (Color figure online)

### 3.3 Fruit Fly Data (CREMI)

Our method was applied on the 3 TEM volumes (resolution  $4 \times 4 \times 40$  nm) of the CREMI challenge [29]. We annotated the training labels for synaptic partners for all 3 volumes given the synaptic partner list provided on the website. All the experimental settings for this experiment remain the same as other except those mentioned in Sect. 3. Out of the 125 training images, we used the first 80 for training and the remaining images for validation. The segmentation used to compute the synaptic direction was generated either by the method of [3].

The performance is only measured in terms of synaptic partner identification task, as the pixelwise cleft error measure is not appropriate for our result (refer to the output provided in Fig. 4(c)). At the time of this submission the our method, which is identified as HCBS on CREMI leaderboard, holds the 2nd place overall (error difference with the first is 0.002) and performed better than both variants of [15] (Table 1).

**Table 1.** Result on CREMI data, lower is better

Method	Submission	CREMI score	FP	FN
HCBS (proposed)	Tr66_80K	0.449	223.000	286.667
IAL [15]	PSP_unar	0.461	266.667	281.000
IAL [15]	PSP_full	0.464	187.333	310.000

## 4 Conclusion

We propose a general purpose synaptic connectivity detector that locates the location and direction of a synapse at the same time. Our method was designed to work on both dyadic and polyadic synapses without any modification to its component techniques. The utilization of deep CNNs for learning location and direction of synaptic communication enables it to be directly applicable to any new dataset without the need for manual selection of features. Experiments on multiple datasets suggests the superiority of our method on existing algorithms for synaptic connectivity detection. One straightforward extension of the proposed two stage method is to enhance the candidate pruning CNN to distinguish between excitatory and inhibitory synaptic connections by adopting a 3-class classification scheme.

## References

1. Jain, V., Seung, S., Turaga, S.: Machine that learn to segment images: a crucial technology for connectomics. *Curr. Opinion Neurobiol.* **20**, 653–666 (2010)
2. Helmstaedter, M.: The mutual inspirations of machine learning and neuroscience. *Neuron* **86**(1), 25–28 (2015)
3. Funke, J., Tschopp, F.D., Grisaitis, W., Singh, C., Saalfeld, S., Turaga, S.C.: A deep structured learning approach towards automating connectome reconstruction from 3D electron micrographs. [arXiv:1709.02974](https://arxiv.org/abs/1709.02974) (2017)
4. Lee, K., Zung, J., Li, P., Jain, V., Seung, H.S.: Superhuman accuracy on the SNEMI3D connectomics challenge. [arXiv:1706.00120](https://arxiv.org/abs/1706.00120) (2017)
5. Januszewski, M., Maitin-Shepard, J., Li, P., Kornfeld, J., Denk, W., Jain, V.: Flood-filling networks. [arXiv:1611.00421](https://arxiv.org/abs/1611.00421) (2016)
6. Parag, T., Ciresan, D.C., Giusti, A.: Efficient classifier training to minimize false merges in electron microscopy segmentation. In: *The IEEE International Conference on Computer Vision (ICCV)*, December 2015
7. Liu, T., Zhang, M., Javanmardi, M., Ramesh, N., Tasdizen, T.: SSHMT: semi-supervised hierarchical merge tree for electron microscopy image segmentation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9905, pp. 144–159. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46448-0\\_9](https://doi.org/10.1007/978-3-319-46448-0_9)
8. Parag, T., Plaza, S., Scheffer, L.: Small sample learning of superpixel classifiers for EM segmentation. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) *MICCAI 2014*. LNCS, vol. 8673, pp. 389–397. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-10404-1\\_49](https://doi.org/10.1007/978-3-319-10404-1_49)
9. Beier, T., Pape, C., Rahaman, N., Prange, T., Berg, S.E., et al.: Multicut brings automated neurite segmentation closer to human performance. *Nat. Methods* **14**, 101–102 (2017)
10. Morgan, J.L., Lichtman, J.W.: Why not connectomics? *Nat. Methods* **10**(6), 494–500 (2013)
11. Denk, W., Briggman, K.L., Helmstaedter, M.: Structural neurobiology: missing link to a mechanistic understanding of neural computation. *Nat. Rev. Neurosci.* **13**(5), 351–358 (2011)
12. Lichtman, J.W., Pfister, H., Shavit, N.: The big data challenges of connectomics. *Nat. Neurosci.* **17**(11), 1448–1454 (2014)

13. Dorkenwald, S., et al.: Automated synaptic connectivity inference for volume electron microscopy. *Nat. Methods* **14**(4), 435–442 (2017)
14. Staffler, B., Berning, M., Boergens, K.M., Gour, A., Smagt, P.V.D., Helmstaedter, M.: SynEM, automated synapse detection for connectomics. *eLife* **6**, e26414 (2017)
15. Kreshuk, A., Funke, J., Cardona, A., Hamprecht, F.A.: Who Is talking to whom: synaptic partner detection in anisotropic volumes of insect brain. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9349, pp. 661–668. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24553-9\\_81](https://doi.org/10.1007/978-3-319-24553-9_81)
16. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
17. Kreshuk, A., et al.: Automated detection and segmentation of synaptic contacts in nearly isotropic serial electron microscopy images. *PLoS ONE* **6**(10), e24899 (2011)
18. Becker, C., Ali, K., Knott, G., Fua, P.: Learning context cues for synapse segmentation. *IEEE Trans. Med. Imaging* **32**(10), 1864–1877 (2013)
19. Kreshuk, A., Koethe, U., Pax, E., Bock, D.D., Hamprecht, F.A.: Automated detection of synapses in serial section transmission electron microscopy image stacks. *PLoS ONE* **9**(2), e87351 (2014)
20. Plaza, S.M., Parag, T., Huang, G.B., Olbris, D.J., Saunders, M.A., Rivlin, P.K.: Annotating synapses in large EM datasets. [arXiv:1409.1801](https://arxiv.org/abs/1409.1801) (2014)
21. Jagadeesh, V., Anderson, J., Jones, B., Marc, R., Fisher, S., Manjunath, B.: Synapse classification and localization in electron micrographs. *Pattern Recognit. Lett.* **43**, 17–24 (2014)
22. Huang, G.B., Plaza, S.: Identifying synapses using deep and wide multiscale recursive networks. [arXiv: 1409.1789](https://arxiv.org/abs/1409.1789) (2014)
23. Roncal, W.G., et al.: VESICLE: volumetric evaluation of synaptic interfaces using computer vision at large scale. [arXiv:1403.3724](https://arxiv.org/abs/1403.3724) (2014)
24. Huang, G.B., Scheffer, L.K., Plaza, S.M.: Fully-automatic synapse prediction and validation on a large data set. [arXiv:1604.03075](https://arxiv.org/abs/1604.03075) (2016)
25. Hou, L., Samaras, D., Kurc, T., Gao, Y., Saltz, J.: ConvNets with smooth adaptive activation functions for regression. In: *Proceedings of the 20th AISTATS* (2017)
26. Mobahi, H.: Training recurrent neural networks by diffusion. [arXiv:1601.04114](https://arxiv.org/abs/1601.04114) (2016)
27. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: *ICCV* (2015)
28. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: *ICCV* (2017)
29. Funke, J., Saalfeld, S., Bock, D., Turaga, S., Perlman, E.: CreMI challenge. <https://cremi.org>
30. Parag, T., et al.: Anisotropic EM segmentation by 3D affinity learning and agglomeration. [arXiv:1707.08935](https://arxiv.org/abs/1707.08935) (2017)
31. Kasthuri, N., Hayworth, K., Berger, D., Schalek, R., et al.: Saturated reconstruction of a volume of neocortex. *Cell* **162**(3), 648–661 (2015)
32. Arganda-Carreras, I., Seung, H.S., Vishwanathan, A., Berger, D.R.: Snemi challenge. <http://brainiac2.mit.edu/SNEMI3D/>