



Synaptic Partner Prediction from Point Annotations in Insect Brains

Julia Buhmann^{1(✉)}, Renate Krause^{1,2(✉)}, Rodrigo Ceballos Lentini¹, Nils Eckstein¹, Matthew Cook¹, Srinivas Turaga², and Jan Funke^{2,3}

¹ Institute of Neuroinformatics UZH/ETHZ, Zurich, Switzerland
{juliab, rekrau}@ini.uzh.ch

² HHMI Janelia Research Campus, Ashburn, USA

³ Institut de Robotica i Informatica Industrial, UPC, Barcelona, Spain

Abstract. High-throughput electron microscopy allows recording of large stacks of neural tissue with sufficient resolution to extract the wiring diagram of the underlying neural network. Current efforts to automate this process focus mainly on the segmentation of neurons. However, in order to recover a wiring diagram, synaptic partners need to be identified as well. This is especially challenging in insect brains like *Drosophila melanogaster*, where one presynaptic site is associated with multiple postsynaptic elements. Here we propose a 3D U-Net architecture to directly identify pairs of voxels that are pre- and postsynaptic to each other. To that end, we formulate the problem of synaptic partner identification as a classification problem on long-range edges between voxels to encode both the presence of a synaptic pair and its direction. This formulation allows us to directly learn from synaptic point annotations instead of more expensive voxel-based synaptic cleft or vesicle annotations. We evaluate our method on the MICCAI 2016 CREMI challenge and improve over the current state of the art, producing 3% fewer errors than the next best method (Code at: <https://github.com/juliabuhmann/syntist>).

1 Introduction

The field of Connectomics studies the reconstruction and connectivity of neurons. So far, only electron microscopy (EM) methods are able to image neural tissue with a resolution high enough to resolve connectivity at a synaptic level. This results in large amounts of data, requiring automated methods in order to be processed in a reasonable amount of time. In order to acquire a full connectivity diagram we need to solve two main problems: First, the neurons have to be segmented to reconstruct their morphology. Second, synaptic connections between neurons have to be identified, since neuron morphology does not provide enough information to estimate connectivity [7].

The segmentation of neurons has received considerable attention by the computer vision community, for recent methods see [3, 6]. Automated synapse detection methods have also seen progress, for instance in [2], where automatically

J. Buhmann and R. Krause—Equal contribution.

© Springer Nature Switzerland AG 2018

A. F. Frangi et al. (Eds.): MICCAI 2018, LNCS 11071, pp. 309–316, 2018.

https://doi.org/10.1007/978-3-030-00934-2_35

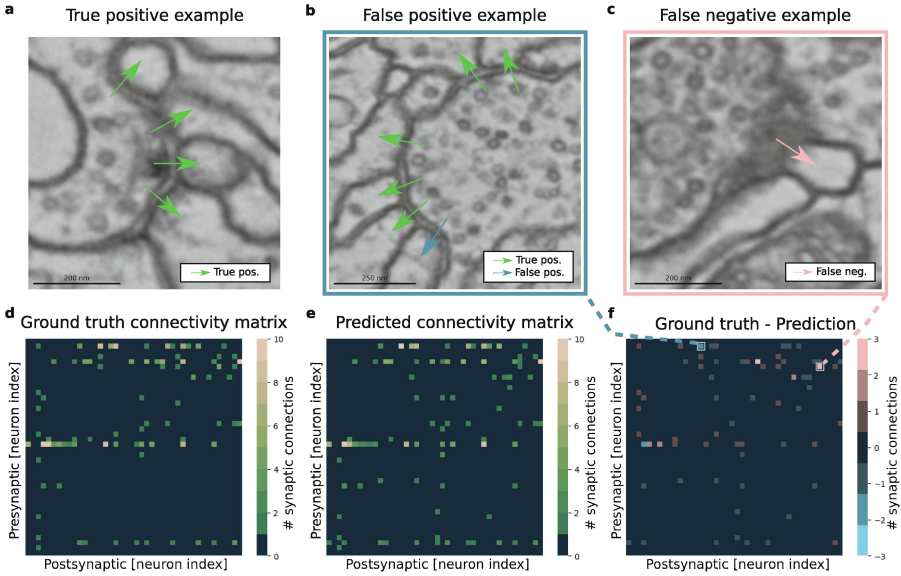


Fig. 1. Qualitative results on the validation set. First row: examples of (a) True Positive (TP), (b) False Positive (FP) and (c) False Negative (FN) in 2D sections. Note that in (b) the predicted partner is likely correct and was overlooked during ground truth annotation while (c) represents an ambiguous case. Second row: Connectivity matrices of (d) ground truth and (e) predicted connectome. Each row and each column in a connectivity matrix corresponds to one neuron. A positive entry in the matrix stands for the number of connections/synapses that exist between the two neurons. In (f) the ground truth matrix is subtracted from the prediction matrix (blue: FP, red: FN). Results shown for the validation set (f-score 0.75) and ground truth segmentation based synaptic partner extraction.

identified synapses allow the analysis of different neuron cell types. Most of the proposed methods identify synaptic *sites* [1, 5, 10–12]. For the overall goal of reconstructing the connectome, however, it is necessary to additionally identify the pre- and postsynaptic *partners*. In mouse, zebrafish, or zebra finch, a synaptic site usually has only one pre- and one postsynaptic partner, whereas in insects such as *Drosophila melanogaster*, synapses form one-to-many connections which are called polyadic synapses (for an example, see Fig. 1a and b).

To identify partners in one-to-one synapses, a common strategy is to first identify contact sites of neuronal segments to extract a set of potential synaptic partners, and second to collect features for them which are then handed over to a classifier to decide if a given candidate is valid [2, 14].

In [4], a similar strategy is used to detect synaptic partners in polyadic synapses found in insects. Specific changes are required to make the method applicable. For instance, while a vesicle cloud is quite indicative for a presynaptic site in songbird [2], the *Drosophila melanogaster* presynaptic site includes

a characteristic T-shaped appearance (called T-bar). In a first step, a convolutional neural network is trained to detect T-bars. Subsequent partner prediction is constrained to those identified locations. This method was shown to perform well on isotropic data.

However, for anisotropic datasets, the task of automatically detecting presynaptic locations is potentially harder.

The only currently proposed method to detect synaptic partners in anisotropic datasets with polyadic synapses by [9] requires the neuron segmentation as well as predicted synapse locations. Based on these, the method identifies a set of potential synaptic partners. Then unary and pairwise factors are assigned to all candidate pairs and an Integer Linear Program (ILP) is solved to select correct synaptic partners.

2 Method

A common approach for synaptic partner detection is to first identify potential pre- and postsynaptic point locations while the correct pairing of candidate sites is performed in a second, separate step [2, 9, 14]. Here we propose a method for synaptic partner prediction that fuses these two, formerly independent steps, into one. We achieve this reduction by formulating the problem as a classification task on the space of directed edges between voxels. An advantage of this approach is that potential pre- or postsynaptic sites are only classified as such if there is evidence for a corresponding partner. The proposed representation also allows us to learn from synaptic point annotations only, since we do not rely on labeled synaptic features, such as synaptic clefts or vesicle clouds.

In the following subsections, we describe how we (1) represent neuron connectivity by directed long-range edges, (2) train a U-Net to classify those edges, and (3) post-process the U-Net output to directly obtain synaptic partners.

2.1 Directed edges for synaptic partner representation

In order to constrain the learning problem we fuse synaptic site detection and partner identification in a single step. To this end we consider a directed graph $G(V, E)$, where each voxel is represented as a vertex $v \in V$ and we encode relations between voxels as directed edges $e_{i \rightarrow j} = (v_i, v_j) \in E$.

If we were to connect each voxel v_i in a given volume with each other voxel $v_{j \neq i}$, we could identify synaptic partners by selecting the subset of those edges $e_{i \rightarrow j} = (v_i, v_j)$ whose vertices v_i and v_j lie inside a pre- and corresponding postsynaptic region respectively, pointing from one site to the other. However, in such a scheme the number of candidate edges we would need to consider scales as $|V|^2$, which is infeasible and not necessary.

We thus limit the space of candidate edges to lie within a biologically plausible range using the prior that pre- and postsynaptic sites are separated from each other by a certain distance. That is, for each voxel v_i , we only consider n_e edges with associated relative offsets $\mathbf{r}_j \in \mathbb{Z}^3$, $j = 1, \dots, n_e$. The number

of edges n_e trades off the coverage rate of synaptic partners. For example, if we use too few edges, there might be synaptic partners that are missed by our representation. In order to select a minimal set of sufficient edges, we perform a grid search on a training and validation set over the number of candidate edges n_e , the length of the vectors $|\mathbf{r}_j|$, and radius of pre- and postsynaptic region r_{syn} such that we obtain a coverage rate of 100% of synaptic partners. We denote the set of edges that we find with the proposed method E^{syn} .

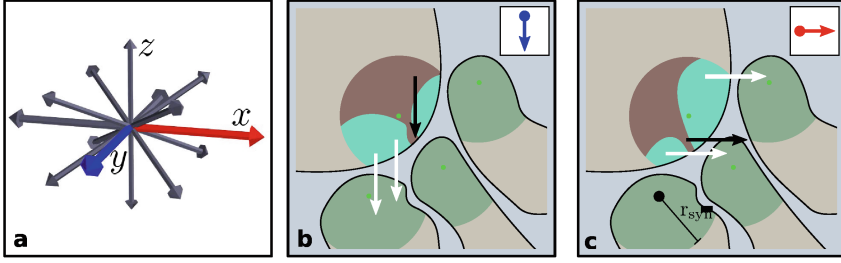


Fig. 2. Illustration of synaptic partner representation. **(a)** 3D rendering of the 14 offset vectors \mathbf{r}_j found with grid search on the CREMI dataset. **(b, c)** One presynaptic neuron with its 3 postsynaptic partners. An example is shown for two offset vectors **(b)** $\mathbf{r}_1 = (0, 120, 0)$, **(c)** $\mathbf{r}_2 = (120, 0, 0)$ in the x-y plane, their direction shown in the upper right corner and marked in **(a)** in blue and red. Green dots represent synaptic point annotations, with their pre- (brown) and post- (green) synaptic location area. White arrow: selected edge $s(e_{i \rightarrow j} = 1)$, black arrow: unselected edge $s(e_{i \rightarrow j} = 0)$. Teal colored regions correspond to the resulting activation map: each voxel v_i is colored teal if the corresponding edge $e_{i \rightarrow j}$ is selected, i.e., voxel v_j lies inside the postsynaptic area.

2.2 Edge classification

We train a 3D U-Net to classify for each voxel v_i which of its directed edges $e_{i \rightarrow j} \in E^{\text{syn}}$ is synaptic or non-synaptic. That is, we predict an n_e dimensional vector \mathbf{y}^i for each voxel in the input volume. Each entry \mathbf{y}_j^i then represents the score $s(e_{i \rightarrow j} = 1)$ that edge $e_{i \rightarrow j}$ is synaptic. Consistent with the search for E^{syn} , we annotate synaptic regions for training by expanding the ground truth point annotation of pre- and postsynaptic sites into a ball of radius r_{syn} , shown in brown/green in Fig. 2. Note that we constrain each synaptic region to lie within its corresponding neuron. An illustration of pre- and postsynaptic regions, selected and unselected edges for two offset vectors \mathbf{r}_1 and \mathbf{r}_2 in the x-y plane is shown in Fig. 2.

2.3 Synaptic partner extraction

We formalize the notion of a connectome by introducing a directed connectome graph $G_c = (S, A)$, with vertices $s \in S \subset P(V)$ representing neuron segments as

subsets of voxels $v \in V$ and candidate synapses between neurons $a \in A \subset P(E)$ as subsets of edges $e \in E$, where $P(\cdot)$ is the power set. We extract candidate synapses $a_{k \rightarrow l}$ from the edge predictions in the following way: we use an underlying neuron segmentation (this can be either ground truth or automatically generated) to collect all edges e that start in segment s_k and end in another segment $s_{l \neq k}$. We only consider edges as evidence for a synapse a if the predicted edge score $s(e = 1)$ is above a certain threshold t_1 . In order to account for neuron pairs that have multiple synapses at different locations, we introduce a separate candidate synapse $a_{l \rightarrow k}^m \subset a_{l \rightarrow k}$ for each set of edges whose target voxels form a connected component in segment s_k . A confidence score for candidate synapses $a_{l \rightarrow k}^m$ is obtained by calculating the sum over all edge scores per synapse:

$$\text{confidence}(a_{l \rightarrow k}^m) = \sum_{e \in a_{l \rightarrow k}^m} s(e) \quad (1)$$

A candidate synapse $a_{l \rightarrow k}^m$ is finally selected if $\text{confidence}(a_{l \rightarrow k}^m) > t_2$. The parameters t_1 and t_2 are used to control precision and recall of the method. Finally we extract single locations for the identified partners by calculating the center of mass of the start and end points of synapses $a_{l \rightarrow k}^m$. Synaptic partners are thus represented by a single pre- and a single postsynaptic location. In cases of reciprocal synapses, we remove the synapse with lower predicted evidence.

Error Metric. We evaluate our method using the synaptic partner identification metric proposed by the MICCAI 2016 CREMI challenge¹. First it uses Hungarian Matching to match each predicted synaptic pair $a_{l \rightarrow k}^m$ to a ground truth pair within a tolerance distance d . It further requires that the underlying segment IDs of ground-truth and predicted pair match. This means that a slight shift of the location is only tolerated if the underlying voxel still corresponds to the correct segment.

3 Results

CREMI Dataset. We tested our method on the publicly available serial section EM dataset from the MICCAI 2016 CREMI challenge¹. The dataset consists of six cutouts from an adult *Drosophila melanogaster* brain, divided into three training volumes (A, B and C) and three test volumes (A+, B+ and C+), with a size of $1250 \times 1250 \times 125$ voxels each and a highly anisotropic resolution of $4 \times 4 \times 40$ nm. We split the three training volumes into a training set (A, B and half C) and a validation set (the other half of C). The validation set is used for model optimization and qualitative analysis of the results. The ground truth for the training set includes volumetric neuron labels and synaptic partner annotations. Each partner is represented with a pre- and postsynaptic location (as shown in Fig. 1a). We use the training volumes to carry out a grid search in order to obtain parameters for our synaptic representation (discussed in 2.1).

¹ <https://cremi.org>.

We find $r_{\text{syn}} = 100$ nm and 14 edges $e_{i \rightarrow j}$ per voxel and the following offset vectors: $\mathbf{r}_1, \dots, \mathbf{r}_{14} = (0, 0, \pm 80), (\pm 120, 0, 0), (0, \pm 120, 0), (\pm 40, \pm 60, \pm 40)$ (shown in Fig. 2a).

Architecture. Our architecture is based on the U-Net architecture proposed by [13]. It has been shown to be well suited to a number of applications in biomedical imaging and neuron segmentation in particular (e.g. [3]). We use the same U-Net architecture as in [3] and we do not carry out any further experiments with regards to the hyperparameters of the network. We use three down- and three up-sampling layers where each layer consists of two convolutions with kernel size $(3 \times 3 \times 3)$. We downsample via max-pooling operations of size $(3 \times 3 \times 1)$ and upsample via transposed convolutions of the same size. Lastly, we apply a convolutional layer with kernel size 1 and a sigmoid activation function. Here, each output voxel has a context of $x = 212, y = 212, z = 28$ voxel (848, 848, 1120 nm). We use the Adam optimizer [8] with a learning rate of $\alpha = 5 \cdot 10^{-6}$ for all our experiments.

CREMI: Challenge Results. We include in our training the whole samples A, B, and C in order to make use of all available training data. We carry out a grid search on the same training data to find $t_1 = 0.5$ and $t_2 = 2500$ as the parameters maximizing the f-score. For the synaptic partner extraction we use an automatically generated neuron segmentation (provided by the authors of [3]), in which we manually mask missing sections and sections with artifacts to exclude them from the synaptic partner extraction. Results for the CREMI challenge are summarized in Table 1. We obtain a 3% gain in the f-score over the previous challenge leader. Looking at the total number of FN and FP this translates to a clear improvement in accuracy of $\sim 19\%$. Note that the f-score is averaged over all three samples equally, although sample A+, B+ contain presumably fewer synapses (the actual TP numbers are not revealed in the challenge). This explains the discrepancy of the $\sim 3\%$ gain in accuracy for the f-score and the 19% gain in accuracy for the absolute numbers.

Table 1. Results on the CREMI challenge for our method (lr U-Net) compared to the next best method (PSP unar). Highlighted in bold are our method and best values per column. Reported measures are (1) f-score, (2) FP and FN averaged over the three test datasets and (3) FN and FP summed over the three test datasets. For a detailed description of the used error metric, see Section 2.3

Method	f-score	avg. FP/FN	abs. FP/FN
PSP unar	0.539	266/ 281	800/ 843
lr U-Net	0.553	175 /314	526 /942
lr U-Net training	0.669	223/185	668/554

Qualitative Results. Since the ground truth is not public we cannot qualitatively analyze our results on the test set. Instead we summarize our observations for the validation set (f-score = 0.75). We find some errors in the ground truth

annotations, where some synapses have not been annotated (FP) or at least are of ambiguous nature (see Fig. 1b and c). Furthermore we observe that our method has problems with very small postsynaptic neurons or synapses that lie in the cutting plane (not shown). Figure 1(d-f) show a visualization of our results in form of a connectivity matrix. Note that the overall appearance of the predicted matrix and the ground truth is surprisingly similar despite the comparably low precision and recall (0.78, 0.73).

4 Discussion

The main advantage of our method is the ability to directly predict synaptic partners. Using a state-of-the-art deep learning architecture, we currently lead the CREMI challenge. Nevertheless, the comparably low overall performance (f-score: 0.55) demonstrates how challenging the task remains. Although we can confidently attribute some “errors” to incorrect labels or ambiguous cases, there are many observed cases that are truly incorrect.

Our current synapse representation has the shortcoming of using a discretized approach: sparse, long range edges. Although we guarantee that all synapses in training and validation set are captured, synapse statistics of other parts of the dataset potentially vary. Our method might thus benefit from a continuous formulation, in which we predict x, y, z values for the direction of the postsynaptic partner. We thus propose to rephrase the classification task as a regression task in future work.

Acknowledgements. This work was funded by the SNF grants P2EZP2_165241 and 205321L_160133.

References

1. Becker, C., Ali, K., Knott, G., Fua, P.: Learning context cues for synapse segmentation. *IEEE Trans. Med. Imaging* **32**(10), 1864–1877 (2013)
2. Dorkenwald, S., et al.: Automated synaptic connectivity inference for volume electron microscopy. *Nat. Methods* **14**(4), 435 (2017)
3. Funke, J., et al.: Large scale image segmentation with structured loss based deep learning for connectome reconstruction. *IEEE Trans. Pattern Anal. Mach. Intell.* (early access). <https://doi.org/10.1109/TPAMI.2018.2835450>
4. Huang, G.B., Scheffer, L.K., Plaza, S.M.: Fully-automatic synapse prediction and validation on a large data set. *arXiv preprint arXiv:1604.03075* (2016)
5. 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)
6. Januszewski, M., Maitin-Shepard, J., Li, P., Kornfeld, J., Denk, W., Jain, V.: Flood-filling networks. *arXiv preprint arXiv:1611.00421* (2016)
7. Kasthuri, N., et al.: Saturated reconstruction of a volume of neocortex. *Cell* **162**(3), 648–661 (2015)
8. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *CoRR abs/1412.6980* (2014). <http://arxiv.org/abs/1412.6980>

9. 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
10. 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)
11. Kreshuk, A.: Automated detection and segmentation of synaptic contacts in nearly isotropic serial electron microscopy images. *PloS One* **6**(10), e24899 (2011)
12. Roncal, W.G., et al.: Vesicle: volumetric evaluation of synaptic interfaces using computer vision at large scale. arXiv preprint [arXiv:1403.3724](https://arxiv.org/abs/1403.3724) (2014)
13. 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
14. Staffler, B., Berning, M., Boergens, K.M., Gour, A., van der Smagt, P., Helmstaedter, M.: Synem, automated synapse detection for connectomics. *eLife* **6**, e26414 (2017)