

RhoanaNet Pipeline: Dense Automatic Neural Annotation

Seymour Knowles-Barley ^{*†} Verena Kaynig [‡] Thouis Ray Jones ^{*‡}
 Alyssa Wilson ^{*} Joshua Morgan ^{*§} Dongil Lee ^{*} Daniel Berger ^{*}
 Narayanan Kasthuri ^{*¶} Jeff W. Lichtman ^{*} Hanspeter Pfister [‡]

November 22, 2016

Abstract

Reconstructing a synaptic wiring diagram, or connectome, from electron microscopy (EM) images of brain tissue currently requires many hours of manual annotation or proofreading (Kasthuri and Lichtman, 2010; Lichtman and Sanes, 2008; Seung, 2009). The desire to reconstruct ever larger and more complex networks has pushed the collection of ever larger EM datasets. A cubic millimeter of raw imaging data would take up 1 PB of storage and present an annotation project that would be impractical without relying heavily on automatic segmentation methods. The RhoanaNet image processing pipeline was developed to automatically segment large volumes of EM data and ease the burden of manual proofreading and annotation. Based on (Kaynig et al., 2015), we updated every stage of the software pipeline to provide better throughput performance and higher quality segmentation results. We used state of the art deep learning techniques to generate improved membrane probability maps, and Gala (Nunez-Iglesias et al., 2014) was used to agglomerate 2D segments into 3D objects.

We applied the RhoanaNet pipeline to four densely annotated EM datasets, two from mouse cortex, one from cerebellum and one from mouse lateral geniculate nucleus (LGN). All training and test data is made available for benchmark comparisons. The best segmentation results obtained gave $V_{F\text{-score}}^{\text{Info}}$ scores of 0.9054 and 0.9182 for the cortex datasets, 0.9438 for LGN, and 0.9150 for Cerebellum.

The RhoanaNet pipeline is open source software. All source code, training data, test data, and annotations for all four benchmark datasets are available at www.rhoana.org

^{*}Department of Molecular and Cellular Biology and Center for Brain Science, Harvard University.

[†]Present Address: Google Inc.

[‡]School of Engineering and Applied Sciences, Harvard University.

[§]Present Address: Washington University in St. Louis

[¶]Present Address: Argonne National Laboratory.

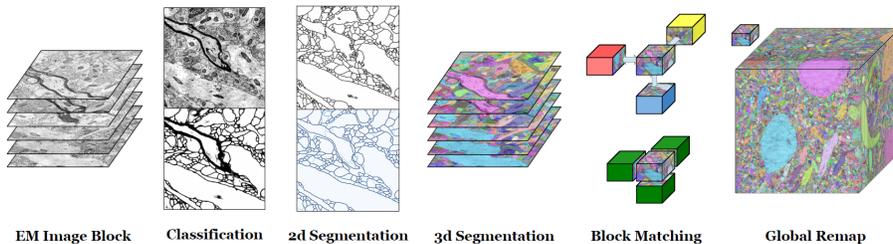


Figure 1: The RhoanaNet segmentation pipeline. First we generate boundary probability maps for each electron microscopy image. These are then used to obtain a 2D segmentation of neuronal regions, followed by region agglomeration across sections to obtain a 3D segmentation. All steps are embarrassingly parallel and can be executed in blocks of image volumes. The 3D segmentations from different blocks are then merged into one global reconstruction volume.

1 Introduction

Reconstructing the wiring diagram of the nervous system at the level of single cell interactions is necessary for discovering the underlying architecture of the brain and investigating the physical underpinning of cognition, intelligence, and consciousness (Kasthuri and Lichtman, 2010; Lichtman and Sanes, 2008; Seung, 2009). Advances in the data acquisition process now make it possible to image millions of cubic microns of tissue with hundreds of TB of raw data (Narayanan Kasthuri et al., 2015). With these techniques, a cubic millimeter of raw imaging data would take up 1 PB of storage and present an annotation project that would be impractical without relying heavily on automatic segmentation methods (Kaynig et al., 2015). Currently the general development of new image segmentation and reconstruction techniques for Connectomics is hindered by the lack of available benchmark data sets. These benchmarks are standard practice in computer vision and machine learning research and driving the development of new methods (Krizhevsky, 2009; Martin et al., 2001). The benefit of benchmarks for Connectomics became evident when the ISBI 2012 competition with one publicly released data set greatly improved the state of the art by opening the field to broader range of computer vision research (Arganda-Carreras et al., 2015).

In this paper we present a complete framework for benchmarking Connectomics reconstructions and apply it to four data sets with annotations (see figure 2), together with an improved version of our automatic segmentation pipeline called RhoanaNet (see figure 1), evaluation metrics and our current benchmark results.

The RhoanaNet pipeline and segmentation proofreading tools Mojo and Dojo (Haehn et al., 2014) are open source software. Source code and data are available online at www.rhoana.org.

2 Methods

We first describe the four benchmark data sets and then our improved automated pipeline and the evaluation metrics.

RhoanaNet Pipeline

Since the original publication of the Rhoana pipeline (Kaynig et al., 2015), we have updated every stage to take advantage of state-of-the-art deep learning techniques, resulting in large improvements in quality and performance. The new RhoanaNet Pipeline consists of five main stages. An overview is given in figure 1. The first stage of the pipeline is membrane classification, where cell membranes are identified in the 2D images producing a membrane probability map. Next, 2D candidate segmentations are generated based on the membrane probability for each image. These segmentations are then grouped across sections into geometrically consistent 3D neuron reconstructions. For this stage, 3D blocks are cropped from the full volume and each block is processed individually to produce a 3D segmentation. An optional cleanup stage removes very small objects and objects completely surrounded by a single object. In the fourth stage, blocks are matched pairwise with neighboring blocks, and overlapping objects are joined. Finally, matched objects are joined globally to produce a consistent segmentation for the full volume.

The modular pipeline approach allows each step to be improved or replaced independently of the rest of the pipeline. This is particularly useful for a direct comparison between methods which only address a part of the pipeline, e.g. updating the membrane classifier or the region agglomeration stage.

In comparison to the original Rhoana pipeline (Kaynig et al., 2015), the membrane classification stage has been updated to use state-of-the-art deep learning techniques, the region agglomeration has been changed from segmentation fusion (Vazquez-Reina et al., 2011) to Gala (Nunez-Iglesias et al., 2014) and the 2D segmentation and pairwise matching stages were updated to be more efficient. In addition, the pipeline was transitioned from MATLAB to a C++ and Python code-base which resulted in run-time performance improvements. The code for our pipeline can be found at www.rhoana.org.

Deep Learning

We used deep learning techniques to generate improved membrane probability maps and 2D segmentations. The Keras deep learning library, (Chollet, 2015) with the Theano back-end (Theano Development Team, 2016), was used to train a U-Net network (Ronneberger et al., 2015) to predict the probability of a pixel in the electron microscopy image representing a cell boundary membrane. The U-Net architecture follows the defaults described by Ronneberger et al. It consists of layer blocks containing three convolutional layers plus either a max-pooling layer for down-sampling or a deconvolutional layer for up-sampling. The

network first has four blocks downsampling the input and then four blocks up-sampling again to the original resolution. Residual connections are built between downsampling and upsampling blocks at matching resolutions. Ronneberger et al. showed that this network architecture is very suitable to membrane detection in electron microscopy images, while still maintaining a good throughput rate. On a GTX 970 GPU the data throughput rate is 1 megapixel per second; about two orders of magnitude faster than the random forest classification on a CPU. The U-Net architecture is sensitive to contrast variations in the images. Therefore, all network training and test data were pre-processed using CLAHE.

2D Segmentation

With improved classification performance from the deep networks the 2D segmentation stage no longer requires a graph-cut gap completion step (Kaynig et al., 2015). Instead, a simple watershed on Gaussian smoothed membrane probability maps is used to generate an over-segmentation, which then serves as input for the region agglomeration stage.

Region Agglomeration

The over-segmentation obtained from the previous step needs to be grouped into geometrically consistent 3D objects of neuronal structures. Gala (Nunez-Iglesias et al., 2014) uses a random forest classifier to predict the probability of two segments belonging to the same object. These scores are then used in an agglomerative clustering scheme to group the segmented regions. The random forest region agglomeration classifier is not only trained to group the initial segments, but also for later stages during the agglomeration phase. Iglesias et al. showed that this form of training significantly improves the clustering result. Compared to the previously employed segmentation fusion, Gala performs with greater accuracy for areas with branching structures such as dendritic spines in the cortex. For volumes with fewer branches, segmentation fusion and Gala perform with about the same accuracy.

Pairwise Block Matching

Pairwise matching was previously solved using Binary Integer Linear Program optimization (Vazquez-Reina et al., 2011). We have replaced this step with a much faster algorithm based on the stable matching algorithm (Gale and Shapley, 1962). Block matching is performed using multiple image planes from the overlapping volume between neighboring blocks. Objects are first matched by the stable matching algorithm for objects with overlaps above a given threshold (usually set to approximately $100nm^2$ per overlapping slice). After the first round of stable matching, any remaining unmatched objects are optionally matched to their largest partners.

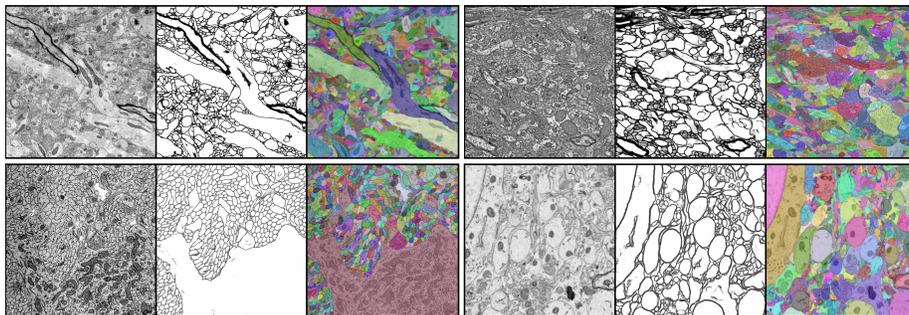


Figure 2: Image summary of the four datasets. Clockwise from top left with image heights: S1 Cortex (12.29 microns), LGN (16.38 microns), Cerebellum (8.19 microns), ECS Cortex (8.19 microns). Each panel consists of a source EM image (left), U-Net membrane probability results (middle) and final segmentation results (right). All images are from test volume data.

Datasets

We provide four benchmark datasets, each of which presents different challenges for segmentation.

S1 Cortex

This dataset was collected from mouse somatosensory cortex (Narayanan Kasthuri et al., 2015) and consists of several ground truth annotated volumes. The data is publicly available on the Open Connectome website (www.openconnectomeproject.org) and has been partly analyzed in an open benchmark competition (brainiac2.mit.edu/SNEMI3D/home). The data consists of 3 fully annotated sub-volumes from the whole S1 data set described in Kasthuri’s paper. This data has a resolution of 6 nm per pixel and a section thickness of about 30 nm. AC4 is a fully annotated volume of size $1024 \times 1024 \times 100$ pixels and corresponds to the training data of the benchmark competition from 2013. In addition we provide a similar sized volume AC3 consisting of $1024 \times 1024 \times 300$ pixels and a large cylindrical volume of a reconstructed dendrite and surrounding structures. The cylindrical volume fits inside a $2048 \times 2048 \times 300$ pixel crop from the S1 dataset. We used half of AC3 and all of AC4 as training data for both membrane classification and agglomeration stages of the pipeline, and the large dendrite volume as test data. All volumes are densely packed with dendrites and axons with many branching structures and small processes such as spine necks which make automatic reconstruction particularly challenging. This image data contains very few artifacts.

Cerebellum

This dataset was collected from developing mouse cerebellum, and contains a mix of parallel fibers and Purkinje cell processes. The parallel fibers generally

travel in the same direction without branches, and therefore present a relatively easy task for automatic reconstruction. However, the developing Purkinje cells consist of many branching processes and contain sub-cellular structures that are difficult to differentiate from the outer membrane of cells. The resolution of this data is about 8 nm per pixel with a section thickness of 30 nm. The data set consists of one fully annotated volume of $1024 \times 1024 \times 100$ pixels. We used the first 50 sections for training and the last 50 sections to evaluate test performance.

LGN

This training and test volume is a small part of a 67 million cubic micrometer volume from mouse lateral geniculate nucleus (LGN) (Morgan et al., 2016). Cell membranes appear densely packed and contrast is lower than in the other datasets. The data also contains some challenging staining artifacts, which are common in electron microscopy data and need to be addressed by automated reconstruction methods. The resolution of this data set is $4nm \times 4nm \times 30nm$. The training volume is $2360 \times 2116 \times 151$ pixels. The test volume is $2360 \times 2116 \times 20$ pixels.

ECS Cortex

This dataset was collected from mouse cortex, with the tissue processed to preserve the extra-cellular space (ECS) normally present in brain tissue. This data looks very different to conventionally stained EM data, as neuronal regions are not densely packed. It is the smallest data set in our collection, but the unique staining and resulting tissue properties make it a very interesting data set to analyze. The resolution of this data is 4 nm per pixel with a section thickness of 30 nm. The training volume is $1536 \times 1536 \times 98$ pixels. A small subsection of this volume consisting of $632 \times 560 \times 40$ pixels has been annotated again by a second person. To train our membrane detection network we used the training images and annotations from both label stacks. The test volume consists of $1536 \times 1536 \times 20$ pixels.

Metrics

It has been difficult to find consensus on the correct metric to use to measure segmentation accuracy in Connectomics. Easy to understand metrics such as “error free path length”, “edit distance” and “number of split and merge errors per cubic micron” are favored by biologists because of the intuitive nature of these metrics. A short description of these metrics provides a good idea of what the resulting numbers mean and how close the segmentation comes to a correct result. Unfortunately these metrics are not robustly defined and require arbitrary decisions to be made about what constitutes an error.

For example, small disagreements between segmentations do not change the overall structure of the 3D reconstruction. Two annotations of the same volume

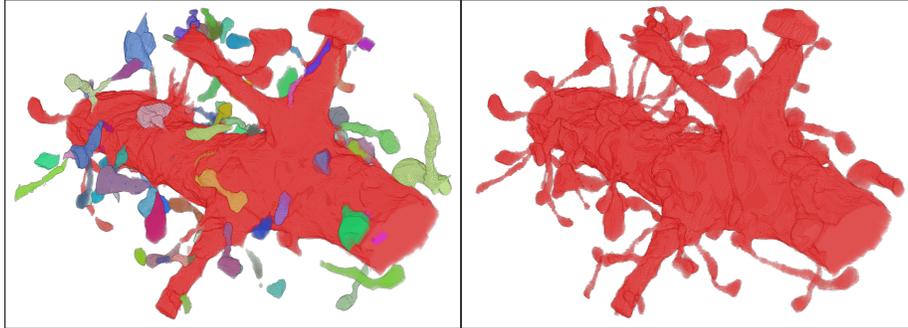


Figure 3: Automatic segmentation results (left) and manual annotations (right) for a $53.4\mu\text{m}^3$ section of dendrite (13.1 μm in length). 131 merge operations (non-red colors) and 23 split operations (not shown) would be required to proofread all errors larger than $0.0054\mu\text{m}^3$.

made by different experts (or made by the same expert at different times) will contain many differences in exact boundary locations, even when the overall structure of the reconstructed objects is the same. Therefore, metrics such as “errors per cubic micron” typically only count errors above a chosen minimum error volume. Similarly, the “error free path length” metric will typically ignore small merge or split errors lasting for a chosen minimum path length or number of annotation nodes. Ignoring these small errors is a reasonable approach to take and necessary when using this type of metric, however it makes it difficult to robustly define the metric and prevent exploitation of tuning parameters in the metric.

To avoid uncertainty introduced by arbitrary decisions, the computer vision community favors metrics that count all errors in a contingency table and assign different weights to each error depending on the size of the error, and complexity of the ground truth segmentation. Rand index, Adjusted Rand index and variation of information (VI) are examples of this type of metric. Unfortunately these metrics result in a number which is difficult to interpret intuitively and does not provide an obvious link to how close the segmentation is to the ground truth or how much time it would take to correct the segmentation manually or with a semi-automated proof reading software.

Arganda-Carreras et al. addressed this problem by defining two evaluation metrics which on the one hand provide the rigorous definition needed for benchmarking, and on the other hand provide some intuition about the ratio of split and merge errors in the segmentation. The two metrics defined in their paper (Arganda-Carreras et al., 2015) are based on the popular Rand and VI score and therefore named V^{Rand} and V^{Info} , with the term Info referring to the information theoretic background of the definition of variation of information. For both metrics Arganda-Carreras et al. separate these metrics into split- and merge-focused sub-metrics ($V_{\text{split}}^{\text{Rand}}$, $V_{\text{merge}}^{\text{Rand}}$ and $V_{\text{split}}^{\text{Info}}$, $V_{\text{merge}}^{\text{Info}}$) which have a

higher score if the segmentation contains fewer split (resp. merge) errors. As a summary metric they suggest the F-score, the harmonic mean between the split and the merge scores. Both metrics are normalized to a range between 0 and 1 with a higher value indicating a better segmentation. The evaluation done by Arganda-Carreras et al. shows that the ranking obtained with both metrics is not necessarily the same. While both metrics are normalized, they show different sensitivity to region sizes. The V^{Rand} emphasizes correct segmentation of large structures, while V^{Info} penalizes erroneous segmentation of smaller structures.

Another interesting point is the influence of the background label. If pixels labeled as background either in the ground truth or the automatic segmentation are excluded from the evaluation, automated segmentations with wide background margins tend to produce more favorable scores. To avoid tuning of this unwanted behavior, it is now standard practice to ignore the background labels from the ground truth segmentation, and grow all regions from the automated segmentation until no pixels are labeled background anymore. We follow this practice and report scores for V^{Rand} and V^{Info} for all data sets as obtained from the RhoanaNet Pipeline segmentation. Note that extra-cellular space is considered background for the ECS dataset.

Another challenge for the Connectomics community is the difficulty in comparing techniques across different species, sample preparation techniques, imaging modalities, and ground-truth annotation methods. It is impossible to meaningfully compare results using any metric if the source data is not the same, and given the complexity of the image processing pipelines it is difficult to identify where one method performs better than another. Fortunately, open datasets and challenges such as the ISBI 2012 and 2013 neuron segmentation challenges, which used the S1 data set, provide a starting point for direct comparisons between segmentation methods. The RhoanaNet Pipeline goes one step further, enabling direct comparison of whole pipelines as well as methods addressing specific parts of the pipeline.

3 Results

Here we demonstrate the RhoanaNet image processing pipeline on the four EM datasets discussed above. Deep neural networks were trained for each dataset individually, using the same network structure and hyper parameters each time. Agglomeration random forest training was performed on sub-volumes from the training data for each dataset individually, and cross-validation was used to choose the best random forest. A range of segmentation agglomeration levels were used to measure $V_{\text{F-score}}^{\text{Info}}$ on test data as shown in Figure 4. Full V^{Rand} and V^{Info} results for segmentations maximizing $V_{\text{F-score}}^{\text{Info}}$ are summarized in Table 1.

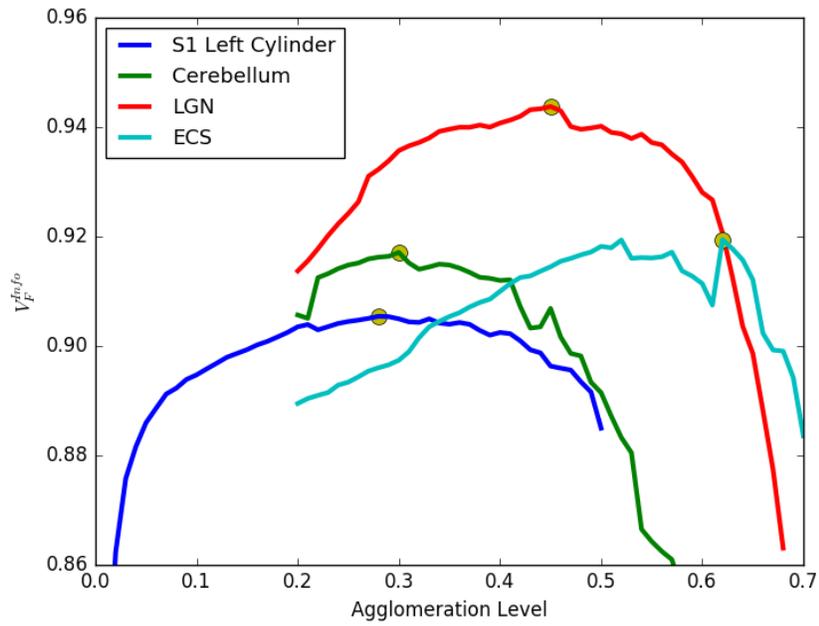


Figure 4: Segmentation metric $V_{F\text{-score}}^{\text{Info}}$ shown for a grid search over agglomeration levels. Maximal $V_{F\text{-score}}^{\text{Info}}$ results are highlighted with yellow dots with full results shown in Table 1.

S1 Cortex

This dataset contained very little noise, and Gala was able to connect most of the branching structures common in this data, including some spine necks, see Figure 3. A curiosity of this data set is that the training volume AC3 contains very little myelin. We therefore first trained on the whole training data (AC3 and AC4), and then fine tuned the network by restricting the training set to only images from AC4. We found that this training method lead to a slight increase in false positive detections on mitochondria, but a a good increase in true positive detection of myelinated cell boundaries. The decision to fine tune the network and all parameter tuning was performed on a validation set consisting of 10 images from AC3 and 10 images from AC4. No parameter tuning or additional training was performed based on the results on the test volume.

Cerebellum

The training set for this data is very small, but proved to be sufficient to train the U-Net to a satisfactory level. We follow the approach described by Ronneberger et al. (Ronneberger et al., 2015) and use random rotations, flips, and non-linear deformations for data augmentation.

LGN

Despite lower contrast and the presence of some artifacts, U-Net training provided good boundary predictions for this dataset and the best agglomeration achieved the highest $V_{F\text{-score}}^{\text{Info}}$ of all datasets at 0.9438.

ECS Cortex

This data set contains annotations from two different neurobiologists. Unfortunately the smaller set of annotations is restricted to a cropped volume of $632 \times 560 \times 40$ pixels. This size is smaller than the default configuration of the U-Net, which takes input patches of size 572×527 pixels. We therefore slightly reduced the size of the input patch for the U-Net to 540×540 pixels for this data set. This input size is still significantly larger than the context evaluated per pixel, thus it is unlikely to influence the accuracy of the network, but it slightly reduces the throughput performance during predictions.

4 Discussion

We presented the Rhoana pipeline for dense neuron annotation, updated to use deep learning U-Nets and region agglomeration using Gala. 3D segmentation results were qualitatively improved by the enhancements, with $V_{F\text{-score}}^{\text{Info}}$ scores ranging between 0.9 and 0.95 for the four datasets, and the pipeline can process very large volumes automatically.

	$V_{\text{split}}^{\text{Rand}}$	$V_{\text{merge}}^{\text{Rand}}$	$V_{\text{F-score}}^{\text{Rand}}$	$V_{\text{split}}^{\text{Info}}$	$V_{\text{merge}}^{\text{Info}}$	$V_{\text{F-score}}^{\text{Info}}$
S1 Cortex	0.7850	0.9216	0.8478	0.9276	0.8843	0.9054
Cerebellum	0.9583	0.9731	0.9656	0.9253	0.9049	0.9150
LGN	0.9162	0.6705	0.7744	0.9590	0.9290	0.9438
ECS Cortex	0.9589	0.6170	0.7509	0.9718	0.8702	0.9182

Table 1: Segmentation Results: Full segmentation metric results for each dataset. F-scores for Rand and Information Theoretic metrics are in bold.

High throughput of EM image data and quality of segmentation results will be very important for the future of Connectomics. Further improvements in segmentation quality and throughput are necessary and informed strategies for proofreading such as (Plaza et al., 2012; Haehn et al., 2014; Plaza, 2014) will be required to minimize human effort. This open source pipeline provides an improvement in both quality and throughput and the open benchmark datasets provide an open and reproducible reference example on which future improvements can be built.

Acknowledgements

Thanks to Linda Xu, students from Masconomet Regional High School and all our summer interns for their careful and diligent annotation work.

Also, thanks to Neal Donnelly, Princeton University, for contributions to the Gala project.

Funding: We gratefully acknowledged support from NSF grants IIS-1447344 and IIS-1607800 and the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBC) contract number D16PC00002, NIH/NINDS (1DP2OD006514-01, TR01 1R01NS076467-01, and 1U01NS090449-01), Conte (1P50MH094271-01), MURI Army Research Office (contract no. W911NF1210594 and IIS-1447786), CRCNS (1R01EB016411), NSF (OIA-1125087 and IIS-1110955), the Human Frontier Science Program (RGP0051/2014), Ruth L. Kirschstein Predoctoral Training Grant 5F31NS089223-02, and nVidia.

References

Arganda-Carreras, I., Turaga, S. C., Berger, D. R., Cirean, D., Giusti, A., Gambardella, L. M., Schmidhuber, J., Laptev, D., Dwivedi, S., Buhmann, J. M., Liu, T., Seyedhosseini, M., Tasdizen, T., Kamentsky, L., Burget, R., Uher, V., Tan, X., Sun, C., Pham, T. D., Bas, E., Uzunbas, M. G., Cardona, A., Schindelin, J., and Seung, H. S. (2015). Crowdsourcing the

- creation of image segmentation algorithms for connectomics. *Frontiers in Neuroanatomy*, 9:142.
- Chollet, F. (2015). Keras. <https://github.com/fchollet/keras>.
- Gale, D. and Shapley, L. S. (1962). College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):9–15.
- Haehn, D., Knowles-Barley, S., Roberts, M., Beyer, J., Kasthuri, N., Lichtman, J. W., and Pfister, H. (2014). Design and Evaluation of Interactive Proof-reading Tools for Connectomics. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):2466–2475.
- Kasthuri, N. and Lichtman, J. W. (2010). Neurocartography. *Neuropsychopharmacology*, 35(1):342–343.
- Kaynig, V., Vazquez-Reina, A., Knowles-Barley, S., Roberts, M., Jones, T. R., Kasthuri, N., Miller, E., Lichtman, J., and Pfister, H. (2015). Large-scale automatic reconstruction of neuronal processes from electron microscopy images. *Medical Image Analysis*, 22(1):77–88.
- Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. *Tech Report*.
- Lichtman, J. W. and Sanes, J. R. (2008). Ome sweet ome: what can the genome tell us about the connectome? *Current opinion in neurobiology*, 18(3):346–353.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. (2001). A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int’l Conf. Computer Vision*, volume 2, pages 416–423.
- Morgan, J. L. L., Berger, D. R. R., Wetzell, A. W. W., and Lichtman, J. W. W. (2016). The Fuzzy Logic of Network Connectivity in Mouse Visual Thalamus. *Cell*, 165(1):192–206.
- Narayanan Kasthuri, A., Jeffrey Hayworth, K., Raimund Berger, D., Eldin Priebe, C., Pfister, H., William Lichtman, J., Kasthuri, N., Lee Schalek, R., Angel Conchello, J., Knowles-Barley, S., Lee, D., Vázquez Reina, A., Kaynig, V., Raymond Jones, T., Roberts, M., Lyskowski Morgan, J., Carlos Tapia, J., Sebastian Seung, H., Gray Roncal, W., Tzvi Vogelstein, J., Burns, R., and Lewis Sussman, D. (2015). Saturated Reconstruction of a Volume of Neocortex. *Cell*, 162:648–661.
- Nunez-Iglesias, J., Kennedy, R., Plaza, S. M., Chakraborty, A., and Katz, W. T. (2014). Graph-based active learning of agglomeration (GALA): a Python library to segment 2D and 3D neuroimages. *Frontiers in neuroinformatics*, 8.

- Plaza, S. M. (2014). Focused Proofreading: Efficiently Extracting Connectomes from Segmented EM Images.
- Plaza, S. M., Scheffer, L. K., and Saunders, M. (2012). Minimizing manual image segmentation turn-around time for neuronal reconstruction by embracing uncertainty. *PloS one*, 7(9).
- Ronneberger, O., Fischer, P., and Brox, T. (2015). *U-Net: Convolutional Networks for Biomedical Image Segmentation*, pages 234–241. Springer International Publishing, Cham.
- Seung, H. S. (2009). Reading the Book of Memory: Sparse Sampling versus Dense Mapping of Connectomes. *Neuron*, 62(1):17–29.
- Theano Development Team (2016). Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688.
- Vazquez-Reina, A., Huang, D., Gelbart, M., Lichtman, J., Miller, E., and Pfister, H. (2011). Segmentation fusion for connectomics. Barcelona, Spain. IEEE.